Grounds and Truthmakers

Habilitationsschrift

zur Erlangung der *venia legendi* im Fach Philosophie

der Fakultät für Geisteswissenschaften

der Universität Hamburg

vorgelegt von

Dr. Stephan Krämer

Hamburg, 2023

Datum der mündlichen Prüfung:	22.11.2023
Gutachter:	Prof. Fabrice Correia
	Prof. Dr. Thomas Krödel

LIST OF ESSAYS

INTRODUCTION

What is the structure of reality? And how does reality connect to our representations of it, making them true, or false, according as they accurately or inaccurately describe reality? In thinking about these questions and articulating candidate answers to them, we naturally find ourselves invoking two closely related concepts: the concept of *ground*, and the concept of *truthmaking*. Ground connects an element of reality to the more fundamental ones that give rise to it, in virtue of which it obtains, thereby imposing an hierarchical structure on reality as whole. Truthmaking connects an element of reality to those elements of thought or language that adequately represent it. The two concepts are connected to each other, moreover, in that truthmaking may be explained in terms of a special case of ground: for an element of reality to make true a representation is for the obtaining of the former to ground the truth of the latter.

This pair of related concepts, of ground and of truthmaking, provides the common theme that unifies the present collection of essays. The collection is divided into three parts. Part I—*Generality, Ground, and Paradox*—consists of three essays dealing with topics that connect my (earlier) research on higher-order quantification to my work on the logic of ground. The bridge between these areas is provided by a family of paradoxes, which may be seen as deriving from circular relationships of aboutness. Part II—*The Logic of Ground*—collects my contributions so far to the study of the specifically logical features of ground. Part III—*Truthmakers, Ground, and Relevance*—finally comprises five essays resulting from my present, more general project of using truthmaker semantics to study various forms of relevance and their applications.

The aim of this introduction is, on the one hand, to situate the essays within the context of the wider debates they engage with, and on the other hand, to bring out recurring themes and ideas, as well as connections between them. The first two sections provide a brief general introduction to contemporary thought about the concepts of ground and truthmaking as they appear in my research. I then devote one section each to the three parts of the collection, in which I first give a more detailed overview of the state of the art concerning the topics that my essays address, and then explain the specific contributions they make to our understanding of these topics.

1. Grounds

To a rough approximation, ground here—and in the contemporary debate I am contributing to[1]—is the notion of something being the case *because*, or *in virtue of* some other thing or things being the case. More specifically, ground is supposed to be a distinctively *metaphysical* determination relation, partly analogous to but distinct from the relation of *causal* determination. Like causation, grounding is intimately connected to *explanation*. While the exact nature of the connection is controversial, it is generally accepted that in many cases where some facts ground another, we can explain why the latter fact obtains by citing the former facts—just as in many cases where some events cause another, we can explain why the latter event occurred by citing the former events.

Here are some examples of claims that are naturally construed as asserting relationships of ground in the intended sense:

(1) The shirt is red in virtue of being scarlet.
(2) Her smile exists because she is smiling.
(3) Their actions were wrong because they needlessly inflicted harm.
(4) That he is male, that he never married, and that he is of marriageable age together make him a bachelor.
(5) The singleton set of Socrates exists because Socrates exists.
(6) There are tables in virtue of there being simples arranged table-wise.
(7) The ball is red and round in part because it is round.
(8) The ball is red or round because it is red.
(9) She is in pain because her c-fibres are firing.
(10) It is true that snow is white because snow is white.

(Some, perhaps all of these, are controversial. The point is not that these claims are all obviously true, but that the questions they concern under their most natural interpretation are questions of ground.)

The examples may be used to illustrate some important basic features of ground. Firstly, there is a distinction between *full* and *partial* ground. In (7), the fact that the ball is round is not, on its own, sufficient to make it the case that the ball is red and round; it is only a partial ground of the latter fact. In contrast, in (1), the fact that the shirt is scarlet suffices, all by itself, for making it the case that the shirt is red; it is therefore a full ground of the latter fact. Note that while (7) explicitly indicates, by the use of 'in part', that it merely expresses a relationship of partial ground, it is plausible that the use

[1] The contemporary debate about grounding began with Kit Fine's 2001 paper 'The Question of Realism' and intensified significantly around 2010; see e.g. Schaffer (2009); Rosen (2010). However, although the contemporary debate has started relatively recently, as a number of authors have highlighted, philosophical interest in and use of the notion of ground, if not always under this label, has a long and distinguished history; see e.g. (Correia and Schnieder, 2012: §2) as well as the contributions on the history of grounding in Raven (2020).

of ‘because’, even without this qualifier, is compatible with the grounding in question being merely partial. In some of the other examples, it may be controversial whether the claims are plausible when understood as asserting full ground, or only when understood as asserting partial or full ground.

Secondly, grounding connects one fact—the fact that is grounded—to one *or more* other facts[2]—the facts jointly grounding the former fact. For instance, in (4), that he is a bachelor is said to be (fully) grounded by several facts taken together: the fact that he is male, the fact that he is unmarried, and the fact that he is of marriageable age. A fact being grounded by several facts together here must be sharply distinguished from it being grounded by their conjunction. For the conjunctive fact that the ball is red and the ball is round is plausibly taken to be jointly grounded by the fact that the ball is red and the fact that ball is round. But it is not plausibly taken to be grounded by their conjunction, i.e. by itself.

So far, I have spoken of ground as a relation between facts. It is controversial, though, if this is strictly accurate. For one thing, it is controversial if ground is, strictly speaking, a relation. Relations are the kinds of things expressed by polyadic predicates, like ‘is taller than’ or ‘is located next to’, forming a sentence when combined with two or more singular terms. The verb ‘ground(s)’ belongs to that grammatical category; its use is thus suggestive of a view of ground as a relation. On the other hand, grounding is often conveyed by means of ‘because’, which is a *connective*, similar to ‘and’ and ‘or’, forming sentences when applied to sentences. Its use would thus seem to no more commit one to seeing ground as a relation between facts than the use of ‘and’ commits one to a relation of conjunction between states of affairs or some such. In regimenting claims of ground, we thus face a choice: to what grammatical category should our basic expressions for ground belong? If we choose a predicate, we treat ground as strictly a relation, if we choose a connective, we don’t. Although the issue won’t matter too much for most of our purposes below, I side with those who prefer the second option. Following Fine, I use ‘$<$’ for full ground—thus regimenting e.g. (4) by ‘He is male, he is unmarried, he is of marriageable age $<$ he is a bachelor’—and ‘$\prec$’ for partial ground—thus regimenting e.g. (7) by ‘the ball is round $\prec$ the ball is red and the ball is round’.[3] Informally, for ease of exposition, I shall continue to speak of ground as a relation.

It is common, moreover, to distinguish between *worldly* conceptions of ground, and conceptual, or *representational* conceptions of ground. The distinction is straightforward to explain if ground is assumed to be relation. For we may then ask what kinds of

---

2 Indeed, it has been argued that some facts are grounded by no facts—or perhaps by the empty collection of facts—where this is different from a fact not being grounded at all. We can set this issue of possible so-called zero-grounding aside here; for discussion, see (Fine, 2012c: p. 47f).

3 (Fine, 2012c: §1.4); for discussion of the issue, see also (Correia and Schnieder, 2012: §3.1).

things it relates. One view is that grounds relates facts, conceived as worldly items, constituents of reality, rather than representations thereof. A different view is that grounding relates truths, i.e. accurate representations of reality. The first view embraces a worldly conception of ground, the second adopts a representational conception.[4] Now while this way of explicating the distinction depends on taking ground to be a relation, it is generally accepted that some form of the distinction also applies independently of that assumption. Exactly how it should be understood is neither completely clear nor completely uncontroversial, but the following explanation from Fine, although somewhat rough and preliminary, seems to be widely accepted:[5]

> A statement represents the world as being a certain way. We may therefore distinguish between the way it represents the world as being and how it represents the world as being that way. The worldly content of the statement is just a matter of the way it represents the world, while the conceptual content is also a matter of how it represents that content. The worldly conception of ground is one that is blind to anything other than factual content, while the conceptual conception of ground is one that also takes into account the representation of the factual content. Thus the worldly conception will presumably not distinguish between P and $P \wedge P$, since these two statements represent the world as being the same and so, just as it would be incorrect to say that P grounds P, it would be incorrect, under the worldly conception of ground, to say P grounds $P \wedge P$. On the conceptual conception of ground, by contrast, it will be perfectly acceptable to say that P grounds $P \wedge P$ since the one representation will be true in virtue of the other. (Fine, 2017b: pp. 685f)

Given the distinction between worldly and representational conceptions of ground, it is a further question whether they should be seen as rivals. My own view is that they should not, and that they simply describe different relations of ground, each of them interesting and important in its own right. Some of my essays below contribute to the theory of a worldly ground, some to the theory of representational ground.

Whether conceived in worldly or representational terms, ground is usually thought of as *factive* in the sense that if $P \prec Q$, then both P and Q are the case. Thus, since snow is white, it is clear that snow's being white or not white is *not* grounded in snow's being not white—though it would be so grounded, if snow were not white. It may be argued, however, that there is also a *non-factive* understanding of ground, on which neither ground nor groundee are required to obtain or be true. On this understanding, that

4 Correia (2010: pp. 256f) was the first to explicitly make this distinction.

5 I say more about this issue in (Krämer, 2021a: pp. 1667f), included below.

snow is white or not is grounded in snow's being white, *and* in snow's not being white, because both states of affairs—obtaining or not—are appropriately related to snow's being white or not; in this case, as disjunct to disjunction. Somewhat metaphorically, we may think of the notion of non-factive ground as obtained from the factive one by subtracting the requirement of factivity for both relata.[6]

Most of my own work on ground is devoted to the project of delevoping a satisfactory *logic* of ground. Following Fine (2012c: p. 54), we may distinguish between the *pure* and *impure* logic of ground. The pure logic of ground consists of structural principles about the various notions of ground, in which no reference is made to the internal structure of the relata of ground. The structural principles of ground usually ascribed to ground may be regarded as flowing from two driving thoughts. The first thought, already mentioned above, is that ground imposes a hierarchical structure, and thereby a form of *order* on reality. This is reflected in the assumption that partial ground has the defining features of a so-called *strict partial order*, namely transitivity—if $P \prec Q$ and $Q \prec R$ then $P \prec R$—and asymmetry—if $P \prec Q$ then $Q \nprec P$.[7] The second thought is that ground implies a strong form of *relevance*: in order for some fact to help ground another, it needs to be wholly relevant to the latter fact's obtaining. This is reflected in the failure of a principle of weakening, or *monotonicity*: from the assumption that $\Gamma < P$, it does not follow that $\Gamma, Q < P$, for arbitrary Q.[8] These principles, and most other principles of the pure logic of ground, are equally plausible under a worldly and a representational conception of ground.

The impure logic of ground consists of principles that concern the interaction of notions of ground with logical concepts such as conjunction, disjunction, existential quantification, etc. Here, too, there is a general guiding idea that underlies the various principles being considered in the literature: that logically complex facts obtain in virtue of the logically simpler constituents from which they are constructed by means of some logical operation. For instance, it is prima facie plausible that any conjunctive fact $P \wedge Q$ is grounded by its conjuncts. As can be seen from the quoted passage from Fine, with regard to these impure logical principles, the distinction between worldly and representational grounding carries greater significance. For if we accept the conjunction principle just stated in full generality, we are committed to the claim that P grounds $P \wedge P$, which is implausible, according to Fine, under a worldly conception of ground.

Still, the impure logics of worldly and representational ground plausibly have many of their most distinctive features in common. In particular, under both conceptions,

6 On the idea of non-factive ground and its relation to factive ground, see esp. (Fine, 2012c: p. 48ff)).

7 *Full* ground plausibly satisfies these closely related principles: If $\Gamma < P$ and $\Delta, P < Q$ then $\Gamma, \Delta < Q$—a strong form of transitivity—and if $\Gamma, P < Q$ then $\Delta, Q \nless P$—a strong form of asymmetry.

8 The first, and most influential in-depth study of the pure logic of ground is Fine (2012d).

ground is *hyperintensional*: statements which are equivalent within classical logic may differ with respect to what they ground, and what they are grounded by. For a very simple example, consider the statement that snow is white or not white. Plausibly, it is grounded by snow's being white. But the fact that snow is white or not white is logically equivalent to the fact that grass is green or not green. This latter fact, however, is *not* plausibly grounded by the fact that snow is white—that fact is wholly irrelevant to grass being green or not. Ground's hyperintensionality is the source of some of the main technical challenges we face in developing an adequate formal theory of ground. As part of such a theory, we need an adequate formal *semantics* for ground. In particular, we need a theory of what we may call *ground-theoretic content*: those aspects of the meaning of sentences that are relevant to the truth-conditions of grounding statements. But the most familiar and well-understood formal conception of content—the view of content as given by a set of possible worlds—is blind to hyperintensional distinctions, and therefore inadequate as a conception of ground-theoretic content. We therefore need a replacement for possible world semantics, which enables us to distinguish between contents in a more fine-grained way, and thereby allows us to assign different contents to, for example, the statement that snow is white or not and the statement that grass is green or not. It is in this context that truthmaking enters the scene in my research.

2. Truthmakers

'The idea of truthmaking', as Fine (2017c: p. 556) puts it, 'is the idea of something on the side of the world—a fact, perhaps, or a state of affairs—verifying, or making true, something on the side of language or thought—a statement, perhaps, or a proposition.' In the modern debate, truthmaker theory has for the most part been pursued as a project in *metaphysics*. Specifically, many philosophers have advocated its use as a *guide* to metaphysics, the idea being that a good method for inquiry into the constituents of reality is to see what kinds of statements we consider true, and then to ask what sorts of things on the side of the world might make them so.[9] More recently, the idea of truthmaking has instead been used as a guide to *semantics* rather than metaphysics, the idea being that we gain an adequate understanding of language by asking how items in the world might make its sentences true.[10] It is in this latter role, as the central concept

9 The contemporary truthmaker debate in metaphysics may be regarded as beginning with the seminal Mulligan et al. (1984), but the philosopher whose name is connected most strongly with this programme of truthmaking-centered metaphysics is David Armstrong; see e.g. Armstrong (1997, 2004).

10 See especially Fine (2017a,b,c) for introduction and development of the general framework. Further publications applying the framework in various areas are mentioned below. Although it is only recently that truthmaker semantics has received much attention, a semantics of this sort was first developed already in van Fraassen (1969). Moreover, situation semantics in the tradition of Barwise and Perry (1983) may also be regarded as a form of truthmaker semantics, though based on an importantly less discriminating, monotonic notion of truthmaking.

of a framework for semantic theorizing, that the idea of truthmaking appears in much of my work.

The core elements of truthmaker semantics are best explained against the background of the more familiar framework of possible world semantics. In possible world semantics, the content of a declarative sentence is represented by the set of possible worlds at which the sentence is true. This account of content makes use of two concepts: that of a *possible world*, and that of (a sentence) being true at (a possible world). A lot can be said about what a possible world might be taken to be, but for its role in semantic theorizing, it is mainly two properties that are crucial: that it is *consistent* in the sense that at most one of any sentence and its negation is true at it, and that it is *complete* in the sense that at least one of any sentence and its negation is true at it. The notion of being true at can be explicated in modal terms. Roughly, a sentence is true at a possible world w iff necessarily, if things are as w has it, then the sentence is true. Within the framework of possible world semantics, the semantics of an expression is then given by its effect on the set of worlds at which the sentences in which it occurs are true. The semantics of conjunction and disjunction, for example, are given by the observations that a conjunction $P \wedge Q$ is true at a world w iff both P and Q are true at w, and a disjunction $P \vee Q$ is true at w iff at least one of P and Q is true at w.

In truthmaker semantics, the content of a declarative sentence is represented by the set of *states* that *make* the sentence true. So, first of all, the appeal to possible worlds is replaced by an appeal to a more general category of *states*. Like possible worlds, states are required to be specific, i.e. non-disjunctive: a state cannot make true a disjunction without making true at least one of its disjuncts. Unlike possible worlds, states need not be worlds, i.e. complete: states typically leave open the truth-value of many statements. Unlike possible worlds, states need not be possible: at least in many applications of truthmaker semantics, we also allow for inconsistent states, such as states making true a contradiction. And unlike possible worlds, states are (non-trivially) ordered by *part-whole*—one state may be properly contained in another—and two states may be *fused* to form a bigger state containing both of them as parts.

Second of all, the appeal to the purely modal notion of being true at is replaced by a much more fine-grained notion of truthmaking. That notion, as I have hinted above, is closely related to the notion of ground: roughly speaking, for a state to make true a statement is for the state's obtaining to ground the statement's being true, to *bring about* the truth of the statement. Note that truthmaking is like *non-factive* rather than factive ground: false, even contradictory statements may have truthmakers, and non-obtaining, perhaps even impossible states may be truthmakers of a given statement. For a statement to be true, it is not enough that it *have* a truthmaker, it must have an *obtaining* truthmaker. So by allowing for states as truthmakers of contradictions, we do not thereby automatically allow for true contradictions; we may still hold that no truthmaker of a contradiction could ever obtain.

The connection between ground and truthmaking is reflected in a number of shared logico-structural features. In particular, just like Γ must be wholly relevant to P for it to be the case that $\Gamma < P$, a state s must be *wholly relevant* to a given statement P for s to be a truthmaker of P. As a result, just like ground, truthmaking in the present sense is non-monotonic: combining a truthmaker of a statement with some other state to form a bigger state does not in general yield another truthmaker of the statement in question.

Within the framework of truthmaker semantics, the semantics of an expression e is then given by its effect on the set of states that make true the sentences in which e occurs. Roughly at least, the semantics of conjunction and disjunction, for example, may be given by the observations that a conjunction $P \wedge Q$ is made true by a state s iff s is the mereological fusion of a state making true P and a state making true Q, and that a disjunction $P \vee Q$ is made true by s iff at least one of P and Q is made true by s.[11] To account for negation, we may also invoke a notion of a state making *false* a given statement, and take a state to make true (false) $\neg P$ just in case it makes false (true) P.

Considering the truthmakers of statements rather than (just) the possible worlds at which they are true allows us to make much more fine-grained distinctions. Consider again the pair of logical truths mentioned earlier. That snow is white or not, by the clause for disjunction, is made true only by states that make it true that snow is white—such as the state of snow being white—and by states that make it true that snow is not white—such as perhaps the state of snow being green.[12] That grass is green or not in turn is made true only by states that make it true that grass is green—such as the state of grass being green—and by states that make it true that grass is not green—such as perhaps the state of grass being white. So even though both statements correspond to the same set of possible worlds, namely the set of all possible worlds, they correspond to distinct and indeed disjoint sets of truthmakers. Moreover, it is very natural to suspect that we can explain the semantics of ground in terms of truthmaking: in a first approximation, the idea would be that snow's being white grounds snow's being white or not because every truthmaker of the former statement also makes true the latter, while grass's being green does not ground snow's being white or not because the truthmakers of grass being green do not make it true that snow is white or not. This is roughly the idea underlying

11 I say roughly because, for certain purposes at least, one may wish to fine-tune these clauses somewhat. For instance, one may also wish to allow arbitrary fusions of truthmakers of P and Q as truthmakers of $P \vee Q$, and one may wish to hold that if a given state both has a part that is a truthmaker of P and is a part of a truthmaker of P, then the state itself is also a truthmaker of P.

12 If we fine-tune the clause for disjunction as indicated in the previous footnote, then this line of reasoning needs some fine-tuning as well, but the essential point remains unaffected by these adjustments.

the truthmaker semantics for ground proposed by Fine (cf. his 2012d; 2012c; 2017b), which we shall consider in more detail later on.

But the theory of ground is by no means the only application of the framework of truthmaker semantics. It has also been used to develop novel semantic accounts of counterfactual conditionals (Fine (2012a,b)), partial truth (Fine (msb)), verisimilitude (Fine (2021)), subject matter (Yablo (2015); Fine (msa)), relevant entailment (Fine (2016); Jago (2020)), statements of permission and obligation (Anglberger et al. (2016); Fine (2018); Yablo (2011); Rothschild and Yablo (202x)), to name but a few examples. In my own work (included below), I have applied the framework especially to the study of various forms of *relevance* relations, such as explanatory relevance (Krämer and Roski, 2017), evidential relevance (2017b), epistemic relevance in the context of belief revision (2022), and logical relevance in the context of a novel, relevance-sensitive understanding of what it is for a truth to be *the whole truth* (202xb; 2023).

3. Part I—Generality, Ground, and Paradox

One of the most important and fundamental facts about human thought and language is our ability to *generalize*, and to express and communicate generalizations using quantificational vocabulary like 'all', 'some', 'most', etc. The behaviour and properties of expressions of generality accordingly constitute central and much-discussed topics across a range of philosophical and neighbouring disciplines, including logic, semantics, and metaphysics.

Quantificational expressions of generality come in different grammatical varieties. The most common and familiar type of quantifiers—singular first-order quantifiers—are associated with the position of singular terms. This kind of quantification occurs, for example, in the following statement: Tom is a singer and Tom is a guitarist, therefore *someone* is both a singer and a guitarist—namely, *Tom*. But not every kind of quantification is of this sort. Thus, consider the following statement: Tom is famous and Amy is famous, therefore Tom is *something* that Amy is, too—namely *famous*. Here we are quantifying into the position of an *adjective* rather than a singular term. This is an instance of (singular) *second-order* quantification. This type of quantification is the topic of my book 'On What There Is For Things To Be' (2014), a revised version of my doctoral dissertation, in which I defend the view that second- and higher-order quantification is a legitimate, intelligible, and sui generis form of quantification which is not reducible to first-order quantification. Following on from my doctoral work, I have worked on the implications of this view for a number of broadly logical *paradoxes*, unified by the feature that they appear to essentially involve some form of *circular* relationships of *aboutness*.

The essays making up the first part of the present collection have resulted from this research. ‘Everything, and Then Some’ (2017a) concerns the threat of paradox for the possibility of absolute generality, i.e. quantification over absolutely everything. Certain ways of responding to this threat make essential use of irreducibly higher-order quantification; the paper shows that the prospects for such a response are less bright than one might have hoped. Circular aboutness also threatens natural principles linking *ground* to aboutness, given that grounding is widely held not to permit circles—the resulting puzzles of ground are the topic of my ‘A Simpler Puzzle of Ground’ (2013) and ‘Puzzles’ (2020). One of my contributions here is to show that the use of higher-order quantification allows us to formulate particularly minimalistic versions of the ground-theoretic paradoxes, which depend on very few assumptions.

3.1. **Paradoxes of Generality.** In most ordinary contexts, the generalizations we make are—implicitly or explicitly—*restricted*: they do not concern literally everything, but just some things that are of interest in the situation at hand. It would seem, however, that not all contexts are like that. When the atheist pronounces that there is no God, she does not merely wish to say of some restricted set of contextually relevant things that none of them is a God. In her context, everything is relevant, and she means to say of absolutely everything without exception that it is no God. Similar observations apply to many of the claims we make in logic, metaphysics, or mathematics. When the logician says that everything is self-identical, when the nominalist says that everything is concrete, and when the mathematician says that nothing is a member of the empty set, they do not intend to rule out merely that the members of some restricted domain of objects D are self-distinct, abstract, or a member of the empty set, while allowing for exceptions to their generalizations as long as they are excluded from D. Call the view that this sort of absolute, unrestricted generality is possible *generality absolutism*.

It turns out that generality absolutism is threatened by paradoxes akin to Russell’s paradox. Russell’s paradox (in one version) arises from the prima facie attractive assumption that every substitution-instance of the following schema is true:

(R) $\exists x \forall y$ (y exemplifies x iff Fy)

If we instantiate ‘Fy’ with ‘y does not exemplify y’, we obtain the statement that $\exists x \forall y$ (y exemplifies x iff y does not exemplify y). But this statement leads to a contradiction in classical logic. Suppose r is a verifying instance of the existential quantification, so $\forall y$ (y exemplifies r iff y does not exemplify y). Now given that we may instantiate the universal quantification with r, we obtain that r exemplifies r iff r does not exemplify r. This statement is equivalent to the contradiction that r both exemplifies and does not exemplify r.

More generally, we may note that every instance of the following schema is inconsistent within classical logic:

(C) $\exists x \forall y\, (Rxy \text{ iff } \neg Ryy)$

Since (R) has substitution-instances that instantiate (C), the validity of the schema (R) is inconsistent with classical logic.

Now how does the claim that absolutely general quantification is possible lead to paradox? The problem—in the version under discussion in my paper—arises when the quantifiers of a language L are assumed to generalize, among other things, over the totality of possible *interpretations* of that very language L—an assumption that seems mandatory if we wish to hold that L's quantifiers express *absolute* generality. For given that assumption, we can argue for the validity of a schema structurally analogous to (R). The difference is that while (R) implies the existence of properties with the specified exemplification conditions, this schema implies the existence of certain kinds of *interpretations* of L:

(I) $\exists i \forall y\, (P \text{ applies to } y \text{ under } i \text{ iff } Fy)$

If we instantiate 'Fy' with 'P does not apply to y under y', we obtain the statement that $\exists i \forall y\, (P$ applies to y under i iff P does not apply to y under $y)$: another instance of the inconsistent schema (C).

Perhaps the most common type of response to this difficulty is to postulate some form of *hierarchy* of languages, each capable only of generalizing over interpretations of languages at lower levels of the hierarchy, and never over all of its own interpretations. In place of (I), one may then use a version in which the universal quantifier is restricted to a range excluding the problematic 'Russell interpretation'. At first glance, however, this comes at the cost of denying generality absolutism: at no stage of the hierarchy can we talk about absolutely everything.

'Everything and Then Some' examines an idea, influentially proposed by Timothy Williamson (2003; 2013), for a version of the hierarchical solution that is compatible with generality absolutism. The idea is to make use of a hierarchy of quantifiers of ever higher *grammatical types* or *orders*. The first, i.e. lowest order of quantification is quantification into name position. Second-order quantification then is quantification into the position of predicates, i.e. sentence-forming operators on names, and more generally $(n+1)^{th}$ order quantification is quantification into the position of sentence-forming operators on the type of expression associated with n^{th}-order quantification. The point of using this grammatical hierarchy is that we may view the ranges of different orders of quantification as *incommensurable*, and thereby deny that the range of first-order quantifiers includes that of second-order quantifiers while resisting the conclusion that the range of first-order quantifiers is restricted.

My paper places Williamson's incommensurability thesis under scrutiny. It turns out that Williamson is pushed towards a view of the higher-order predicates and quantifiers as forming a *cumulative* type-hierarchy, where the higher types strictly subsume the lower types. I show that as a result of this, there is a significant sense in which the different ranges of quantifiers are commensurable after all. This result makes doubtful whether Williamson's view is still a form of generality absolutism worthy of the name, or whether it is rather a form of generality relativism in spirit. A number of subsequent papers have engaged with the ideas and arguments of the paper in some detail; see especially Florio and Jones (2021); Button and Trueman (2021); Florio and Jones (202x).

3.2. **Paradoxes of Ground.** As mentioned before, one may see the paradox of generality presented above as arising in part due to circular relationship of aboutness: roughly speaking, under the problematic Russell interpretation, the application conditions of the given predicate—which determine what the predicate is about—seem to be defined in part by reference to themselves. The fact that wide-ranging generalizations quickly lead to forms of self-reference or circular aboutness also lies at the core of a number of paradoxes in the theory of ground, first presented by Kit Fine in his paper 'Some Puzzles of Ground' (2010). These paradoxes shows that individually plausible principles about ground jointly lead to a contradiction. Some of the relevant principles are structural, or *pure* in the sense explained above, and reflect the picture of ground as imposing an hierarchical order on reality; specifically, the principles that ground is transitive and asymmetric, and hence irreflexive: nothing strictly grounds itself. Other relevant principles are *impure*, and concern in particular the interaction between ground and quantificational expressions.

In its original, Finean form, one of the relevant paradoxes can be summarized as follows. Firstly, it is plausible that true existential quantifications are grounded by their true instances (E). For example, that Fine is a philosopher grounds that someone is a philosopher. Secondly, it is plausible that if the proposition that p is true, then this is so because p, i.e. that p grounds that the proposition that p is true (T). Combining these principles with the assumption that grounding is transitive yields a violation of the principle that grounding is irreflexive. For clearly, some proposition is true. Hence, the proposition that (some proposition is true) is true, and by (T), this is so because some proposition is true. But the statement that the proposition that (some proposition is true) is true is then a true instance of the existential quantification that some proposition is true. So by (E) some proposition is true because the proposition that (some proposition is true) is true. By transitivity, some proposition is true because some proposition is true, in violation of irreflexivity.

In this version, the paradox employs several assumptions that, while plausible, are not uncontroversial, and hence seem to provide avenues for blocking the paradox. In particular, it assumes the existence of the relevant propositions, the principle (T), transitivity, as well as irreflexivity and the principle (E). The central contribution of 'A Simpler Puzzle of Ground' is to show that if we replace the assumption of a domain of propositions by the use of higher-order quantification into sentence position, we no longer need the assumption of transitivity or the principle (T): the (higher-order analogue of) (E) and irreflexivity are then in direct conflict. As a result of this, various attempted responses to Fine's puzzle do not apply to this version, and may therefore be argued to violate the principle that, roughly speaking, a uniform collection of paradoxes should be solved in a uniform way. Although my shortest paper by far, 'A Simpler Puzzle of Ground' has received a fair amount of attention; substantial discussions of its central points may be found especially in Woods (2018); Fritz (2020); Lovett (2020b); deRosset (2021); Goodman (2022).

The puzzles just described—Fine's, and my simplified variant—are not isolated problems, but belong to a larger family of related paradoxes of ground, all of which exploit circular forms of aboutness. They differ from one another in the means employed to achieve circular aboutness and in the specific impure principles of ground they assume. For instance, instead of existential quantification, we may build a suitable circle using either universal quantification or devices of singular self-reference, similarly as in the notorious Liar paradox: 'this statement is false'. The third and final article of part I is a survey article, written for the *Routledge Handbook of Metaphysical Grounding* (Raven, 2020), providing an overview of this entire family of ground-theoretic paradoxes. I first provide a systematic overview of the numerous versions of these puzzles, their interrelation, and the assumptions on which they depend. I then articulate a set of plausible desiderata by which attempted solutions to the puzzles may be evaluated. Finally, I briefly present the different solutions to the puzzles proposed in the literature and offer preliminary assessments of which desiderata they satisfy or fail to satisfy.

4. Part II—The Logic of Ground

Contemporary work on the logic of ground began around 2010, with the pioneering contributions Batchelor (2010), Correia (2010), Fine (2010, 2012c,d), Rosen (2010), and Schnieder (2011).[13] These approaches differ somewhat in their scope and target; I shall here focus on Correia's and Fine's work since it has been the most influential, and is the most important for my own work.

---

[13] The earliest systematic and formal study of the logical features of ground is due to Bernard Bolzano (1837). On Bolzano's groundbreaking contributions to this field, see esp. Roski (2017), Roski and Schnieder (2022).

Correia (2010) is the first detailed development of a formal logic of ground, explicitly targeting a *worldly* conception of ground. Aside from the familiar logical vocabulary, the language Correia uses has two specifically ground-theoretic expressions. One corresponds to our '$<$', it connects list of sentences with a single sentence and expresses the notion of full ground. The other is a two-place sentential connective '$\approx$' for factual equivalence; it is used to express that the facts described by the sentences on either side are identical, so that the sentences are freely interchangeable in the scope of the grounding operator.

The axioms and rules of Correia's logic fall into three groups. The first comprises structural principles for ground and for factual equivalence. The second consists of principles involving factual equivalence and the truth-functional connectives. It includes, for example, the principle that $P \approx P \wedge P$, in line with Fine's contention above that a worldly conception of ground should be blind to the difference between P and $P \wedge P$. In the third group we have principles governing the interaction of the grounding operator and the truth-functional connectives. The basic idea here is that logically complex facts built up using truth-functional operations like conjunction and disjunction are always grounded, according to systematic principles by logically simpler facts in which these operations do not occur. In a first approximation, we might state two such principles as follows:

($\wedge$) If P and Q, then $P, Q < P \wedge Q$
($\vee$) If P, then $P < P \vee Q$

Under the worldly conception of ground, these aren't quite right yet. For ($\wedge$) yields the result that if P, then $P, P < P \wedge P$, from which we may infer, given $P \approx P \wedge P$, that $P < P$: every fact grounds itself. Similarly, from ($\vee$) we obtain $P < P \vee P$ if P, from which we may again infer, given the plausible assumption that $P \approx P \vee P$, that $P < P$. For this reason, Correia's logic merely includes restricted versions of ($\wedge$) and ($\vee$), which respect the irreflexivity of ground.

Correia then goes on to provide an algebraic semantics for the logic and proves that his axioms and rules are sound and complete with respect to that semantics. While his result is very significant and constitutes an important milestone in the development of the logic of ground, the overall account is not completely satisfactory. One limitation of Correia's account is that his semantics is not very illuminating. In particular, it does not give us an independent grip on the question what logical principles ground should be taken to satisfy, since it simply postulates, principle by principle, that the facts used to interpret the formal language form a grounding hierarchy that obeys exactly the principles endorsed in the logic.

Strikingly, however, it turns out that the much more natural and illuminating truthmaker semantics for ground that was independently developed in Fine (2012d) matches

Correia's account exactly (cf. (Fine, 2017b: pp. 685ff)). Specifically, say that $\Gamma \leq Q$—read: Γ *weakly* fully grounds Q—iff every truthmaker of $\bigwedge \Gamma$ is a truthmaker of Q.[14] Then say that $\Gamma < Q$—Γ *strictly* fully grounds Q—iff $\Gamma \leq Q$, and for all P among Γ, there is no Δ such that $Q, \Delta \leq P$. Less formally: Γ strictly fully grounds Q if Γ weakly fully grounds Q, and Q does not help weakly ground any fact among Γ. If we now assume two closure principles about the set of truthmakers of any given statement P, this account of ground is in exact alignment with the logic proposed by Correia. The first principle is that the set of truthmakers is always *closed under fusion*: if some states make true P, then so does their fusion. The second assumption is that the set of truthmakers is *convex*: any state that is both part of some truthmaker of P and that has some truthmaker of P as part is itself a truthmaker of P.

A remarkable feature of this account of ground, as highlighted by Fine (cf. 2017b: p. 686), is that ground turns out to be definable in what seem to be purely *logical* terms, namely disjunction, conjunction, and factual equivalence: $\Gamma \leq Q$ holds just in case $\bigwedge \Gamma \vee Q \approx Q$. So ground, as seen through the lens of the truthmaker account, turns out to be a logical rather than a metaphysical notion.

The first pair of papers included in this part of the collection push for certain refinements of this account of worldly ground. 'A Note on the Logic of Worldly Ground' (Krämer and Roski, 2015), co-authored by Stefan Roski, argues against certain ground-theoretic principles corresponding to the principle of convexity for truthmaking. This paper also marks the beginning of my more general research project on notions of *relevance*—the focus of part III of this collection—by tracing the problems it identifies for Fine's original account to an inadequate view of the kind of relevance characteristic for grounding. My 'Singular Troubles with Singleton Socrates' (2021b) shows that owing to its narrowly logical conception of ground, the present approach is prevented from capturing certain important kinds of instances of worldly ground—such as the grounding of the existence of the singleton set of Socrates by the existence of Socrates. It then proposes a generalization of the account under which ground is no longer a purely logical notion, but extends to such distinctively metaphysical cases of ground.

The second pair of papers in this part contribute to the theory of the representational conception of ground. As I have mentioned, under a worldly conception of ground, the rules describing the interaction between ground and the truth-functional operators are somewhat complicated: the simple rules ($\wedge$) and ($\vee$) stated above need to be restricted in certain ways, on pain of generating violations of the asymmetry of ground. Under a representational conception of ground, however, it is natural to adopt a more fine-grained view of ground-theoretic content. We may then regard a truth P as ground-theoretically inequivalent to the corresponding self-disjunction and -conjunction $P \wedge P$

[14] $\bigwedge \Gamma$ is the conjunction of all the statements in Γ.

and $P \vee P$. Relatedly, we may endorse very simple, unrestricted rules describing the grounds of truth-functionally complex truths. In pursuing this approach, the question arises whether the truthmaker semantics for ground can somehow be modified so as to yield a suitably more fine-grained conception of content. 'Towards a Theory of Ground-Theoretic Content' (2018) and its sequel 'Ground-Theoretic Equivalence' (2021a) develop such a modification of truthmaker semantics, and an account of (what I would now describe as[15]) representational ground based on it.

4.1. **Worldly Ground.** In 'A Note on the Logic of Worldly Ground', Stefan Roski and I identify a problem for the impure logic of worldly ground proposed in Correia (2010), and hence for the truthmaker semantics for ground corresponding to that logic: it implies certain implausible claims of partial grounding. To vary the central case of the paper, it implies, for example, that the fact that Arvo Pärt is a composer helps ground the fact that there are German composers—even though Pärt is Finnish, not German, which would seem to prevent him playing any (suitable[16]) role in bringing it about that there are German composers. After presenting Correia's account and highlighting some of its attractive features (§2), we show that a principle of factual equivalence endorsed by Correia—the distributivity of $\vee$ over $\wedge$, $A \vee (B \wedge C) \approx (A \wedge B) \vee (A \wedge C)$—leads to implausible grounding claims of the sort just described (§3). We go on to present an alternative derivation of the same result from a related purely structural principle of *convexity* for ground, to the effect that $\Gamma < C$ and $E < C$ imply that $\Delta < C$ whenever $\Gamma \subseteq \Delta \subseteq E$ (§4). Finally, we suggest a diagnosis of the problem by tracing it back to an overly simplistic conception of the kind of *relevance* characteristic of ground (§5). Sections 3 and 5 represent my contribution to the paper, sections 2 and 4 are Roski's.

It should be mentioned that the problem we describe was also discovered, independently and at the same time, by Correia himself (2016), who then carried out some of the technical work required for an improved worldly logic and a truthmaker semantics for that logic, specifically concerning the appropriate logic of factual equivalence. My forthcoming 'Truthmaker Equivalence' (202xa) expands on Correia's results. While Correia thus agreed with our criticism, some authors have subsequently tried to defend his original account against our apparent counter-examples; cf. (Lovett, 2020a: p. 23) and (Elgin, 2021: pp. 13ff).

15 As I pointed out above, it is not completely clear how exactly the worldly/representational divide should be understood. In the first of the two papers, I emphasize that under *some* natural ways to explicate the distinction, the account I develop may still be classified as worldly. In the second, I explicitly state a number of alternative ways to understand the distinction, and explain how my account is classified under each of them. I am now inclined to think that under the most useful way(s) of drawing the worldly/representational distinction, the approach comes out as representational.

16 His compositions might have inspired some Germans to become composers, of course, but that does not make his being a composer *ground-theoretically* relevant to there being German composers.

'Singular Troubles with Singleton Socrates' develops a different objection to Fine's truthmaker semantics for worldly ground, and proposes a modification that avoids the difficulty. In a nutshell, the problem may be described as follows. Under the usual structural assumptions of asymmetry and transitivity, the grounds (if any) of a given fact form a tree-like structure, with the fact in question as its root. But it turns out that Fine's truthmaker semantics imposes a certain additional constraint on the tree structures that the set of grounds of a fact can exhibit in that it requires proper branching at every node. That is, whenever a given fact has one immediate full ground, it has at least two. I argue that this assumption is problematic: there is a range of plausible cases in which a fact has just one immediate full ground. For example, it is quite plausible that the fact of the existence of a given singleton set has as its sole immediate full ground the existence of the set's sole member. I then develop a modification of the truthmaker account of ground which makes room for these grounding structures, and prove that the modification leaves unchanged the plausible parts of the logic of ground obtained under the original Finean semantics.

As indicated above, we may connect the point here to the observation that the original truthmaker account renders ground a *logical* notion. Closely related to this observation is the fact, highlighted in the paper, that the account recognizes exactly two methods by which we may progress from a given fact to a strict ground. One consists in the decomposition of a conjunctive fact into its conjuncts, the other to the specification of a disjunctive fact by showing which of its disjuncts obtains. The problem is that a number of plausible *metaphysical* grounding links do not seem to be of either of these types. The link between the existence of Socrates and the existence of {Socrates} provides a particularly simple example, but the general phenomenon may well be more wide-spread and thereby more significant than the example of singleton sets perhaps suggests. For instance, on certain meta-ethical views, it may be very plausible to hold that each normative fact is fully grounded by descriptive facts, while denying that a normative fact is either a conjuction or disjunction of descriptive facts. Such a view can then be captured within my modified account of worldly ground, but not within the original, simpler account proposed by Fine.[17] One might therefore see Fine's original account as adequate relative to a restricted notion of exclusively *logical* worldly grounding, and my modification as adequate for the wider and more general notion of *metaphysical* worldly grounding, which comprises the logical instances as a special case.

[17] I should also note that these problems are already anticipated to some extent in Correia's 2010. Correia's *Reduction Theorem* is in effect a variant on the logical definition of ground described above, and Correia himself points out that it seems to be in tension with grounding connections of the sort that obtains between the existence of Socrates and his singleton set; cf. (Correia, 2010: pp. 271f).

4.2. **Representational Ground.** My other contribution to the logic of ground is a novel conception and theory of content, more fine-grained than Fine's truthmaker theory of content. Based on this theory, we can give a semantics of ground that validates the simplest principles, like (∧) and (∨), about the grounds of truth-functionally complex truths. The fundamental idea is to enrich the truthmaker framework by taking into account not just what states make a given statement true, but also—in a sense to be explained—*how* the statement may be made true by those states. Consider the disjunction that the ball is blue or red. Both the state of the ball being blue and the state of the ball being red are truthmakers of this disjunction. But there is an intuitive difference between the two states concerning how they make the disjunction true: the state of the ball being blue makes it true *by* making true the disjunct that the ball is blue, whereas the state of the ball being red makes it true *by* making true the disjunct that the ball is red. Crucially, by appealing to these ways, or *modes* of truthmaking, we can distinguish even between the claim that the ball is red, and the claim that the ball is red or the ball is red. While both are made true only by the state that the ball is red, it is only the second claim that this state makes true *by* making it true that the ball is red.

'Towards a Theory of Ground-Theoretic Content' introduces the basic idea and develops a formal implementation that I call the *mode-ified* truthmaker theory of content. On the basis of that theory, ground may then be defined by taking Γ to ground P iff making true Γ is one of the modes in which P may be made true. I go on to prove that the resulting account validates all the structural principles of Fine's (2012d) pure logic of ground as well as the simple rules for the truth-functional connectives proposed in Fine (2012c), including (∧) and (∨) above. The problem of formulating a complete system of axiom and rules for this semantics is still open. 'Ground-Theoretic Equivalence' addresses the question of just how fine-grained a view of ground-theoretic content we can and should take under the mode-ified truthmaker approach. I describe three successively stricter standards of equivalence, determine sound and complete axiomatizations of their respective logics, and offer some reasons for favouring the intermediate one.

5. Part III—Truthmakers, Ground, and Relevance

Ground and truthmaking, as I have stressed, share the feature that they impose a strong requirement of *relevance* on the connection between the things they relate. For some facts to ground another, the former must not include anything irrelevant to the latter. Every part of every one of the grounding facts must play a part, so to speak, in bringing about that the grounded fact obtains. And similarly for truthmaking: for a state to make true a given proposition, every part of that state must play a part in making the proposition true.

The shared connection to relevance is responsible for other features shared by ground and truthmaking. It is the requirement of relevance that makes both ground and truthmaking *non-monotonic*: expanding either a full ground or a truthmaker of a fact or truth P in arbitrary ways does not in general yield another ground or truthmaker of P, since the added element may be irrelevant to P. And it is the requirement of relevance that makes both relations *hyperintensional*. For relevance requires a connection in *content* between the things relevant to each other, while it is well-known that logical equivalence does not ensure such a connection: that snow is white if snow is white and that the weather is cold if the weather is cold are both logical truths, and hence logically equivalent, but lack a connection in content.

However, ground and truthmaking are far from the only philosophically significant concepts with a tight link to relevance. Others include explanation, causation, confirmation, justification or evidential support, as well as certain conceptions of entailment: all these relations seem to be such that for things to instantiate them is for them to be relevant to each other in a certain way. These observations suggest a number of intriguing questions: Is there a general, and to some extent unified phenomenon of relevance in play here? Do all or most of them also exhibit the hyperintensionality and non-monotonicity of ground and truthmaking? To what further commonalities, if any, does the shared connnection to relevance give rise? And might the framework of truthmaker semantics provide an adequate formal background within which to study these relations and their similarities and interconnections?

Beginning with the reflections on ground-theoretic relevance in Krämer and Roski (2015), these sorts of questions have provided the overarching theme and context for most of my subsequent work and led me to widen the scope of my studies to a broader range of relevance relations beyond ground and truthmaking. The essays in this final part of the collection have all resulted from this work: 'Difference-Making Grounds' (Krämer and Roski, 2017) discusses *explanatory* relevance and its connection to ground. 'The Whole Truth' (202xb) and 'That's It! Hyperintensional Total Logic' (2023) contribute to our understanding of *logical* forms of relevance by developing a relevance-sensitive conception of the notion of the whole truth and related notions of totality. 'A Hyperintensional Criterion of Irrelevance' (2017b) and 'Mighty Belief Revision' (2022) concern *epistemic* forms of relevance. The first defends a hyperintensional, truthmaker-based account of what it means for evidence to be irrelevant to a given hypothesis. The second uses the truthmaker framework to develop a novel account of rational belief revision based on a hyperintensional conception of the update, i.e. the new information triggering a revision of a subject's beliefs.

5.1. **Explanatory Relevance: Difference-Making.** Grounding is closely related to a particular kind of explanation.[18] Sometimes, at least, when some fact is grounded by some other facts, we can give a metaphysical explanation of why the former fact obtains by citing its grounds. When this is so, the grounding facts are not just ground-theoretically but also *explanatorily* relevant to the grounded fact. But does every grounding connection give rise to an adequate explanation? Or is it only a special class of grounds that are explanatorily relevant to what they ground? 'Difference-Making Grounds', co-authored with Stefan Roski, argues for the latter option. We proceed by considering an influential thought concerning the relation of causal and explanatory relevance—that it is mainly a certain subclass of the causal factors of an event that are explanatorily relevant to the latter, namely those that *make a difference* to the caused event's occurring—and then transposing it to the ground-theoretic context.

We first define a notion of difference-making partial grounds in rough analogy to the explication of difference-making causal factors proposed by Strevens (2008) (§2). We then determine some of its formal features under standard assumptions about the logic of ground, most importantly its failure to inherit the *transitivity* of partial ground (§§3-4). On the basis of this discussion, we then suggest that—much as in the causal case—it is mainly the difference-making partial grounds that help explain the grounded fact, and use this idea to rebut an argument, due to Schaffer (2012), against taking partial ground to be transitive (§5), while criticizing Schaffer's own preferred reaction to that argument (§6). We end with some suggestions on how to further develop and refine the proposed account of grounding-explanatory relevance (§7). An appendix proves that under any of the standard assumptions about the logic of ground, there are purely logical cases both of difference-making and non-difference-making partial grounding. My contribution to the paper consists of §§3, 4 and 6 as well as the appendix, the other sections represent Roski's work. Since the publication of the paper, our conception of difference-making has received some uptake and discussion both within the grounding literature (Makin, 2019; Richardson, 2020, 2021; Woods, 2018) and elsewhere, in the philosophy of science (Loew and Hüttemann, 2022), meta-ethics (Wodak, 2020; Väyrynen, 2018), and the general theory of relevance (Yablo, 202x).

5.2. **Logical Relevance: The Whole Truth.** A statement is a truth when it describes its subject matter accurately. Among the accurate descriptions of a subject matter, we may furthermore distinguish between those that *exhaustively* describe that subject matter, and those that describe it only partially. The former may be regarded as the *whole* or *complete* truth with respect to the subject matter in question. The notion of a complete

[18] This is accepted and emphasized by just about every participant in the debate; cf. e.g. the opening passage of Fine (2012c). Glazier (2020) is a helpful survey on the different possible views regarding the exact nature of the connection.

truth finds application in a wide and diverse range of contexts. One example concerns the standard interpretation of questions and answers in natural language. When I am asked who is coming for dinner, the natural interpretation of the question is as a request for an exhaustive list of guests. Relatedly, if I answer that Bill, Bob, and Sarah are coming, the natural interpretation of my answer is that these are all the guests coming for dinner. Similar observations apply in most, if not all cases of question-answer discourse. The standard interpretation of such discourse thus implicitly invokes the notion of a truth that is complete with respect to the subject matter introduced by the relevant question.

Another example arises in metaphysics, when we ask what *kind* of truth may be *absolutely* complete, i.e. complete with respect to the most comprehensive, all-comprising subject matter. Thus, the doctrine of *physicalism* may be explicated, in a first approximation, as the claim that the conjunction of all physical truths is absolutely complete. At second glance, though, it is clear that this explication of physicalism is not quite right: it makes physicalism clearly false. For if physicalism is true, we may suppose, then there are no demons, and so it is a truth that there are no demons. But then a truth could not be complete, could not be the whole truth, without at least entailing the truth that there are no demons. It is not at all plausible, however, that the conjunction of all physical truths entails, by itself, the absence of non-physical items such as demons. An attractive way to refine the explication of physicalism to avoid this difficulty makes use of a so-called *totality operator*, which may be glossed by natural language locutions like 'and that's it'. The idea, first proposed by Chalmers and Jackson (2001: pp. 317ff), is that adding to the conjunction of all physical truths: *and that's it*, serves to exclude the possibility of there being further, non-physical truths not entailed by the physical truths. Note that a similar fix is required in the case of the dinner-question. To say Bill, Bob, and Sarah are coming leaves open the possibility that so is Barack Obama, and thus does not stand a chance of qualifying as the whole truth regarding the matter of our dinner guests. To say that Bill, Bob, and Sarah are coming, *and that's it* (as far as the matter of our dinner guests goes) is not open to the same complaint. Thus, a totality operator is also of use in explicating the content of purportedly exhaustive answers to questions.

Given the theoretical significance of the notion of the whole truth and the related concept of a totality operator, a formally precise account of these notions and their general logical features is desirable. 'The Whole Truth' and 'That's it! Hyperintensional Total Logic' develop at least part of such an account. In the former, I argue at length that in many contexts—in particular, when non-contingent subject matters are at issue—we need a hyperintensional notion of completeness on pain of trivializing the concept. Simply put, if the whole truth P about a subject matter is necessary, then it is intensionally equivalent to a triviality of the form $Q \vee \neg Q$. So on an intensional conception, the latter must count as the whole truth if the former does. A better, more discriminating conception may be obtained by imposing relevance requirements: roughly speaking, I propose to regard a proposition as a complete truth only if every part of the subject matter in question is relevant to making the proposition true. I go on to sketch a formal implementation of this proposal within the framework of truthmaker semantics. The second paper develops the proposal in more detail, defines two complementary *that's it*-operators to be used in articulating complete truths, and determines their logics.[19]

5.3. **Epistemic Relevance: Evidence and Belief Revision.** The forms of relevance characteristic of ground, truthmaking, and (relevant) logical entailment may plausibly be regarded as *mind-independent* and *absolute*. That is, they are ways to be relevant to whether something—a fact, a truth, a conclusion—obtains or is the case, with no immediate connection to, far less dependence on, our knowledge or beliefs about the matter. But there are also important forms of relevance—forms of *epistemic* relevance, we might say—that are ways to be relevant specifically to what we may know or should believe about a given issue. Such forms of relevance are then mind-dependent in at least one important sense, and they will typically be relative to an overall background system of beliefs. The final two essays included in this collection, 'A Hyperintensional Criterion of Irrelevance' and 'Mighty Belief Revision', are devoted to the study of epistemic forms of relevance.

Given any hypothesis, we may intuitively classify available evidence as either evidentially relevant or evidentally irrelevant (relative to a given subject's background knowledge or beliefs) to the hypothesis. What does it mean for evidence to be irrelevant? The most popular answer appeals to the probabilistic framework of Bayesian epistemology: a piece of evidence is irrelevant to a given hypothesis just in case the probability of the hypothesis given the evidence is equal to the prior probability of the hypothesis. An obvious and familiar difficulty for this claim is that a piece of evidence may in part support a hypothesis and in part speak against it.[20] In such a case, the net effect, as it were, may still be to leave the probability of the hypothesis unchanged. But there seems to be a clear sense in which such evidence is still relevant. Certainly, such evidence seems relevant in a way in which evidence no part of which bears in any way on the hypothesis is not relevant. A natural idea is thus to count evidence as irrelevant only if *no part* of it is relevant in the Bayesian, probabilistic sense. But to make this proposal precise, we then need to explain what it is for something to be a part of a given piece of evidence, and this turns out to be difficult. The first, destructive part of 'A Hyperintensional Criterion of Irrelevance' considers the explication proposed by Gemes (2007) and shows that

[19] These papers are too recent to have garnered attention in the published literature, but a substantial response to 'The Whole Truth' from Kit Fine will appear alongside my paper in the same collection.

[20] John Maynard Keynes appears to have been the first to explicitly point out the difficulty; cf. (Keynes, 1929: p. 79).

his proposal still misclassifies relevant evidence as irrelevant. The second, constructive part, argues that we can improve on Gemes' view by using the hyperintensional, truthmaker-based notion of the parts of a statement developed by Fine (2016). I then propose a further refinement of the resulting view using a notion of a *helpful* part of a proposition, which is very closely related to the notion of a difference-making partial ground defined in Krämer and Roski (2017).[21]

Just as we may ask, given some hypothesis, what information is relevant to whether we should accept the hypothesis, we may also ask, given some information, which beliefs it is relevant to. In particular, we may ask which existing beliefs a given piece of new information might make it rational to give up. This is the central question of a theory of rational belief revision. The goal, in developing such a theory, is to articulate the general, broadly formal constraints on the rational ways to revise one's beliefs in the light of new information. The most influential attempt to do this is the so-called AGM theory of belief revision, so named after its authors Carlos Alchourrón, Peter Gärdenfors, and David Makinson (1985). A key assumption of this approach is that both (ideally rational[22]) belief systems and updates—pieces of information triggering a revision of a subject's beliefs—are individuated intensionally. That is, rational methods of belief revision are blind to any differences that may obtain between logically equivalent belief systems or logically equivalent updates. In 'Mighty Belief Revision', I argue that this assumption should be given up at least with respect to the update. Motivated by consideration of a particular puzzle case in which AGM appears to give the wrong results, I propose a novel conception of the update as *mighty*—as encoding both what the subject learns *might* be the case as well as what they learn must be the case—which requires a hyperintensional standard for individuating updates. I go on to implement the proposal in the form of a precise definition of rational belief revision functions within the framework of truthmaker semantics and prove that, modulo hyperintensionality, the resulting account validates counterparts of the usual AGM principles and provides a satisfactory account of the puzzle case.

21 Historically, the present paper was written first, and 'Difference-Making Ground' resulted from the idea that the notion of a helpful part could help deflect Schaffer's attack on the transitivity of ground.

22 Since their aim is to describe the general properties of rational ways to revise one's beliefs, theories of belief revision usually work with a highly idealized conception of the epistemic agents under consideration as maximally rational and logically perfect reasoners.

References

Alchourrón, C. E., P. Gärdenfors, and D. Makinson (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic 50*, 510–530.

Anglberger, A., J. Korbmacher, and F. Faroldi (2016). An Exact Truthmaker Semantics for Permission and Obligation. In O. Roy, A. Tamminga, and M. Willer (Eds.), *Deontic Logic and Normative Systems, 13th International Conference, DEON 2016*, pp. 16–31. London: College Publications.

Armstrong, D. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.

Armstrong, D. (2004). *Truth and Truthmakers*. Cambridge: Cambridge University Press.

Barwise, J. and J. Perry (1983). *Situations and Attitudes*. Cambridge, MA: The MIT Press.

Batchelor, R. (2010). Grounds and Consequences. *Grazer Philosophische Studien 80*, 65–77.

Bolzano, B. (1837). *Wissenschaftslehre*. Aalen: Scientia.

Button, T. and R. Trueman (2021). Against Cumulative Type Theory. *The Review of Symbolic Logic*, 1–43.

Chalmers, D. J. and F. Jackson (2001). Conceptual Analysis and Reductive Explanation. *The Philosophical Review 110*, 315–361.

Correia, F. (2010). Grounding and Truth-Functions. *Logique et Analyse 53(211)*, 251–279.

Correia, F. (2016). On the Logic of Factual Equivalence. *Review of Symbolic Logic 9(1)*, 103–122.

Correia, F. and B. Schnieder (2012). Grounding: An opinionated introduction. In F. Correia and B. Schnieder (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality*, pp. 1–36. Cambridge: Cambridge University Press.

deRosset, L. (2021). Hollow Truth. *The Philosophical Review 130(4)*, 533–581.

Elgin, S. (2021). The Semantic Foundations of Philosophical Analysis. *Review of Symbolic Logic FirstView*, 1–21.

Fine, K. (2001). The Question of Realism. *Philosopher's Imprint 1*, 1–30.

Fine, K. (2010). Some Puzzles of Ground. *Notre Dame Journal of Formal Logic 51(1)*, 97–118.

Fine, K. (2012a). Counterfactuals without Possible Worlds. *The Journal of Philosophy 109(3)*, 221–246.

Fine, K. (2012b). A Difficulty for the Possible Worlds Analysis of Counterfactuals. *Synthese 189*, 29–57.

Fine, K. (2012c). Guide to Ground. In B. Schnieder and F. Correia (Eds.), *Metaphysical Grounding*, pp. 37–80. Cambridge: Cambridge University Press.

Fine, K. (2012d). The Pure Logic of Ground. *The Review of Symbolic Logic 5(1)*, 1–25.

Fine, K. (2016). Angellic Content. *Journal of Philosophical Logic 45*(2), 199–226.

Fine, K. (2017a). A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic 46(6)*, 625–674.

Fine, K. (2017b). A Theory of Truthmaker Content II: Subject Matter, Common Content, Remainder and Ground. *Journal of Philosophical Logic 46(6)*, 675–702.

Fine, K. (2017c). Truthmaker Semantics. In B. Hale, C. Wright, and A. Miller (Eds.), *A Companion to the Philosophy of Language, 2nd ed.*, pp. 556–577. Chichester: Wiley Blackwell.

Fine, K. (2018). Compliance and Command ii, Imperatives and Deontics. *The Review of Symbolic Logic 11(4)*, 634–664.

Fine, K. (2021). Verisimilitude and Truthmaking. *Erkenntnis 86*, 1239–1276.

Fine, K. (msa). Review of Steve Yablo's Aboutness. Unpublished manuscript.

Fine, K. (msb). A Theory of Partial Truth. To appear in a volume for David Kaplan.

Florio, S. and N. K. Jones (2021). Unrestricted Quantification and the Structure of Type Theory. *Philosophy and Phenomenological Research 102(1)*, 44–64.

Florio, S. and N. K. Jones (202x). Two Conceptions of Absolute Generality. *Philosophical Studies forthcoming*.

Fritz, P. (2020). On Higher-Order Logical Grounds. *Analysis 80(4)*, 656–666.

Gemes, K. (2007). Irrelevance: Strengthening the Bayesian criteria. *Synthese 157(2)*, 161–166.

Glazier, M. (2020). Explanation. In M. J. Raven (Ed.), *The Routledge Handbook of Metaphysical Grounding*, pp. 121–132. Routledge.

Goodman, J. (2022). Grounding Generalizations. *Journal of Philosophical Logic Online First*.

Jago, M. (2020). Truthmaker Semantics for Relevant Logic. *The Journal of Philosophical Logic 49*, 681–702.

Keynes, J. M. (1929). *A Treatise on Probability*. Macmillan.

Krämer, S. (2013). A Simpler Puzzle of Ground. *Thought: A Journal of Philosophy 2(2)*, 85–89.

Krämer, S. (2014). *On What There Is For Things To Be*. Frankfurt: Klostermann.

Krämer, S. (2017a). Everything, and Then Some. *Mind 126(502)*, 499–528.

Krämer, S. (2017b). A Hyperintensional Criterion of Irrelevance. *Synthese 194(8)*, 2917–2930.

Krämer, S. (2018). Towards a Theory of Ground-Theoretic Content. *Synthese 195(2)*, 785–814.

Krämer, S. (2020). Puzzles. In M. J. Raven (Ed.), *The Routledge Handbook of Metaphysical Grounding*, pp. 271–282. Routledge.

Krämer, S. (2021a). Ground-Theoretic Equivalence. *Synthese 198*, 1643–1683.

Krämer, S. (2021b). Singular Troubles with Singleton Socrates. *Philosophy and Phenomenological Research 103(1)*, 40–56.

Krämer, S. (2022). Mighty Belief Revision. *Journal of Philosophical Logic 51*, 1175–1213.

Krämer, S. (2023). That's It! Hyperintensional Total Logic. *Journal of Philosophical Logic Online First*.

Krämer, S. (202xa). Truthmaker Equivalence. In S. Leuenberger and A. Rieger (Eds.), *Themes from Alan Weir*. Synthese Library.

Krämer, S. (202xb). The Whole Truth. In F. Faroldi and F. van de Putte (Eds.), *Kit Fine on Truthmakers, Relevance, and Non-Classical Logic*. Springer.

Krämer, S. and S. Roski (2015). A Note on the Logic of Worldly Ground. *Thought: A Journal of Philosophy 4(1)*, 59–68.

Krämer, S. and S. Roski (2017). Difference-Making Grounds. *Philosophical Studies 174*, 1191–1215.

Loew, C. and A. Hüttemann (2022). Are We Free to Make the Laws? *Synthese 200(4)*.

Lovett, A. (2020a). The Logic of Ground. *The Journal of Philosophical Logic 49*, 13–49.

Lovett, A. (2020b). The Puzzles of Ground. *Philosophical Studies 177*, 2541–2564.

Makin, M. (2019). Rigid/Non-rigid Grounding and Transitivity. *Inquiry 62(2)*, 136–150.

Mulligan, K., P. Simons, and B. Smith (1984). Truth-Makers. *Philosophy and Phenomenological Research 44*, 287–321.

Raven, M. J. (Ed.) (2020). *The Routledge Handbook of Metaphysical Grounding*. Routledge.

Richardson, K. (2020). Grounding Pluralism: Why and How. *Erkenntnis 85*, 1399–1415.

Richardson, K. (2021). Grounding is Necessary and Contingent. *Inquiry 64(4)*, 453–480.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale and A. Hoffmann (Eds.), *Modality: Metaphysics, Logic, and Epistemology*, pp. 109–136. Oxford: Oxford University Press.

Roski, S. (2017). *Bolzano's Conception of Grounding*. Frankfurt am Main: Klostermann.

Roski, S. and B. Schnieder (Eds.) (2022). *Bolzano's Philosophy of Grounding*, Oxford. Oxford University Press.

Rothschild, D. and S. Yablo (202x). Permissive Updates. In F. Faroldi and F. van de Putte (Eds.), *Kit Fine on Truthmakers, Relevance, and Non-Classical Logic*. Springer.

Schaffer, J. (2009). On What Grounds What. In D. J. Chalmers, D. Manley, and R. Wassermann (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 347–383. Oxford: Oxford University Press.

Schaffer, J. (2012). Grounding, Transitivity, and Contrastivity. In B. Schnieder and F. Correia (Eds.), *Metaphysical Grounding*, pp. 122–138. Cambridge: Cambridge University Press.

Schnieder, B. (2011). A Logic for "Because". *The Review of Symbolic Logic 4(3)*, 445–465.

Strevens, M. (2008). *Depth*. Cambridge Mass.: Harvard University Press.

van Fraassen, B. (1969). Facts and Tautological Entailment. *Journal of Philosophical Logic 66(15)*, 477–487.

Väyrynen, P. (2018). Normative Commitments in Metanormative Theory. In J. Suikkanen and A. Kauppinen (Eds.), *Methodology and Moral Philosophy. Routledge Studies in Ethics and Moral Theory*, pp. 193–213. Routledge.

Williamson, T. (2003). Everything. *Philosophical Perspectives 17*, 415–65.

Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Wodak, D. (2020). Redundant Reasons. *Australasian Journal of Philosophy 98(2)*, 266–278.

Woods, J. (2018). Emptying a Paradox of Ground. *The Journal of Philosophical Logic 47*, 631–648.

Yablo, S. (2011). A Problem about Permission and Possibility. In A. Egan and B. Weatherstone (Eds.), *Epistemic Modality*, pp. 270–294. Oxford University Press.

Yablo, S. (2015). *Aboutness*. Princeton: Princeton University Press.

Yablo, S. (202x). Relevance Without Minimality. In A. Egan, P. van Elswyk, and D. Kindermann (Eds.), *Unstructured Content*. Oxford University Press.

Everything, and Then Some

Stephan Krämer
University of Hamburg
stephan.kraemer@uni-hamburg.de

On its intended interpretation, logical, mathematical and metaphysical discourse sometimes seems to involve absolutely unrestricted quantification. Yet our standard semantic theories do not allow for interpretations of a language as expressing absolute generality. A prominent strategy for defending absolute generality, influentially proposed by Timothy Williamson in his paper 'Everything' (2003), avails itself of a hierarchy of quantifiers of ever increasing orders to develop non-standard semantic theories that do provide for such interpretations. However, as emphasized by Øystein Linnebo and Agustín Rayo (2012), there is pressure on this view to extend the quantificational hierarchy beyond the finite level, and, relatedly, to allow for a cumulative conception of the hierarchy. In his recent book, *Modal Logic as Metaphysics* (2013), Williamson yields to that pressure. I show that the emerging cumulative higher-orderist theory has implications of a strongly generality-relativist flavour, and consequently undermines much of the spirit of generality absolutism that Williamson set out to defend.

1. Introduction

In ordinary discourse, most of our quantifications are restricted to a set of contextually relevant objects. I say 'There is no beer', meaning not that there is absolutely no beer in the entire universe, but that there is no beer in my fridge, and thus no contextually relevant beer. In logical, metaphysical and mathematical discourse, in contrast, we often seem to generalize without any such restrictions. In typical utterances of 'Nothing has contradictory properties', 'Everything is self-identical', 'Everything is either abstract or concrete', or 'Nothing is a member of the empty set', it seems, absolutely nothing is excluded as contextually irrelevant.

However, the appearance that absolute generality can thus be expressed comes under pressure from a number of theoretical considerations.[1] The one that is most important for our purposes is that in our most successful and best understood semantic theories,

[1] For an excellent overview of the debate, including a discussion of other arguments against the possibility of absolute generality, see Rayo and Uzquiano (2006a), as well as Florio (2014).

Mind, Vol. 126 . 502 . April 2017
doi:10.1093/mind/fzv187 Advance Access publication 30 September 2016

quantification is always interpreted with respect to some set that constitutes the presumed domain of discourse. Since there is no universal set, these semantic theories do not allow for absolute generality. Some philosophers hold that this apparent limitation of such theories cannot be overcome in a satisfactory manner, and have therefore embraced *(Generality) Relativism*, the view that, initial appearances notwithstanding, there can be no such thing as absolutely general discourse.[2]

In his paper 'Everything' (2003), Timothy Williamson mounts a forceful defence of the opposition to Relativism, that is, *(Generality) Absolutism*. As part of this defence, he proposes an alternative kind of semantics in which object-language quantifiers need not be interpreted as restricted to a set.[3] The crucial move that enables him to avoid this limitation is the employment of *higher-order* quantification in the metalanguage.[4] As Williamson notes, the obvious generalization of his proposal requires the use of at least the whole finite hierarchy of orders of quantification. Moreover, as Øystein Linnebo and Agustín Rayo (2012) have shown, given a number of plausible further assumptions, we need to countenance even quantification of *transfinite* and *cumulative* orders. And when Williamson returns to the issue in his *Modal Logic as Metaphysics* (2013), he explicitly avails himself of such quantifiers.

The purpose of this paper is to show that the move to transfinite, cumulative orders of quantification undermines much of the spirit, if not the letter, of the Absolutist picture that Williamson wishes to defend. Moreover, since the relevant components of Williamson's view are difficult or impossible to reject given his basic approach of using higher-order resources to defend absolute generality, the

[2] Sometimes in the debate over absolute generality a distinction is made between a *metaphysical* question—roughly, whether there is an all-inclusive domain of discourse—and a linguistic or *availability* question—roughly, whether we could quantify over such a domain, if it existed (see Rayo and Uzquiano 2006a, p. 2, where they credit Kit Fine with having first emphasized this point). The distinction, if it can be made, is not of particular importance for our present purposes. As I understand the distinction, our discussion concerns the availability question throughout.

[3] It may be objected that familiar truth-theoretic semantics formulated in the Davidson-Tarski style have no need to interpret object-language quantifiers as restricted to a set. That is, of course, correct; the difficulty arises when we try to formulate a theory that specifies the truth conditions of the object-language sentences for arbitrary interpretations of the language. We shall come back to this point shortly.

[4] The same kind of move has been employed for a similar objective by George Boolos (1985) and, further developing Boolos's suggestions, by Rayo and Uzquiano (1999).

argument also presents a severe challenge to any view following this general approach. I first present the current state of the debate. §2 describes the problem that Absolutists face when trying to formulate an adequate semantics for absolutely general quantification. §3 presents Williamson's higher-orderist solution to the problem, and then explains both the notion of quantification at transfinite and cumulative orders and why Williamson avails himself of such quantifiers. §4 develops my argument that the cumulative higher-order picture has implications strongly reminiscent of Relativism. I first state the argument in informal terms and indicate the obstacles to formalizing it in the non-cumulative higher-orderist's canonical language. I then show how the introduction of cumulative resources allows us to overcome these obstacles. The crucial bit of cumulative ideology that we need is a higher-order, cumulative analogue of the notion of identity. That notion is discussed in more detail in §5. Finally, I ask to what extent the problems Williamson raises for Relativism have counterparts that apply to his own cumulative higher-orderist version of Absolutism. §6 points out that cumulative higher-orderism faces a similar difficulty to standard Relativism concerning the adequate formulation of its apparent Relativist commitments. §7 turns to the criticism of standard Relativism that its restricted interpretations often seem to constitute weird misinterpretations of relevant object-language discourse. I show that this criticism also applies on the cumulative higher-orderist picture, though in a somewhat less dramatic form. I conclude in §8 that Williamson's cumulative higher-orderist Absolutism is a lot closer to Relativism than it first appears.

2. The semantic challenge for absolutism

A peculiar feature of the dispute over absolute generality is that it is not very easy to say just what it takes to be an Absolutist, or a Relativist, who is worthy of the name.[5] Take Absolutism. It is not enough to say 'It is possible to quantify over absolutely everything', or even

(1) I am now quantifying over absolutely everything.

The reason is that *everyone* can say that, since any utterance of (1) will express a truth. For whatever the quantifier phrase 'absolutely everything' ranges over in such an utterance, what the speaker then claims

[5] On this issue, see Williamson (2003, §V), Florio (2014, p. 2), Rayo and Uzquiano (2006a, pp. 2f.), and the references given there.

is merely that they are quantifying over everything in that range, which cannot fail to be true. However, if the quantifier phrase ranges over only a restricted domain, the claim made by the speaker, though true, is irrelevant to the spirit of Absolutism, if perhaps not to its letter. It is only if the quantifier phrase ranges over absolutely everything that the speaker makes a claim that is relevant to the debate, such that accepting it makes one an Absolutist proper.

If the speaker is to count as an Absolutist, therefore, we shall require them to back up their utterance of (1) by insisting, in a metalinguistic utterance, on an interpretation of their utterance on which it generalizes over everything. And while the speaker remains firmly on the safe side with their utterance of (1), once they back it up in this way in the metalanguage, as Williamson shows, they face a real threat of contradiction. First, let us shift our focus slightly by considering formal languages instead of utterances of English sentences like (1). Presumably, if it can be consistently maintained of an utterance of (1) that it generalizes over everything, then it is possible to specify a formal language of which it can be consistently maintained that its quantifiers range over everything. We shall therefore assume that the Absolutist commits to the following claim:

(GA) It is possible to specify a formal language containing quantifiers that, from the point of view of a suitable metalanguage, range over absolutely everything.

A language $\mathcal{L}'$ is a suitable metalanguage for a formal language $\mathcal{L}$ if it is possible to develop an adequate semantic theory for $\mathcal{L}$ in $\mathcal{L}'$.

A second assumption we shall make has to do with the *kind* of semantic theorizing about the object language that we want to be possible in our metalanguage. We shall assume that for any formal language it is possible to construct what we may call a *generalized semantics* for that language.[6] A generalized semantics for a language is a semantic theory which provides an inductive characterization of the truth conditions of the object-language sentences, relative to arbitrary ways of interpreting (at least) their non-logical constants. (A generalized semantics is therefore exactly the kind of theory we require as the basis of a broadly model-theoretic theory of logical conse-

[6] The way I set up the dispute here is indebted to Linnebo (2006) and Linnebo and Rayo (2012), from which I have also borrowed my terminology.

quence.) Our second assumption can thus be described as a principle of *Semantic Optimism*:

(SO) For any formal language $\mathcal{L}$, it is in principle possible to construct a theory adequately specifying the truth conditions of sentences of $\mathcal{L}$ relative to arbitrary ways of interpreting $\mathcal{L}$'s non-logical constants.

This principle is by no means trivial. Still, in the debate about absolute generality, it is standardly assumed that the principle enjoys a high degree of initial plausibility (see, for example, Linnebo 2006, p. 150; Linnebo and Rayo 2012, pp. 276f.). For the purposes of this paper, I shall therefore take it for granted. The challenge for the Absolutist is then to specify a formal language, and to formulate a generalized semantics for that language, so that from the point of view of the metalanguage, the quantifiers of the object language range over absolutely everything. Unfortunately, there is a powerful argument purporting to show that this cannot be done.[7]

Suppose that the Absolutist has described a formal language $\mathcal{L}_1$ whose (first-order) quantifiers they wish to maintain express absolute generality. The Absolutist must now formulate a generalized semantics for $\mathcal{L}_1$ in a suitable metalanguage. A generalized semantics is supposed to specify the truth conditions of the sentences of $\mathcal{L}_1$ relative to arbitrary ways of interpreting (the non-logical constants of) $\mathcal{L}_1$. It therefore needs to generalize over (things that model) ways of interpreting $\mathcal{L}_1$. For simplicity, assume that 'i' is a metalanguage variable that ranges over whatever the semantics uses to model ways of interpreting $\mathcal{L}_1$.

Now take some monadic predicate P of $\mathcal{L}_1$. It seems very plausible that we can in principle use any contentful monadic predicate of the metalanguage to interpret P. If so, then any instance of the following *Comprehension Schema for Predicate Interpretations*, obtained by replacing 'F' with such a metalanguage predicate, should be true.

(CPI Informal) Under at least one way of interpreting the non-logical constants of $\mathcal{L}_1$, P applies to all and only the Fs in the range of $\mathcal{L}_1$'s quantifiers.

The semantics should then imply a suitable regimentation of (CPI Informal) in the metalanguage. For brevity, we add copies '$\forall_o$' and '$\exists_o$' of $\mathcal{L}_1$'s first-order quantifiers to the metalanguage. We may then

[7] Except for minor details, the argument to follow is Williamson's (2003, pp. 425ff.).

formulate a suitable regimentation of (CPI Informal) as follows, using 'applies$_i$' to abbreviate 'applies under interpretation i':[8]

(CPI Formal 1) $\exists i \forall_o x$ (P applies$_i$ to $x \leftrightarrow Fx$)

The Absolutist also has to include in their theory an expression of the claim that, from the point of view of the semantics, $\mathcal{L}_1$'s quantifiers range over absolutely everything. An obvious way to formalize that claim is:

(GA$_1$) $\forall y \exists_o x\ x = y$

However, these assumptions jointly entail a contradiction. For consider the following instance of (CPI Formal 1), obtained by substituting '$\neg P$ applies$_x$ to x' for 'Fx':[9]

(2) $\exists i \forall_o x$ (P applies$_i$ to x iff $\neg P$ applies$_x$ to x)

Assuming that some interpretation i' verifies (2), we have:

(3) $\forall_o x$ (P applies$_{i'}$ to x iff $\neg P$ applies$_x$ to x)

Now crucially, by (GA$_1$), the range of '$\forall_o x$' comprises everything. So, in particular, it includes the interpretation i'. We may therefore instantiate (3) with i' to obtain:

(4) P applies$_{i'}$ to i' iff $\neg P$ applies$_{i'}$ to i'

But (4) is equivalent in classical logic to an explicit contradiction.

The Relativist maintains that it is the Absolutist's claim (GA$_1$) which is to blame for the contradiction, and concludes that there is more than is dreamt of by $\mathcal{L}_1$'s quantifiers:

(MORE$_1$) $\exists y \forall_o x\ x \neq y$

To be an Absolutist, in contrast, one must keep (GA$_1$), and thus find some other way out. The next section explains how Williamson proposes to do that.

3. Higher-orderist absolutism

Our above regimentation of (CPI Informal) by (CPI Formal 1) embodies a tacit assumption to the effect that the informal talk of ways of interpreting an object-language predicate is to be understood in *first-order* quantificational terms. After all, (CPI Formal 1) uses a first-order quantifier '$\exists i$' presumed to range over objects of some sort that are

[8] The variable 'i' is not allowed to occur free in a substituend for 'F'.

[9] Note that since 'i' in 'applies$_i$' was assumed to be an ordinary first-order variable, we can legitimately put the first-order variable 'x' in its place.

identified with, or taken to represent, ways of interpreting object-language predicates. Williamson suggests that this tacit assumption is mistaken; it is more plausible, according to him, to represent ways of interpreting predicates by means of a *second-order* variable (2003, pp. 452ff.). In effect, his proposal has us replace (CPI Formal 1) with a second-order analogue, in which the first-order variable '*i*' and the quantifier binding it have been replaced by a second-order variable and quantifier:[10]

(CPI Formal 2) $\exists I\, \forall_o x\ (P \text{ applies}_I \text{ to } x \leftrightarrow Fx)$

The semantic predicate 'applies$_I$' may be thought of as defined by:

(Df. applies$_I$) $P \text{ applies}_I \text{ to } x \leftrightarrow I(P, x)$

Note that from the standard (full) axiom scheme of comprehension for second-order logic, we obtain every instance of the schema

(5) $\exists I\, \forall_o x\ (I(P,x) \leftrightarrow Fx)$

in which the expression replacing '*F*' does not contain '*I*' free. This in turn guarantees the validity of (CPI Formal 2).

At least at first glance, it seems to me, this approach to the semantic challenge for Absolutism—call it the *higher-orderist* approach—has a lot to be said for it. What the Russell-paradoxical argument from (2) to (4) shows, we might say, is that there are always strictly more ways to interpret a predicate with respect to a domain of objects than there are objects in that domain, and thus ways to interpret a first-order variable with respect to that domain.[11] As a result, not every way to interpret a predicate with respect to a given domain can be represented by an object in that domain. The Relativist concludes from this that there must always be objects outside any domain with respect to which a predicate can be interpreted. The higher-orderist approach, in contrast, concludes that it is a mistake to try to represent ways of interpreting a predicate by the values of first-order variables. Instead, the thought goes, we should use second-order variables for that purpose. For crucially, there is no obstacle to holding that there are always at most as many ways to interpret a predicate with respect to a domain of objects as there are ways to interpret a predicate variable with respect to that domain.

[10] Similarly as for (CPI Formal 1), instantiation of '*F*' with a predicate in which '*I*' is free is not allowed.

[11] The talk of domains is not to be taken too literally. There is no need to assume that the objects that a predicate is interpreted as true or false of are members of some further object that we call a domain.

It is important to note, however, that the proposal depends on a very specific view—call it *higher-orderism*—of second- and higher-order quantification. Before stating the view, let me clarify my talk of quantifiers of first, second and higher orders. For present purposes, that classification is to be thought of in *syntactic* terms. On this understanding, what makes the familiar quantifiers '$\exists x$' and '$\forall x$' of $\mathcal{L}_1$ first-order is that they bind variables that stand in the syntactic position of singular terms. What we shall call *second-order* quantifiers are then quantifiers that bind variables standing in the syntactic position of expressions that form sentences when combined with one or more singular terms as their arguments, that is, ordinary *predicates*. A *third-order* quantifier, by analogy, is a quantifier binding variables that take the position of expressions forming sentences when combined with ordinary predicates as their arguments. The hierarchy extends in the obvious way to quantifiers and variables of order *n* for any finite *n*.

Higher-orderism can now be defined as the conjunction of the following three theses concerning this syntactic hierarchy. (1) Quantification of any finite order is a legitimate linguistic device. (2) Quantifications of a given order are not in general paraphrasable by quantifications of a lower order. (3) For any finite *n*, adequate semantic clauses for *n*th-order quantifications themselves employ *n*th-order quantifiers of the metalanguage in the way standard clauses for first-order quantifiers employ first-order quantifiers of the metalanguage.[12]

To get an idea of why the Williamsonian proposal requires the full strength of higher-orderism, note first that if second-order quantification were in general paraphrasable by first-order quantification, then in particular (CPI Formal 2) could be paraphrased in first-order terms. But such a paraphrase would reintroduce the inconsistency engendered by (CPI Formal 1). Moreover, given (SO), we can develop a generalized semantics for our second-order metalanguage. If we were to do that in a first-order metalanguage, perhaps by construing second-order quantifiers as ranging over properties, we should again run into a version of Russell's Paradox. Indeed, by a higher-order analogue of the above Russell-style argument, it can be shown that a generalized semantics for a second-order language cannot be given even in a second-order language: we need to use third-order

[12] I am not claiming that the three theses are independent, but it is a non-trivial question what entailment relations may obtain between them, and so it seems best to stay neutral with respect to that question in characterizing higher-orderism.

resources.[13] So Williamson's higher-orderist approach depends on the legitimacy and irreducibility of quantification of every finite order.

The irreducibilism embodied in higher-orderism has important implications for how we can read second- and higher-order quantifications in natural language. For example, at first glance the second-order quantification '$\exists X\ \forall x\ \neg Xx$' might naturally be read as 'some property is had by no object'. However, although often useful and appropriate for heuristic purposes, such a reading cannot be considered strictly adequate on the higher-orderist view. The reason is that the English quantifier here used to interpret the formal higher-order quantifier is itself *first-order*. This can be seen from the fact that if we ask for a witnessing instance of the quantification, grammar demands that the answer consist in a singular noun phrase like 'the property of being self-distinct' rather than a predicate.

It is controversial whether one can translate second-order quantification into natural language in a way that fits higher-orderism.[14] For our purposes, it does not matter too much, for even if this is possible, it seems clear that natural languages do not provide us with the resources needed to appropriately translate quantifiers of arbitrarily high finite orders. So higher-orderists must hold that we can, at least in principle, somehow come to understand third-, fourth- and higher-order quantification without the benefit of a translation into vocabulary that we independently understand.[15] In what follows, to avoid excessive formalism, and for heuristic purposes, we shall make use of various natural language constructions to approximate the higher-orderists' intended interpretation of their vocabulary, bearing in mind, though, that these may occasionally yield a slightly misleading picture of the higher-orderist view. We shall grant, moreover, both the truth of higher-orderism and the claim that higher-orderism provides sufficient expressive resources to formulate a generalized semantics for any language of finite order, that is, such that for some finite n, it contains no quantifiers of an order higher than n.

[13] Strictly speaking, the situation is slightly more complicated; it is examined in detail in Rayo (2006). The essential point remains, however. Semantic optimism forces the higher-orderist up the hierarchy of metalanguages, and if they are to permit the formulation of a generalized semantics in accordance with the higher-orderist approach, the metalanguages must include quantificational devices of ever increasing finite orders.

[14] George Boolos (1984, 1985) famously proposed a translation using English plural quantification; an alternative, more predicational reading was first suggested by Arthur Prior (1971, ch. 3) and recently developed in more detail by Rayo and Yablo (2001).

[15] See Williamson (2003, pp. 457ff.), Linnebo (2006, pp. 152ff.).

So far, I have merely described the explicit key commitments of the higher-orderist defence of absolute generality as described in, for example, Williamson (2003) and Rayo (2006). We now turn to two crucial further claims that Williamson endorses, implicitly at least, in his recent book (2013). The *first* claim is that in addition to the languages of finite order that we have already canvassed, there is also a legitimate language that contains quantifiers of *every* finite order. With respect to the ordering of languages of higher and higher finite orders, this language would occupy the level of the first limit ordinal ω, so we may call it $\mathcal{L}_\omega$. Williamson uses a language like this as his favoured background language for metaphysical theorizing, and so is obviously committed to considering such a language legitimate.[16]

The question whether the higher-orderist defence of absolute generality on its own is committed to this claim is more difficult to answer. Certainly, the legitimacy of $\mathcal{L}_\omega$ does not follow *logically* from anything the higher-orderist has said so far. Nevertheless, it would seem prima facie quite implausible to disallow it.[17] For every bit of vocabulary we find in $\mathcal{L}_\omega$ has already been deemed legitimate, since it is also found in some 'successor' language in the higher-orderist's hierarchy. And it is hard to see how pooling all these individually coherent and legitimate linguistic resources together into a single language could somehow fail to produce an equally coherent and legitimate language. The burden of proof would therefore seem to lie with anyone wishing to deny the legitimacy of the limit language $\mathcal{L}_\omega$.

If $\mathcal{L}_\omega$ is allowed, then by (SO), it is possible to give a generalized semantics for it. The *second* crucial claim of Williamson's is that such a semantics can be stated using quantification of *transfinite* orders. Note first that higher-orderism, as defined above, does not by itself provide sufficient resources to formulate a semantics for $\mathcal{L}_\omega$ in accordance with the higher-orderist approach. For on this approach, to interpret expressions of order *n*, we need to make use of an interpretation variable of at least order *n*. But every kind of variable higher-orderism provides us with belongs to some finite order *n*. And for every finite order *n*, since $\mathcal{L}_\omega$ contains quantifiers and variables of every finite order, it contains variables of order $n + 1$. So no interpretation variable of a finite order can be used to interpret all the expressions of $\mathcal{L}_\omega$.

[16] The modal higher-order language ML_P proposed by Williamson (2013, ch. 5) contains quantifiers of every finite order, and is thus relevantly like $\mathcal{L}_\omega$.

[17] On this point, see also Linnebo and Rayo (2012, pp. 275f.) and Rayo (2006, pp. 246ff.).

What Williamson does, therefore, in specifying a semantics for a limit language, is avail himself of predicates, and quantifiable predicate variables, of *transfinite* orders.[18]

Aside from their transfinite character, Williamson's predicates and predicate variables of transfinite order have a further striking feature that sets them apart from any of the expressions we have encountered so far. That feature concerns what expressions they accept as arguments. Usually, a predicate or predicate variable accepts as arguments expressions of the next lower order: a first-order predicate accepts singular terms, a second-order predicate accepts first-order predicates, and so on. This rule cannot extend to our new predicates, though. For these are of order ω, and there is no next lower order to ω. Instead, these predicates accept expressions of *any* lower—that is finite—order in their argument place. I shall therefore describe these new predicates as *syntactically cumulative*. It turns out that a syntactically cumulative ω-order predicate variable is just what is needed, and just what Williamson uses, for an interpretation variable in a generalized semantics for $\mathcal{L}_\omega$.[19]

Let me stress that for the purposes of this paper, the admission of *cumulative* resources is the important point. The admission of $\mathcal{L}_\omega$ and transfinite orders of quantification matter only in virtue of their bearing on this point. And while the above consideration provides perhaps the most principled and compelling case for the legitimacy of cumulativity, it should be noted that even independently of limit languages and transfinite orders of quantification, higher-orderists may have reason to be sympathetic to this claim, for since they allow quantification of every finite order, they are already committed to a fairly liberal standard for admissible linguistic devices. And from a logico-mathematical point of view, cumulative resources make perfect sense

[18] See Williamson (2013, pp. 236ff.). I do not know whether the move to transfinite orders is strictly the *only* way to give a generalized semantics for a language like $\mathcal{L}_\omega$, as seems to be suggested by the discussion in Linnebo and Rayo (2012, p. 275, and appendix B). For all I know, it might also be possible to use separate interpretation variables for every order, which could then themselves all be of finite order. However, we would then have to relativize satisfaction to infinitely many parameters, so we should have to introduce predicates with infinitely many argument places, as well as quantifiers binding an infinite set of variables. I do not know what a semantics that is adequate for this kind of infinitary language would have to look like. However, I cannot entirely rule out that there might be a coherent version of higher-orderism that allows $\mathcal{L}_\omega$ but no transfinite orders of quantification.

[19] The construction is given in Williamson (2013, pp. 236ff.). A more detailed presentation and discussion of similar constructions is given in the appendices of Linnebo and Rayo (2012).

and have a principled and well-behaved logic.[20] So it is not at all clear that there is a defensible standard of intelligibility that could serve to rule out cumulative devices without at the same time ruling out quantifiers of very high finite orders.

4. Cumulative higher-orderism and the spirit of absolutism

Let us return from the dizzy heights of ω-order quantification to the first and simplest stage in the higher-orderist's hierarchy of languages and semantic theories: the generalized semantics for the first-order language $\mathcal{L}_1$, formulated in a second-order metalanguage. Recall the sentence that started us off on the whole higher-order journey, our formalization of the claim that $\mathcal{L}_1$'s quantifiers are, from the point of view of the semantics for $\mathcal{L}_1$, absolutely general:

(GA_1) $\forall y \, \exists_o x \; x = y$

This section develops an argument that according to cumulative higher-orderism, even though the higher-order semantics for $\mathcal{L}_1$ includes (GA_1), it does not make the quantifiers of $\mathcal{L}_1$ absolutely general. More precisely, I argue that for a cumulative higher-orderist, there is a good sense in which the following holds:

(MORE) From the point of view of the higher-order semantics for $\mathcal{L}_1$, there is more than is quantified over in $\mathcal{L}_1$.

I shall begin by sketching the argument in informal terms.

In the setting of the original, first-order semantics for $\mathcal{L}_1$, (GA_1) constitutes an adequate formalization of the claim that $\mathcal{L}_1$'s quantifiers are absolutely general. For the quantifier '$\forall y$' ranges over absolutely every bit of reality that is countenanced in that semantics, and so (GA_1) says of absolutely everything countenanced in the semantics that it is in the reach of the quantifiers of $\mathcal{L}_1$. However, that situation changes when the higher-orderist proceeds to extend the language of the semantics by second-order quantifiers. Since the higher-orderist insists that these are in no way reducible to first-order ones, we have to see them as concerned with new bits of reality that are not in the range of the semantics' first-order quantifiers.[21] But then '$\forall y$' no longer

[20] On this point, see also Linnebo and Rayo (2012, p. 278). For an in-depth discussion of the logico-mathematical properties of cumulative higher-order logic, see Degen and Johannsen (2000).

[21] Higher-orderists may well consider my use of the phrase 'new bits of reality' misleading and inappropriate, so I want to emphasize that I am merely offering an informal sketch of an argument to be developed with formal rigour below. (Thanks here to an anonymous referee.)

ranges over absolutely every bit of reality that is countenanced in that semantics. Consequently, (GA_1) is no longer a plausible formalization of the absolute generality of $\mathcal{L}_1$'s quantifiers. Even on the higher-orderist picture, there is *more* than is dreamt of by $\mathcal{L}_1$'s quantifiers—what sets the picture apart from the Relativist's is only that what is more is not in the range of the semantics' first-order quantifiers but their second-order cousins. So the higher-orderist version of Absolutism is not really worthy of the name. In subscribing to (GA_1), it preserves the letter of Absolutism, but in implying (MORE), it gives up on its spirit.

Cast as it is in informal terms, the objection so far inspires limited confidence. After all, the higher-orderist has warned us that informal, natural-language approximations of their higher-order quantifications can be misleading. Can we put the informal objection on a more rigorous footing by reproducing it in a formal setting congenial to higher-orderism? More specifically, can we find a plausible formalization of (MORE) which is a consequence of the higher-orderist semantics?

Here is a somewhat flat-footed argument that we cannot. To say, as the objection alleges, that there is more, on the higher-orderist view, than is included in the range of $\mathcal{L}_1$'s quantifiers, we should have to say that there is something which is *distinct* from everything in that range. The condition of being distinct from everything in that range is expressed by

($\neq_1$) $\forall_o x \ldots \neq x$

Now, to say that there is something satisfying that condition, one has to put a variable into the gap of ($\neq_1$), and bind it by an existential quantifier. If the result is to be well-formed, however, we can only put a *first-order* variable into the gap of ($\neq_1$). But the putative extra bits of reality countenanced by the higher-orderist are supposed to be introduced only by second-order quantifiers. Such a quantifier, however, cannot bind the first-order variable in the gap of ($\neq_1$). So there is no sense in which, according to higher-orderism, there is something *more* than is ranged over by the quantifiers in $\mathcal{L}_1$. Any attempt to even formulate that claim in the higher-orderist's canonical language produces an ill-formed string. In that language, as Williamson puts it, 'quantification into predicate position is simply incommensurable with quantification into name position; the former presents no coherent threat to the absolute generality of the latter' (Williamson 2003, p. 458).

As I have stated it, this line of reasoning depends on the following assumption:

(DIST) A regimentation of (MORE) contradicts the spirit of Absolutism only if it uses ($\neq_1$) to express distinctness from anything in the range of the object language quantifier.

I shall argue that this assumption is implausible, and that once we are allowed to make use of *cumulative* higher-order resources, second-order quantification ceases to be incommensurable with first-order quantification, and does present a serious threat to the latter's absolute generality. In a first step, I show that (DIST) should be rejected even independently of any issues to do with cumulativity. In a second step, I present a number of relatively modest ways to extend the higher-orderist's second-order metalanguage with cumulative vocabulary, and show that in this extended cumulative higher-orderist setting, we can formulate and prove a well-formed regimentation of (MORE) that clashes with the spirit of Absolutism.

It will help if I first set up a system of grammatical types that allows us to describe the syntax of cumulative as well as non-cumulative expressions. For ease of comparison of the resources I employ to those used by Williamson, I base my system on the one he uses (2013, p. 221). It has just one basic type *e*, which is the type of singular terms. Then whenever $t_1, \ldots, t_n$ are types, $\langle t_1, \ldots, t_n \rangle$ is the derived (functional) type of expressions that form sentences when combined with *n* further expressions of types $t_1, \ldots, t_n$, respectively. As a limiting case, we allow $\langle\ \rangle$ as the type of sentences, that is, expressions forming sentences when combined with zero further expressions. We also add a category of cumulative types: whenever $t_1, t_2, \ldots$ are types, $[t_1, t_2, \ldots]$ is the cumulative type including all expressions belonging to any of $t_1, t_2, \ldots$[22] We write '$\langle e* \rangle$' to abbreviate the infinite string '$\langle e \rangle, \langle e, e \rangle, \langle e, e, e \rangle, \ldots$' and similarly for other types. So $[\langle e* \rangle]$ is the cumulative type including every (first-order) predicate, of whatever adicity. Since cumulative types are types, the recursive clause for functional types now also yields new functional types. There is, for example, the functional type $\langle e, [e, \langle e* \rangle] \rangle$ of expressions forming

[22] For the metalanguage of his limit language, Williamson adds only one cumulative type, namely, a cumulative infinite limit type, λ, that comprises exactly the expressions belonging to any finite type (2013, p. 221). For my purposes, it is simpler to use only smaller cumulative types at finite orders. There should be no objection to this. If we can form a cumulative infinite limit type like λ and use expressions of types derived from it, surely we can also form a cumulative type comprising, say, only names and first-order predicates, and use expressions of types derived from it.

sentences when combined with a name as their first argument and an expression of type $[e, \langle e* \rangle]$ as their second argument. Since there are both names and predicates among the latter, such expressions accept, in the same argument place, both names and predicates.

We use this system of types to describe, in the first step, an ordinary, non-cumulative second-order language $\mathcal{L}_2$. It includes, for each type τ among e, $\langle e* \rangle$, a countably infinite stock of constants and variables. We write the constants using lower-case letters from the beginning of the alphabet, augmented with subscripts as required, and marking their type by a superscript. We adopt the same convention for variables, except for choosing letters from the end of the alphabet. $\mathcal{L}_2$ also includes the identity predicate '$=^{\langle e,\, e \rangle}$', the usual quantifier symbols, connectives, and parentheses. Type-superscripts may be omitted if there is no risk of ambiguity, and parentheses may be omitted or added according as readability is improved.

$\mathcal{L}_2$ does not include any cumulative devices, so we may call it a *pure* second-order language. In all relevant ways, it is exactly the kind of language the higher-orderist needs to formulate his generalized semantics for the first-order language $\mathcal{L}_1$. If we extend it by copies '$\forall_o x^e$' and '$\exists_o x^e$' of the quantifiers in $\mathcal{L}_1$, we can formulate this version of (GA_1):

$(GA_1{}^*)$ $\forall y^e \, \exists_o x^e \; x^e = y^e$

Given the requisite amount of syntax and set theory, we can go on to formulate in $\mathcal{L}_2$ a higher-orderist generalized semantics for $\mathcal{L}_1$ that includes $(GA_1{}^*)$.

I shall now argue that, independently of the admissibility of cumulative expressions, (DIST) should be rejected. The reason is that although the identity predicate '$=$' itself does not apply at the level of second- and higher-order quantification, higher-order analogues of that predicate do apply. Using these higher-order analogues of the identity predicate, we can construct higher-order counterparts of $(\neq_1)$ and thereby obtain higher-order analogues of $(MORE_1)$. I maintain that these are similar enough to $(MORE_1)$ that they should be taken to contradict the spirit of Absolutism.

We wish to extend $\mathcal{L}_2$ by non-cumulative higher-order cousins of the identity predicate connecting predicates and predicate variables of a given adicity. These expressions should then belong to the types $\langle \langle e \rangle, \langle e \rangle \rangle$, $\langle \langle e, e \rangle, \langle e, e \rangle \rangle$, etc. We shall officially write them as '$=^{\langle \langle e \rangle,\, \langle e \rangle \rangle}$', '$=^{\langle \langle e,\, e \rangle,\, \langle e,\, e \rangle \rangle}$', etc., but often let context fix the type. What should we take these predicates to mean? For present purposes, we can think of them in one of three ways. Firstly, we can take them as primitive expressions subject to inference rules analogous to the standard rules for identity. Where 'a' and 'b' belong to some type τ among $\langle e\rangle$, $\langle e, e\rangle$, …, and $\Phi^{a/b}$ is the result of replacing zero or more occurrences of 'b' in Φ by 'a':

$(=_1)\ \vdash a =^{\langle\tau,\tau\rangle} a$
$(=_2)\ \vdash a =^{\langle\tau,\tau\rangle} b \rightarrow (\Phi \rightarrow \Phi^{a/b})$

Secondly, we can extend $\mathcal{L}_2$ by third-order variables of type $\langle\langle e\rangle\rangle$ ($\langle\langle e, e\rangle\rangle, \ldots$) and quantifiers binding them. The expression '$x^{\langle e\rangle} = y^{\langle e\rangle}$' can then be taken to abbreviate the corresponding indiscernibility condition: $\forall z^{\langle\langle e\rangle\rangle}(z^{\langle\langle e\rangle\rangle}(x^{\langle e\rangle}) \leftrightarrow z^{\langle\langle e\rangle\rangle}(y^{\langle e\rangle}))$, and similarly for the other types. Thirdly, since $\mathcal{L}_2$ is an extensional language, co-extensiveness implies indiscernibility, so we can take the predicates simply to abbreviate the relevant co-extensiveness condition, so that, for example, '$x^{\langle e\rangle} = y^{\langle e\rangle}$' abbreviates '$\forall z^e(x^{\langle e\rangle}(z^e) \leftrightarrow y^{\langle e\rangle}(z^e))$'. Under both the second and third option, the rules $(=_1)$ and $(=_2)$ are derivable.

Now suppose we include in $\mathcal{L}_2$ a copy '$\forall_o x^{\langle e\rangle}$' of a monadic second-order quantifier of some second-order language. We may wonder whether, from the perspective of $\mathcal{L}_2$, that quantifier expresses an unrestricted, absolute form of monadic second-order generality. Roughly speaking, that is, we may wonder whether there is something in the range of $\mathcal{L}_2$'s monadic second-order quantifier that is distinct from anything in the range of '$\forall_o x^{\langle e\rangle}$'. The informal talk of being distinct from anything in that range here can be regimented by means of '$=^{\langle\langle e\rangle,\langle e\rangle\rangle}$':

$(\neq_2)\ \forall_o x^{\langle e\rangle} \ldots \neq^{\langle\langle e\rangle,\langle e\rangle\rangle} x^{\langle e\rangle}$

By putting a monadic second-order variable in the empty argument place and binding it with an existential quantifier, we obtain

$(\mathrm{MORE}_2)\ \exists y^{\langle e\rangle}\ \forall_o x^{\langle e\rangle}\ y^{\langle e\rangle} \neq^{\langle\langle e\rangle,\langle e\rangle\rangle} x^{\langle e\rangle}$

It seems to me that it would be very implausible to discount this claim as irrelevant to the spirit of Absolutism on the grounds that it does not operate with the notion of *distinctness* as it occurs in (GA_1) and (MORE_1). It seems much more plausible to interpret this sentence as saying, in a sense that is relevant to the spirit of Absolutism, that the quantifier '$\forall_o x^{\langle e\rangle}$' expresses only a restricted form of (monadic, second-order) generality. If so, then (DIST) should be rejected.

Analogues of the ordinary identity predicate allow for the formulation of analogues of (MORE_1) that contradict the spirit of Absolutism.[23]

We may thus turn to the second step in my argument. I shall argue that once we are allowed to use cumulative resources, we can formulate another variant on ($\neq_1$), employing a cumulative higher-order counterpart of the identity predicate. We can then use that expression to formulate a cumulative higher-order version of (MORE_1) as our regimentation of (MORE). Just as we took (MORE_2) to contradict the spirit of Absolutism, it seems to me, we should take that regimentation to do so as well. And since it turns out to be derivable in the semantics of the cumulative higher-orderist, the introduction of cumulative resources into higher-orderism is in this way seen to undermine the spirit of Absolutism.

We form the cumulative second-order language $\mathcal{L}_{2\equiv}$ by extending $\mathcal{L}_2$ with an identity-like symbol of type $\langle [e, \langle e* \rangle], [e, \langle e* \rangle] \rangle$, which we shall write as '≡' to make it easier to distinguish from the previous ones. In contrast to these, '≡' is *cumulative* in both its argument places: it accepts both singular terms and predicates of any adicity as arguments.

What should we take '≡' to mean? For present purposes, we may either take it as a new primitive expression, subject to certain inference rules, or we may take it as being given an explicit definition. If we introduce '≡' as a primitive, it should be governed by whatever rules render it as identity-like as possible, consistent with its non-standard syntax. It seems clear, then, that it should satisfy at least the following two rules. Where '*a*' and '*b*' belong to $[e, \langle e* \rangle]$:

($\equiv_1$) $\vdash a \equiv a$
($\equiv_2$) $\vdash a \equiv b \rightarrow (\Phi \rightarrow \Phi^{a/b})$

(When '*a*' and '*b*' belong to different types, substituting one for the other sometimes produces ill-formed results. The rule ($\equiv_2$) is therefore to be understood as restricted to well-formed instances.[24])

[23] Could a higher-orderist dig their heels in and simply insist that (MORE_2) and its ilk are irrelevant to Absolutism? Although such a view would seem very unattractive, I know of no reason to think that it would have to somehow turn out to be internally incoherent. However, I think that Williamson could not happily resort to such a position, since he himself exploits the analogy between the identity predicate and its higher-order counterparts in a way similar to how I have just used it (e.g. Williamson 2013, pp. 263ff.).

[24] Absent contexts other than '≡' in which expressions of different types can be exchanged without loss of grammaticality, our rules leave open the truth-value of any specific sentences in which '≡' connects expressions of different types. We consider these in more detail in the next section.

If we wish instead to introduce '$\equiv$' by definition, we can extend $\mathcal{L}_{2\equiv}$ by variables of type $\langle[e, \langle e*\rangle]\rangle$ and allow these to be bound by quantifiers. We may then let '$a \equiv b$' simply abbreviate the indiscernibility condition, $\forall z^{\langle[e,\,\langle e*\rangle]\rangle}(z^{\langle[e,\,\langle e*\rangle]\rangle}(a) \leftrightarrow z^{\langle[e,\,\langle e*\rangle]\rangle}(b))$.[25] Given a suitable comprehension scheme for the new quantifiers, the rules ($\equiv_1$) and ($\equiv_2$) are then derivable.[26]

We now use '$\equiv$' to formulate a regimentation of (MORE). Let us include in $\mathcal{L}_{2\equiv}$ copies '$\forall_o x^e$' and '$\exists_o x^e$' of the first-order quantifiers of $\mathcal{L}_1$. We can then formulate a cumulative variant on ($\neq_1$) to express a condition of distinctness from anything in the range of $\mathcal{L}_1$'s quantifier:

($\neq_C$) $\forall_o x^e \ \ldots \ \not\equiv x^e$

Since '$\equiv$' accepts expressions of type $\langle e\rangle$ as arguments, we can now put a monadic second-order variable in its empty argument place in ($\neq_C$), and bind it with an existential quantifier. We then obtain the following regimentation of (MORE):

(MORE$_C$) $\exists y^{\langle e\rangle}\ \forall_o x^e\ y^{\langle e\rangle} \not\equiv x^e$

Given ($\equiv_1$) and ($\equiv_2$), this sentence is derivable by broadly standard Russell-style reasoning in second-order logic.[27] I maintain that, like (MORE$_2$), it is similar enough to (MORE$_1$) to be regarded as contradicting the spirit of Absolutism.

[25] For a related but somewhat more complicated definition of a much more general version of '$\equiv$', see Linnebo and Rayo (2012, §5) and Degen and Johannsen (2000, pp. 149f.).

[26] Absent a meaningful notion of co-extensiveness defined for $[e, \langle e*\rangle]$ and thus names and predicates alike, unlike its non-cumulative cousins, '$\equiv$' cannot be introduced as an abbreviation of a co-extensiveness condition.

[27] *Proof*: We define a first-order predicate '$r^{\langle e\rangle}$' as follows, with 'x^e' ranging over everything$_o$: $r^{\langle e\rangle}(x^e) \leftrightarrow_{\mathrm{df}} \forall x^{\langle e\rangle}(x^{\langle e\rangle} \equiv x^e \rightarrow \neg x^{\langle e\rangle}(x^e))$. Assume for *reductio* that for some$_o$ object r^e, $r^{\langle e\rangle} \equiv r^e$. Now $r^{\langle e\rangle}(r^e) \vee \neg r^{\langle e\rangle}(r^e)$. Assume $r^{\langle e\rangle}(r^e)$. Then by definition, $\forall x^{\langle e\rangle}(x^{\langle e\rangle} \equiv r^e \rightarrow \neg x^{\langle e\rangle}(r^e))$, and so in particular, $r^{\langle e\rangle} \equiv r^e \rightarrow \neg r^{\langle e\rangle}(r^e)$. By assumption, $r^{\langle e\rangle} \equiv r^e$, so $\neg r^{\langle e\rangle}(r^e)$, contradicting our assumption of $r^{\langle e\rangle}(r^e)$. So $\neg r^{\langle e\rangle}(r^e)$. Then by definition, $\neg\forall x^{\langle e\rangle}(x^{\langle e\rangle} \equiv r^e \rightarrow \neg x^{\langle e\rangle}(r^e))$. However, assume $x^{\langle e\rangle} \equiv r^e$. Since $r^{\langle e\rangle} \equiv r^e$, it follows by ($\equiv_1$) and ($\equiv_2$) that $x^{\langle e\rangle} \equiv r^{\langle e\rangle}$. Since $\neg r^{\langle e\rangle}(r^e)$, by another application of ($\equiv_2$), $\neg x^{\langle e\rangle}(r^e)$. So $x^{\langle e\rangle} \equiv r^e \rightarrow \neg x^{\langle e\rangle}(r^e)$. Since $x^{\langle e\rangle}$ was arbitrary, $\forall x^{\langle e\rangle}(x^{\langle e\rangle} \equiv r^e \rightarrow \neg x^{\langle e\rangle}(r^e))$. But then by definition, $r^{\langle e\rangle}(r^e)$. Contradiction. So $r^{\langle e\rangle} \not\equiv r^e$, and since r^e was arbitrary, $\forall_o x^e\, r^{\langle e\rangle} \not\equiv x^e$. (MORE$_C$) follows by existential generalization on $r^{\langle e\rangle}$.

The proof assumes that the comprehension scheme for the second-order quantifiers of $\mathcal{L}_{2\equiv}$ allows impredicative instances, including our new '$\equiv$'. Specifically, the legitimacy of the definition of '$r^{\langle e\rangle}$' and/or the subsequent existential generalization on it depend in effect on this impredicative instance of the comprehension schema: $\exists y^{\langle e\rangle}\forall x^e(y^{\langle e\rangle}(x^e) \leftrightarrow \forall z^{\langle e\rangle}(z^{\langle e\rangle} \equiv x^e \rightarrow \neg z^{\langle e\rangle}(x^e)))$. Could the higher-orderist reject such impredicative instances of comprehension? I think not; such a move would appear to undermine the whole motivation for higher-orderism, since the initial Russell-style argument depends on an impredicative instance of (CPI Formal 1).

5. Cumulative identity

The extent to which cumulative higher-orderism's commitment to (MORE_C) is a departure from the spirit of Absolutism depends on the strength of the analogy between '≡' and '=', and thus (MORE_C) and (MORE_1). Given the unfamiliarity of cumulative devices in general and '≡' in particular, it may not be very easy to get a clear sense of the strength of this analogy. This section therefore examines the behaviour of '≡' and its relation to '=' in more detail. I argue that there is no dissimilarity between '≡' and '=' that could undermine the analogy between (MORE_C) and (MORE_1).

We note first that '≡' shares the distinctive structural features of the identity predicate. In particular, it expresses an *equivalence relation* in the sense that in addition to the reflexivity rule ($\equiv_1$), symmetry and transitivity rules are derivable for '*a*', '*b*', '*c*' in $[e, \langle e* \rangle]$:

($\equiv_3$) $\vdash a \equiv b \rightarrow b \equiv a$
($\equiv_4$) $\vdash (a \equiv b \;\&\; b \equiv c) \rightarrow a \equiv c$

The proofs are exactly analogous to the corresponding proofs for '='.

Admittedly, the reflexivity, transitivity and symmetry of the ordinary identity relation can also be expressed in a non-schematic way by means of quantifications like $\forall x\, x = x$ and $\forall x \forall y (x = y \rightarrow y = x)$. The structural features of '≡' are not thus expressible in $\mathcal{L}_{2\equiv}$. The reason is that we do not have any variables that range over an entire cumulative type. So although we can express the reflexivity of ≡ with respect to type *e* by '$\forall x^e\, x^e \equiv x^e$', and with respect to type $\langle e \rangle$ by '$\forall x^{\langle e \rangle}\, x^{\langle e \rangle} \equiv x^{\langle e \rangle}$', and so on, we cannot express by a single sentence the reflexivity of ≡ *tout court*.

We can remove even this disanalogy by moving to a new language $\mathcal{L}_{2C}$ extending $\mathcal{L}_{2\equiv}$ by what I shall call *semantically cumulative* variables. These are variables that belong to a cumulative type $[t_1, t_2, \ldots]$ without belonging to any of the accumulated types $t_1, t_2, \ldots$ Roughly speaking, they are intended to range over the entirety of values of the variables from the accumulated types. We may use underlined lowercase letters from the end of the alphabet for these variables, and express, for example, the reflexivity of ≡ in $\mathcal{L}_{2C}$ as follows:[28]

[28] Although the introduction of such variables could perhaps in principle be rejected by a cumulative higher-orderist, it is important to see how natural their introduction is once syntactically cumulative expressions have been introduced. For, absent semantically cumulative variables, we have cumulative predicates with application conditions defined for a range of items that cannot be swept out by a single variable. It seems more than odd to think that it should be impossible to add variables that can take values from the entire application range of

($\equiv_5$) $\forall \underline{x}^{[e,\langle e*\rangle]}\ \underline{x}^{[e,\langle e*\rangle]} \equiv \underline{x}^{[e,\langle e*\rangle]}$

With respect to their structural features, then, '≡' and '=' seem exactly analogous.

Moreover, whenever '≡' connects two singular terms, the resulting sentence is true iff the corresponding sentence with '=' is true. More generally, any given sentence in which '≡' connects two expressions of the same non-cumulative type is true just in case the result of replacing '≡' with a suitable non-cumulative identity-like predicate is true. That is, for τ among e, $\langle e*\rangle$:

($\equiv_6$) $\vdash a^\tau \equiv b^\tau \leftrightarrow a^\tau =^{\langle\tau,\tau\rangle} b^\tau$

Let us then turn to the somewhat stranger contexts of '≡' in which it connects expressions of different types; call them *cross-type identifications*. Our derivation of (MORE$_C$) has shown that in conjunction with the rest of the higher-orderist's logic, ($\equiv_1$) and ($\equiv_2$) already have substantive *general* implications concerning cross-type identifications: roughly speaking, at least one item in type $\langle e\rangle$ cannot be identified with any item in type e. The truth-values of all *specific* cross-type identifications, however, are left open by our theory. That is, where τ and σ are different types among e, $\langle e*\rangle$, no sentence of the form $a^\tau \equiv b^\sigma$ is either derivable or refutable from our rules.[29] This is exactly parallel to the situation for the logic of '=': where a and b are distinct singular terms, $a = b$ is neither derivable nor refutable from the logical rules alone. Of course, if we are given some sentences $\Phi(a)$, $\neg\Phi(b)$ as premisses, we can infer $a \neq b$ from them. But in just the same way, for a and b in $[e, \langle e*\rangle]$, and given sentences $\Phi(a)$, $\neg\Phi(b)$ as premisses, we may also infer $a \not\equiv b$ from them. So far, then, no relevant disanalogy between '=' and '≡' has emerged.

the predicates. I shall henceforth assume that semantically cumulative variables are no more problematic than syntactically cumulative expressions.

[29] It is clear that our rules do not permit the derivation of any theorem of the form $a \equiv b$ where a and b are distinct expressions. *A fortiori*, they do not permit the derivation of any such theorem where a and b belong to different types. It may not be as obvious that ($\equiv_2$) does not, as it stands, allow the derivation of any negation of a cross-type identification. But note that so far, '≡' is the *only* cumulative predicate in our language. So any well-formed cross-type instance of ($\equiv_2$) will be a formula like '$a \equiv b \rightarrow (b \equiv c \rightarrow a \equiv c)$', where the substitution in the consequent occurs in the scope of '≡'. As a result, we could only obtain the negation of a cross-type identity from such a premiss given another negation of a cross-type identity to start with. For essentially the same reason, we also cannot derive negations of cross-type identifications appealing to the proposed indiscernibility definition of '≡'.

Perhaps one might be tempted to argue for such a disanalogy along the following lines.[30] A cumulative counterpart of the identity predicate like '≡' is most naturally thought of on the model of the *disjunction* of the relevant non-cumulative identity-like predicates. Now, when we think of '≡' in this way, then it *trivially* produces a false sentence whenever it is fed expressions of different types as arguments. But that makes the crucial claims of distinctness from which (MORE_C) is inferred importantly dissimilar to the claims of distinctness from which (MORE_1) is inferred. For the latter claims of distinctness are not trivial in this way. And since (MORE_C) is now seen to be merely an immediate consequence of a perfectly trivial claim, it would be implausible to consider it as contradicting the spirit of Absolutism. For surely it would be disingenuous to interpret the spirit of Absolutism in such a way that it is straightforwardly inconsistent with a mere triviality.

My response is that the envisaged disjunctive understanding of '≡' is incompatible with how I have introduced the expression, and that an understanding that is in line with how I have introduced '≡' does not trivialize cross-type identifications in the way described. I suggested two alternative ways to introduce '≡'. The first is to take it as primitive, subject to the inference rules ($\equiv_1$) and ($\equiv_2$). These rules do not trivialize cross-type identifications, but tie them to predications involving the relevant terms. As a result, there are non-trivial considerations that can be brought to bear on the question of the truth or falsity of a given cross-type identification.

Consider first the question of the identity between objects *a* and *b*. In order to decide the question, we may ask whether there is a predicate that, under a uniform interpretation, is defined for both *a* and *b*, and applies to *a* but not to *b*. Since the cumulative higher-orderist allows cumulative predicates, a counterpart of that consideration applies to questions of cross-type identity. For instance, to decide whether $a^e \equiv b^{\langle e \rangle}$, we may ask whether there is a predicate of type $\langle [e, \langle e \rangle] \rangle$—whose application conditions are accordingly defined for both a^e and $b^{\langle e \rangle}$—which, under a uniform interpretation, applies to a^e but not to $b^{\langle e \rangle}$.

The second way of introducing '≡' I offered is through an explicit definition in terms of the corresponding indiscernibility condition. This definition does not trivialize cross-type identifications either, but ties them to the existence of discriminating properties. To

[30] Thanks here to Nick Haverkamp.

decide whether a given cross-type identification is true, we may ask whether there is a property of type $\langle[e, \langle e\rangle]\rangle$—whose exemplification conditions are accordingly defined for both a^e and $b^{\langle e\rangle}$—which a^e has but $b^{\langle e\rangle}$ does not. In the same way, the existence of discriminating properties bears on questions of first-order identity: we ask whether there is a property whose exemplification conditions are defined for both a and b and which a has and b does not.

Finally, it is not obvious whether some cross-type identifications might not be true.[31] Consider some way for a thing to be, say, wise, and the corresponding property, wisdom, conceived of as a bona fide object, that is, something properly designated in a formal language by a singular term. It is not absurd to think that if '≡' is to be the closest thing to the identity predicate that is defined for a range comprising both objects and ways for objects to be, then it should be the case that wisdom ≡ is wise.[32] Such a view might allow a cumulative higher-orderist to hold that most of the ways for things to be recognized in $\mathcal{L}_{2\equiv}$ are 'identical' to things in the range of $\mathcal{L}_1$'s quantifier—the exception being the Russell-style ways for things to be. Indeed, a non-classical cumulative higher-orderist who endorses a naive comprehension schema for properties at the cost of some truths of classical logic could in this way argue for the negation of (MORE$_C$):

(GA$_C$) $\forall \underline{y}^{[e, \langle e*\rangle]}\, \exists_o x\; \underline{y}^{[e, \langle e*\rangle]} \equiv x$

Whatever the overall merits or demerits of such a view, it would thereby underwrite a stronger form of Absolutism than is consistent with a classical cumulative higher-orderism such as Williamson's.

I conclude that '≡' is not problematically dissimilar to '=', and that the analogy between (MORE$_1$) and (MORE$_C$) is accordingly strong enough that (MORE$_C$) should be taken to contradict the spirit of Absolutism.

[31] It should be noted that the option of holding things of different types in general distinct is plausible at most for what we may call *pure* types, that is, types that are neither cumulative nor derived from cumulative types. If we allow semantically cumulative variables, we of course obtain some examples of true cross-type identifications. In particular, we should have the result that, for example, $\forall x^e\, \exists \underline{y}^{[e, \langle e*\rangle]}\, x^e \equiv \underline{y}^{[e, \langle e*\rangle]}$. Moreover, for higher-level analogues of '≡' which also apply to expressions of functional types derived from cumulative types, we shall also have true cross-type identifications where neither argument is of a cumulative type. For example, for the counterpart of '≡' in type $\langle[\langle e\rangle, \langle[e, \langle e\rangle]\rangle], [\langle e\rangle, \langle[e, \langle e\rangle]\rangle]\rangle$ we should have that $\forall x^{\langle e\rangle}\, \exists y^{\langle[e, \langle e\rangle]\rangle}\, x^{\langle e\rangle} \equiv y^{\langle[e, \langle e\rangle]\rangle}$.

[32] The details of how nominalizations behave in natural language may lend further support to this idea; see Liebesman (2015, esp. §3).

6. Expressive difficulties

One of Williamson's arguments against Relativism is related to the peculiarity we noted in §2 concerning the most natural ways an Absolutist might attempt to express their view. Prima facie, the obvious way to do that is to utter a sentence like 'It is possible to quantify over absolutely everything'. Now if Absolutism is true, then we can interpret the Absolutist's use of the quantifier phrase 'absolutely everything' in the absolutist way they intend, and under such an interpretation the Absolutist has said something relevant and true. But if Absolutism is false, then we cannot so interpret the utterance, and our Absolutist, far from making a relevant claim that is unfortunately not true, has made a true but irrelevant claim. While that is perhaps a somewhat strange dialectical situation, it seems that at least by their own lights, the Absolutist can express their intended thesis.

Williamson (2003, §5) suggests that the Relativist is in a considerably worse position. The idea is this. Suppose the Relativist attempts to state his view by uttering the negation of the Absolutist's sentence, that is, 'It is impossible to quantify over absolutely everything'. That sentence can express what the Relativist intends it to express only on the assumption that Relativism is false. And in that case, what the sentence expresses is also false, as one would have hoped. But if Relativism is true, the sentence simply says something unintended. Whatever restricted domain the quantifier 'absolutely everything' is interpreted as ranging over, the sentence then says that quantification restricted in that way is impossible. But that is not something the Relativist wants to proclaim to be impossible, and rightly so, for it evidently is not impossible. It seems that if the Absolutists cannot help but say something true, but at least say what they want if their view is right, the Relativists cannot help but say something false, and say what they want only if their view is wrong.

Of course, even if successful, this argument does not show that standard Relativism is incoherent. It shows only that a natural first idea for stating the view is incoherent. Relativists might respond in two ways. They might simply resist the urge to produce a general claim supposed to capture their view, limiting themselves to claims like

(MORE$_1$) $\exists y \forall_o x\ x \neq y$

describing individual languages like $\mathcal{L}_1$ as expressing only a restricted form of generality. Although these claims seem to instantiate a common pattern, crying out for generalization, Relativists might simply reject any suggestion that there is a true generalization of which all these claims are mere special cases.[33] Admittedly, while internally coherent, this kind of quietist position may be less than fully satisfactory. Alternatively, Relativists can try to devise some other way of formulating a general thesis that can capture their view without collapsing into incoherence.[34]

Whether or not Relativists can find a convincing solution to the problem, it appears that cumulative higher-orderists face exactly the same sort of difficulty. Their proposal requires that from any legitimate language, we can move to a metalanguage that includes quantifiers of an order higher than any order of quantification found in the object language. So like Relativists, they seem committed to a kind of inexhausibility thesis; only the higher-orderist's thesis concerns the entire hierarchy of orders of quantification. The most natural way to attempt to express the view is by uttering a sentence like 'It is impossible to quantify over absolutely everything in the entire hierarchy of higher and higher orders of quantification'. Clearly, this sentence is no better off than the Relativist's doomed 'It is impossible to quantify over absolutely everything'.

Of course, this does not show that cumulative higher-orderism is incoherent. It shows only that a natural first idea for stating the relativist component of the view is incoherent. Like Relativists, higher-orderists might respond in two ways. They might simply resist the urge to produce a general claim supposed to capture the relativist element of their view, limiting themselves to claims like

(MORE$_C$) $\exists y^{\langle e\rangle} \forall_o x^e \; y^{\langle e\rangle} \not\equiv x^e$

describing individual languages like $\mathcal{L}_1$ as expressing only a limited form of generality. Although these claims seem to instantiate a common pattern, crying out for generalization, higher-orderists might simply reject any suggestion that there is a true generalization of which all these claims are mere special cases. Again, the view seems to be internally coherent, but to the same extent as its Relativist cousin, it also seems less than fully satisfactory. Alternatively, higher-orderists can try to devise some other way of formulating a general thesis capturing their intended form of relativism. It is not obvious how they might do so, but if they manage it, that would answer the objection just presented. However, it is to be expected

[33] Button (2010) puts forth a view like this.

[34] See Fine (2006) for one proposal.

(though I cannot prove) that if any such means were to be found, it would also give the standard Relativist a way of formulating their thesis without falling into incoherence.

I conclude that, as far as adequate formulation of their overall view is concerned, standard Relativists and cumulative higher-orderists are in a very similar situation.[35]

7. Interpretative limitations

Williamson (2003, pp. 415ff., 435) also criticizes Relativism for painting an unattractive picture of parts of logical, metaphysical and mathematical discourse. As I mentioned in the introduction, such discourse appears to provide us with numerous examples of utterances intended as absolutely general. As a result, the interpretations of such utterances that the Relativist can offer seem, from a pre-theoretic standpoint, quite weird and implausible. In this section, I investigate whether the interpretations offered by the cumulative higher-orderist should be considered similarly weird and implausible.

Let me first try to bring out as clearly as I can why the Relativist's interpretations seem weird.[36] Consider any of the following sentences:

(6) Everything is self-identical.
(7) If something x is identical to something y, then this is necessarily so.
(8) Everything is necessarily identical to something.
(9) Everything is either abstract or concrete.

[35] The present criticism of *cumulative* higher-orderism is somewhat similar to criticisms advanced in Linnebo (2006, §4) against *non-cumulative* higher-orderism. Linnebo argues that the higher-orderist is committed to claims such as that there are infinitely many different kinds of semantic value, but cannot by his own lights give proper expression to these claims. He maintains that this limitation is similar to that afflicting Generality Relativism. In an earlier paper (Krämer 2014), I have offered a response to Linnebo's arguments on behalf of the non-cumulative higher-orderist; unfortunately, it does not carry over to protect cumulative higher-orderism against the above argument.

[36] Although I myself agree that they do seem weird, I have encountered some resistance to this view in discussion. So it may be worth stressing two points. Firstly, it is important to guard against a misunderstanding. As I see it, the issue concerns the plausibility of certain interpretations from a *pre-theoretic* standpoint. Among other things, this means that 'interpretation' here must not be understood in any theoretically loaded sense like *assignment of semantic values*, or some such, but is to be taken in an intuitive, pre-theoretic sense. Secondly, for the purposes of my argument, it does not matter much whether the following considerations do show that the Relativist's interpretations should seem implausible. For my claim will simply be that if they do, then similar considerations establish an analogous, if less drastic, verdict on the cumulative higher-orderist's interpretations.

From the point of view of a Relativist's metalanguage and semantics, utterances of any of (6)–(9) can be interpreted only as restricted to some less than all-inclusive domain *D*. In typical cases, any such interpretation seems to simply be a *mis*interpretation. Consider the formalization

(10) $\forall x\ (Ax \vee Cx)$

of (9) in a formal first-order language $\mathcal{L}_1$. Assume we have formulated a Relativist semantics in a first-order metalanguage $\mathcal{L}_2$, which we may suppose includes the predicates 'is abstract' and 'is concrete' with their usual meanings. Now consider the following two sentences of $\mathcal{L}_2$:

(11) $\forall x$ (x is concrete $\vee$ x is abstract)
(12) $\forall x \in D$ (x is concrete$_D$ $\vee$ x is abstract$_D$)

where *D* is the domain with respect to which we are interpreting $\mathcal{L}_1$, and 'abstract$_D$' and 'concrete$_D$' mean the same as their subscript-free counterparts except that their application conditions are defined only over *D*. Evidently, the generalization expressed by (12), as well as every generalization expressed by a version of (12) in which some other set is referred to in place of *D*, bears an interesting relationship to the generalization expressed by (11). Roughly speaking, (12) is obtained from (11) by replacing the quantifier '$\forall x$' with a proper restriction of it, and replacing the predicates with counterparts having accordingly restricted application conditions. I shall say the generalization expressed by (12) is a *mere restriction* of that expressed by (11), and conversely that the latter is a *mere expansion* of the former.

The important point is that the claim expressed by (12), because it is a mere restriction of that expressed by (11), seems a strange target for metaphysical inquiry, and relatedly, a strange claim to put forth as the upshot of a metaphysical investigation. The claim expressed by (11) seems a much more interesting and natural claim to focus on. At least at a first, and again, perhaps somewhat naive glance, it therefore seems implausible to interpret a metaphysician using the $\mathcal{L}_1$-sentence (10) to put forth (part of) his metaphysical theory as having endorsed the claim expressed by (12), not (11). Of course, if we try to specify, in a metalanguage for our metalanguage, an intended interpretation of (11), it too will turn out to express a mere restriction of a yet more encompassing counterpart. In this way, saying that everything is abstract or concrete becomes something of a metaphysisyphean task.

To illustrate the point another way, imagine a necessitist like Williamson uttering (8). In effect, the Relativist would interpret him as saying that everything *except perhaps some things outside D* is necessarily (identical to) something. This seems very strange. From a naive point of view, one expects the Relativist to then ask our necessitist what his view is with respect to the things outside D. The kind of conversation that would then unfold seems not worth having. N: 'I think that *everything*, in or outside D, is necessarily something.' R: 'Oh, I see, you think that everything in your domain of discourse D' is necessarily something. But what about the things outside D'?' N: 'Yes, them too. Indeed, *everything*, and so in particular everything outside D', is necessarily something.' R: 'I see, so everything in D'' …'

Thankfully, disputes in logic and metaphysics rarely take this shape. But it is not obvious why interpreting speakers in the Relativist's way should not invite this kind of tiring sequence of questions and responses. Relatedly, it is not obvious why the Relativist's interpretations of the pertinent utterances are not cases of strange misinterpretation. So at first glance, it seems that Relativism paints a somewhat disconcerting picture of significant parts of logico-metaphysical inquiry.

However, it seems to me that the same kind of problem also arises for the cumulative higher-orderist, although only in a somewhat less dramatic form. Specifically, for some universal generalizations of $\mathcal{L}_1$, it seems that, from the perspective of the cumulative metalanguage $\mathcal{L}_{2C}$, they express mere restrictions of more encompassing claims. This is so in particular for $\mathcal{L}_1$-generalizations involving only logical vocabulary.

Consider the $\mathcal{L}_1$-sentence

(13) $\forall x\ x = x$

On the cumulative higher-orderist's semantics given in $\mathcal{L}_{2C}$, it expresses the claim that

(14) $\forall x^e\ x^e = x^e$

Now, as we have seen, using cumulative resources, we can express a version of the notion of identity that is defined both for objects and for ways for objects to be. That notion of $\equiv$ seems to relate to the ordinary notion of identity in much the same way as the notion of abstractness relates to that expressed by 'abstract$_D$' above. The ordinary identity predicate '$=$' means the same as '$\equiv$', except that its application conditions are defined only for objects.

Moreover, the semantically cumulative quantifier '$\forall \underline{x}^{[e,\ \langle e*\rangle]}$' of $\mathcal{L}_{2C}$ seems to relate to '$\forall x^e$' as expansion to restriction in much the same way as '$\forall x$' above does to '$\forall x \in D$'. Indeed, from the perspective of $\mathcal{L}_{2C}$, we could also say that (13) expresses the claim that[37]

(15) $\forall \underline{x}\ (\mathrm{OBJ}(\underline{x}) \rightarrow \underline{x} \equiv_{\mathrm{OBJ}} \underline{x})$

where '$\mathrm{OBJ}(\underline{x})$' is defined by '$\exists y^e\ \underline{x} \equiv y$', and thus applies to all and only objects, and '$\underline{x} \equiv_{\mathrm{OBJ}} \underline{y}$' is defined by '$\underline{x} \equiv \underline{y}$' except that it is defined only for objects.

In this way, we can formulate mere expansions for any given universal generalization in $\mathcal{L}_1$ as long as it contains only first-order predicates that have natural higher-order analogues. Clearly, this is not the case for *every* generalization. For example, it does not seem as though the abstract–concrete distinction has any obvious counterparts at higher orders of quantification. Nevertheless, it holds for quite a few interesting and contentious logico-metaphysical theses, such as the necessity of identity and existence. More cases show up when we consider a second-order language as object language. They include, for example, the claim that nothing has contradictory properties, which we can formulate as $\forall x^e \neg\exists x^{\langle e\rangle}(x^{\langle e\rangle}(x^e)\ \&\ \neg x^{\langle e\rangle}(x^e))$, or the conclusion of the Russell Paradox, which may be expressed as $\neg\exists z^{\langle e,\, e\rangle}\forall x^{\langle e\rangle}\exists x^e \forall y^e\ (x^{\langle e\rangle}(y^e) \leftrightarrow z^{\langle e,\, e\rangle}(x^e, y^e))$.[38]

It would be an exaggeration to claim in conclusion that with respect to the difficulty of interpreting apparently absolutely general discourse in logic, mathematics or metaphysics, a cumulative higher-orderist is in as worrying a position as the Relativist. The cases in which they may be thought to interpret speakers in an implausibly restricted fashion are less widespread than those arising for Relativism, and it is not obvious that the restricted interpretations are as implausible-looking as those given by the Relativist. Nevertheless, the difference appears to be one of degree, and a smaller one than one might have expected.[39]

[37] For readability, I have here dropped the type-superscript $[e, \langle e* \rangle]$ from the underlined, cumulative variables.

[38] A detailed assessment and comparison of the seriousness of the problem for Relativism and cumulative higher-orderism is beyond the scope of this paper, but the following consideration may indicate a reasonable way to start. Suppose someone makes an utterance about a topic *t*, and consider the distinctive 'domain-expanding' moves then canvassed by Relativism and cumulative higher-orderism, respectively. We may then ask whether the additional expressive resources available after the move enable us to exclude ways for the world to be *with respect to t* that we could not previously exclude. If it turns out that the range of cases in which this is so is significantly smaller for cumulative higher-orderism than it is for Relativism, this would seem to speak in favour of the former view. (Thanks to an anonymous referee for suggesting this way of viewing the matter.)

[39] There may be further examples of problems arising for higher-orderism which seem akin to problems of Generality Relativism. In particular, Linnebo and Rayo (2012, §7) argue that

8. Conclusion

The attempt to develop a generalized semantic theory in a first-order language for an object language that expresses absolute generality runs into a version of Russell's Paradox. Standard Relativism concludes that absolute generality is impossible. Williamson (2003) criticizes that view on a number of grounds, such as that it is by its own lights not properly expressible, and that it yields an unsatisfactory picture of parts of logical, metaphysical and mathematical inquiry. To avoid Relativism, Williamson appeals to a hierarchy of higher and higher orders of quantification, which, in his recent book, he extends to transfinite and cumulative orders of quantification. I have argued that the emerging view has implications strongly reminiscent of the distinctive commitments of standard Relativism, and thus gives up on much of the spirit of Absolutism. Moreover, versions of the two criticisms of Relativism just mentioned also apply on the cumulative higher-order view.

Cumulative higher-orderism is not thereby shown to be wrong, of course. (Indeed, I incline to think it is correct.) What I hope to have shown, however, is that it does not let us be as much of a Generality Absolutist as we might have thought and hoped. Even according to cumulative higher-orderism, there is not just everything. There is everything, and then some.[40]

higher-orderists are committed to a strictly open-ended type-theoretic hierarchy, and that this appears to face problems similar to the Relativist's open-ended ontological hierarchy. It may be worth stressing that the claims and arguments of the preceding two sections are independent of the issue of open-endedness. They turn crucially on the *cumulative* nature of the type-theoretic hierarchy rather than its size. Specifically, it is thanks to cumulative types that a regimentation of the distinctive relativist claim (MORE) can be formulated and proved in cumulative higher-order logic, and that what appear to be mere expansions of lower-order generalizations become expressible at higher cumulative orders. (Thanks to an anonymous referee for pressing me on this.)

[40] Earlier versions of this paper were presented at the Hamburg Summer School 2012, at which Timothy Williamson presented his book, *Modal Logic as Metaphysics*, and at a workshop on Typed and Untyped Approaches in Semantics in Oslo. I am grateful to the audiences at both occasions for helpful discussion, and especially to Timothy Williamson for extremely helpful and encouraging comments. Many thanks are also due to Nick Haverkamp, Stefan Roski, Benjamin Schnieder, Richard Woodward, and all fellow members of the Hamburg *phlox* research group, as well as two referees and the editor of *Mind* for helpful comments and criticism. The research for this paper was carried out while I was employed in the DFG/ANR funded project Nominalizations, and I gratefully acknowledge that support.

References

Boolos, George 1984: 'To Be is to be a Value of a Variable (or to be Some Values of Some Variables)'. *Journal of Philosophy*, 81, pp. 430–49.
——1985: 'Nominalist Platonism'. *Philosophical Review*, 94, pp. 327–44.
Button, Tim 2010: 'Dadaism: Restrictivism as Militant Quietism'. *Proceedings of the Aristotelian Society*, 110, pp. 387–98.
Degen, Wolfgang, and Jan Johannsen 2000: 'Cumulative Higher-Order Logic as a Foundation for Set Theory'. *Mathematical Logic Quarterly*, 46, pp. 147–70.
Fine, Kit 2006: 'Relatively Unrestricted Quantification'. In Rayo and Uzquiano 2006b, pp. 20–44.
Florio, Salvatore 2014: 'Unrestricted Quantification'. *Philosophy Compass*, 9(7), pp. 441–54.
Krämer, Stephan 2014: 'Semantic Values in Higher-Order Semantics'. *Philosophical Studies*, 168, pp. 709–24.
Liebesman, David 2015: 'Predication as Ascription'. *Mind*, 124, pp. 517–69.
Linnebo, Øystein 2006: 'Sets, Properties, and Unrestricted Quantification'. In Rayo and Uzquiano 2006b, pp. 149–78.
Linnebo, Øystein and Agustín Rayo 2012: 'Hierarchies Ontological and Ideological'. *Mind*, 121, pp. 269–308.
Prior A. N. 1971: *Objects of Thought*. Edited by P. T. Geach and A. J. P. Kenny. Oxford: Clarendon Press.
Rayo, Agustín 2006: 'Beyond Plurals'. In Rayo and Uzquiano 2006b, pp. 220–54.
Rayo, Agustín and Gabriel Uzquiano 1999: 'Toward a Theory of Second-Order Consequence'. *Notre Dame Journal of Formal Logic*, 40(3), pp. 1–11.
——2006b: 'Introduction'. In Rayo and Uzquiano 2006b, pp. 1–20.
——2006a: *Absolute Generality*. Oxford: Oxford University Press.
Rayo, Agustín, and Stephen Yablo 2001: 'Nominalism through De-Nominalization'. *Noûs*, 35(1), pp. 74–92.
Williamson, Timothy 2003: 'Everything'. *Philosophical Perspectives*, 17, pp. 415–65.
——2013: *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

THOUGHT
A JOURNAL OF PHILOSOPHY
THE NORTHERN INSTITUTE OF PHILOSOPHY

Thought ISSN 2161-2234

ORIGINAL ARTICLE

A Simpler Puzzle of Ground

Stephan Krämer

University of Hamburg

Metaphysical grounding is standardly taken to be irreflexive: nothing grounds itself. Kit Fine has presented some puzzles that appear to contradict this principle. I construct a particularly simple variant of those puzzles that is independent of several of the assumptions required by Fine, instead employing quantification into sentence position. Various possible responses to Fine's puzzles thus turn out to apply only in a restricted range of cases.

Keywords grounding; circularity; puzzle; sentential quantification; Fine

DOI:10.1002/tht3.77

In the recent debate on metaphysical grounding and its logic,[1] it has generally been accepted that grounding is irreflexive, i.e., nothing grounds itself, and that true existential quantifications are grounded in their true instances. Call this second principle *EG* (Existential Grounding). Kit Fine (2010) has shown that given a number of plausible auxiliary assumptions, *EG* yields counterexamples to irreflexivity.[2] It turns out that an extremely simple derivation of a counterexample to irreflexivity from *EG* is available if we (i) take grounding to be expressed by an operator on sentences (or lists of sentences), (ii) avail ourselves of (non-substitutional[3]) quantification into sentence position, and (iii) assume that *EG* extends to this kind of quantification. It is safe to assume that

$$(1)\ \exists p\ p.$$

which is a theorem of any standard logic with sentential quantification. Now note that (1) is a true instance of itself, so that by *EG*, we immediately obtain a case of self-grounding (writing '$\prec$' for partial[4] grounding):

$$(2)\ \exists p\ p \prec \exists p\ p.$$

Fine's own arguments differ from this one in that they do not rely on sentential quantification, but instead quantify first-order over facts, sentences, or propositions, and additionally appeal to a number of principles relating grounding to these kinds of entities. A Finean counterpart to the present puzzle might go as follows.[5] In place of (1), we assume that some proposition is true (read the variables as restricted to propositions):

$$(\text{F1})\ \exists x\ x\ \text{ is true.}$$

Correspondence to:: E-mail: stephan.kraemer@uni-hamburg.de

Next, we assume a claim of Propositional Existence: that there is such a thing as the proposition that some proposition is true. Letting '$<p>$' abbreviate 'the proposition that p':

$$\text{(PE) } \exists y\, y = <\exists x\, x \text{ is true}>.$$

Now let 's' abbreviate '$\exists x\, x$ is true', i.e., sentence (F1). Using the suitable instance of a Truth-Introduction principle

$$\text{(TI) If } p \,\&\, \exists x x = <p>, \text{then } <p> \text{ is true.}$$

we infer from (F1) and (PE):

$$\text{(F2) } <s> \text{ is true.}$$

Next, we make use of the following plausible principle, which is often associated with Aristotle (cf. e.g., Schnieder 2006, 35f):

$$\text{(A) If } <p> \text{ is true then } p \prec <p> \text{ is true.}$$

We infer from (F2) and the relevant instance of (A):

$$\text{(F3) } s \prec <s> \text{ is true.}$$

But '$<s>$ is true' is a true instance of '$\exists x\, x$ is true', i.e. 's'. So from (F2) and *EG* we obtain:

$$\text{(F4) } <s> \text{ is true} \prec s.$$

The claims (F3) and (F4) jointly violate the principle that grounding is asymmetric. If we assume moreover the Transitivity of Grounding

$$\text{(TG) If } p \prec q \text{ and } q \prec r, \text{ then } p \prec r.$$

we obtain a violation of irreflexivity.

The only responses to this Finean argument that straightforwardly apply to my version of the puzzle as well are to reject irreflexivity or *EG*. Some possible reasons for denying (F1) may also motivate a rejection of (1), but since both moves appear extremely unattractive I shall set them aside. Since none of the other auxiliary assumptions (PE), (TI), (A), and (TG) are used in my argument, rejecting any of these will not by itself help with the latter.[6]

I shall now briefly comment on the assumptions (i)–(iii), stated at the beginning of the paper, that are required in my, but not Fine's argument. I have nothing interesting to say here about the claim (i) that grounding is adequately expressed by a sentential operator. Like Fine and many others, I accept the claim, but it is not uncontroversial—dissenters include, e.g., Rosen (2009) and Audi (2012).

Re (ii): If one thinks that quantification into sentence position is meaningless one can of course very simply reject the purported counterinstance to irreflexivity on that ground. If one thinks that quantification into sentence position is meaningless unless understood as merely abbreviating first-order quantification over propositions, the argument will turn out to be a notational variant on its Finean counterpart, tacitly relying on the same premises. Both claims are highly controversial, though, so it is interesting to see what happens if they are denied. Moreover, given (i), there are special reasons for friends of grounding not to dismiss quantification into sentence position. For some interesting and important structural theses about grounding—like irreflexivity, transitivity, and well-foundedness—then seem most naturally expressed by means of such quantification. The most natural statement of transitivity, for instance, is as the claim that $\forall p \forall q \forall r\ ((p \prec q \,\&\, q \prec r) \rightarrow p \prec r)$ rather than the schematic (TG) above. For in contrast to the schema, the quantificational claim can be properly embedded, and its import does not inappropriately depend on the linguistic resources available to instantiate it. By way of analogy, expressing the irreflexivity of grounding by saying that every instance of '$p \prec p$' is false seems as unsatisfactory as stating the reflexivity of identity by saying that every instance of '$a = a$' is true, instead of saying that $\forall x\, x = x$.

Re (iii): Here is a natural statement of a rule capturing *EG* for the case of sentential quantification ('α' stands for an arbitrary sentence, '$\phi(\ )$' for a suitable sentential context):

$$\text{(EG-S)}\ \ \text{From}\ \ \phi(\alpha)\,,\text{infer}\ \ \phi(\alpha) \prec \exists p\ \phi(p).$$

The notion of (1) being an instance of itself may perhaps seem weird, so let me try to briefly dispel that impression. The instance of (EG-S) that legitimizes the move from (1) to (2)—i.e.

$$\text{(X)}\quad \text{From}\ \exists p\ p, \text{infer}\ \exists p\ p \prec \exists p\ p.$$

can be obtained by putting '$\exists p\ p$' for α and the null context, i.e. nothing, for ϕ. Note that it seems perfectly appropriate to take (EG-S) as allowing the null instantiation of ϕ, for an analogous reading is required for the corresponding rule of existential generalization—from $\phi(\alpha)$, infer $\exists p\ \phi(p)$—if it is to legitimize the clearly valid inference to '$\exists p\ p$' from an atomic sentence letter as a premise. Moreover, on the standard conception of how to construct an instance of a given quantification, we proceed by deleting the initial quantifier phrase and then systematically replacing, in the remaining expression, the formerly bound variable by an expression of the suitable grammatical category. Deleting the '$\exists p$' from '$\exists p\ p$' yields 'p', and systematically replacing 'p' in 'p' by the sentence '$\exists p\ p$' then yields '$\exists p\ p$', as desired.

So let us look for reasons to reject (EG-S) while still accepting a rule capturing *EG* for the case of first-order quantification. *Simply* rejecting (EG-S) and insisting that *EG* holds for first-order quantification is unsatisfactory, for the motivation usually offered for *EG* is not specific to the case of first-order quantification. The principle is sometimes held to command intuitive support (cf. e.g., Rosen 2009, p. 117); to the extent that it does so, it seems to me, it does so equally for any sort of quantification. More substantively, the principle is sometimes motivated by appeal to the highly plausible principle of disjunctive grounding—if p, then $p \prec (p \vee q)$—and the analogy between existential quantification and disjunction (cf. e.g., Schnieder 2011, p. 460; Fine 2012a, p. 60). This analogy extends to the case of sentential quantification.

Can we give reasons to reject (X) while retaining a *restricted* version of (EG-S)? The only potentially plausible suggestion I can think of is to blame the *impredicativity* exhibited by (X). Thus we might say, firstly, that in the sense that is relevant to *EG*, being an instance of a quantification is not a purely syntactic matter. Rather, the expression generalized upon also has to satisfy a semantic condition: roughly, that of determining, or picking out, a value in the range of the corresponding existential quantifier. Secondly, we say that a sentence that itself contains a given sentential quantifier does not determine a value in the range of that quantifier. A simple implementation of that idea restricts (EG-S) to cases in which α is free of sentential quantifiers; a less restrictive option is to introduce a hierarchy of sentential quantifiers and postulating a version of (EG-S) for each of them, requiring in each case that α contain only quantifiers lower in the hierarchy.

A rejection of impredicative definition is also one of the possible motivations to reject (PE) in Fine's arguments. So in this way the predicativists' response to Fine's puzzles may also generalize to my version. But if my contention is correct and impredicativity provides the only potentially plausible ground for rejecting instance (X) of (EG-S), then there are simply no analogues available to the rejection of (TI), (A), or (TG) above. A wholesale rejection of *EG*, a ban on at least certain sorts of impredicative instances of *EG*, and the admission of counterexamples to irreflexivity then are the only options left.

Acknowledgments

I would like to thank Fabrice Correia, Kit Fine, Stefan Roski, Benjamin Schnieder, Alex Skiles, and an anonymous referee for this journal for helpful feedback and discussion.

Notes

1 I am thinking in particular of Fine (2010, 2012a, 2012b), Schnieder (2011), Correia (2011), and Rosen (2009).

2 Fine's paper also establishes analogous points for a counterpart to *EG* for universal quantifications.

3 This qualification will be taken as understood henceforth.

4 All of the grounding claims to follow are standardly taken to hold for both partial and complete grounding. For simplicity, take '$\prec$' as binary; arguably, for some purposes we need an operator that takes an arbitrary number of sentences in the left-hand argument place (cf. e.g., Fine 2012a, 46f).

5 Fine does not state or discuss this exact argument, which is something of a hybrid of Fine's Particular Argument for Facts and his Universal Argument for Propositions. The latter rests on essentially the same premises as the present one, except that in place of (F1), it assumes that every proposition is either true or false. See Fine (2010, section 5).

6 Correia (forthcoming), in the context of this puzzle, argues that irreflexivity should be given up. Skiles (unpublished manuscript) advocates rejecting (A) as well as my assumption (i) from the beginning of the paper. I should also note that Correia (2011, 7f) formulates a version of the Finean puzzle that, like mine, employs sentential quantification, but also assumes a version of (A) and of (TI).

References

Audi, P. "A Clarification and Defense of the Notion of Grounding," in *Correia/Schnieder*, 2012, 101–21.

Correia, F. "From Grounding to Truth-Making: Some Thoughts," in *Papers dedicated to Kevin Mulligan* (online), edited by A. Reboul, 2011. Available at http://www.philosophie.ch/kevin/festschrift/Correia-paper.pdf

Correia, F. "Logical Grounds." *Review of Symbolic Logic* (forthcoming)

Correia, F. and Schnieder, B. *Metaphysical Grounding*. Cambridge: CUP, 2012.

Fine, K. "Some Puzzles of Ground." *Notre Dame Journal of Formal Logic* **51**.1 (2010): 97–118.

———. "Guide to Ground," in *Metaphysical Grounding*, edited by Correia, F. and Schnieder, B. Cambridge: CUP, 2012a, 37–80.

———. "The Pure Logic of Ground." *Review of Symbolic Logic* **5**.1 (2012b): 1–25.

Rosen, G. "Metaphysical Dependence: Grounding and Reduction," in *Modality: Metaphysics, Logic, and Epistemology*, edited by B. Hale and A. Hoffman. New York: OUP, 2009, 109–36.

Schnieder, B. "Truth-Making *Without* Truth-Makers." *Synthese* **152** (2006): 21–47.

———. "A Logic for "Because"." *Review of Symbolic Logic* **4** (2011): 445–65.

18
PUZZLES

Stephan Krämer

Ground seems to satisfy certain *structural* principles, like irreflexivity, asymmetry, and transitivity (cf. Chapter 17). Ground also seems to satisfy certain principles concerning its interaction with *logical concepts* like disjunction and existential quantification, such as the principle that a true existential quantification is grounded by its true instances (cf. Chapter 14). Finally, ground seems to satisfy some principles linking it to concepts like that of a fact, proposition, or truth, which might be called broadly logical concepts. An example is the principle that if the proposition that *P* is true, then it is true because *P*. (As is common in the literature, I sometimes use "because" to convey metaphysical grounding.) We are faced with a *logical puzzle of ground* when the combination of a number of principles of these sorts, given standard logical inference rules, yields an absurd conclusion.

Kit Fine, in his 2010 paper "Some Puzzles of Ground", first showed that there are logical puzzles of ground. To give you a taste, here is a sketch of one. It is very plausible that some proposition is true. If so, it seems we may infer that the proposition *that some proposition is true* is true. By the principle stated above, this proposition is true *because* some proposition is true. But by the principle that an existential quantification is grounded by any true instance, we *also* have that some proposition is true *because* (among other things) the proposition *that some proposition is true* is true. Combining these two results, we obtain a violation of the asymmetry of ground.

This chapter reviews the variety of logical puzzles of ground that have been identified in the literature, describes the solutions that have been proposed, and indicates what the main challenges are that these solutions face. I begin by introducing relevant notation as well as the key concepts and principles that will subsequently be used in formulating the puzzles before turning to the puzzles themselves. In principle, there is a huge number of different derivations of contradictions from the relevant principles about ground. Many of them are essentially alike, so that any reasonable solution to one will immediately provide a solution to the other. Some of them exhibit more substantial differences, however, and I shall try to describe all the substantially different types of puzzles. I then briefly discuss what desiderata we might impose on adequate solutions to the puzzles before I finally turn to the solutions themselves. Many of these, once developed in detail, involve a fair bit of formal machinery. I shall mainly attempt to convey the basic philosophical ideas underlying and motivating the technical work; readers keen on the details will have to consult the primary texts.

Background

To state the logical puzzles of ground clearly and concisely, we employ the usual symbolism and write "$P \prec Q$" to say that P partially grounds Q. Note that "$\prec$" combines with *sentences*. In this way, the notation fits with the *operator* view of grounding rather than the predicate view (see the editor's introduction for an explanation of the distinction). However, both the puzzles and the solutions to be discussed could be transposed, with minor adjustments, to the predicational setting.

The logical puzzles of ground all invoke certain principles of the so-called *Pure Logic of Ground* (Fine 2012b), comprising purely structural rules for ground. The most important ones are that partial ground is irreflexive, asymmetric, and transitive. We shall formulate these principles in the form of inference rules:

Irreflexivity	$A \prec A$	$/ \bot$
Asymmetry	$A \prec B, B \prec A$	$/ \bot$
Transitivity	$A \prec B, B \prec C$	$/ A \prec C$

(Inference rules are stated in the form "Premise, Premise, . . ./Conclusion", allowing the empty list of premises as a special case. "$\bot$" stands for absurdity.) *Irreflexivity* follows from *Asymmetry*, and *Asymmetry* is derivable given *Irreflexivity* and *Transitivity*, so either of *Irreflexivity* and *Asymmetry* could be dropped as a basic rule. However, since some puzzles require only *Irreflexivity* and some only *Asymmetry*, it is best to state them separately.

Next, several principles from the *Impure Logic of Ground*, linking ground to the logical connectives and quantifiers, play an essential role in the puzzles:

$\exists$-*Grounding*	$B(a)$	$/ B(a) \prec \exists x\, B(x)$
$\forall$-*Grounding*	$\forall x\, B(x)$	$/ B(a) \prec \forall x\, B(x)$
$\vee$-*Grounding*	A	$/ A \prec A \vee B$

Most of the puzzles also invoke some means of talking about either sentences or things suitably related to sentences, such as propositions or facts. Our focus here will be mainly on puzzles that talk about propositions rather than sentences or facts; where the differences between these options matter, this will be made explicit. For brevity, instances of "$\langle P \rangle$" are used as shorthand for corresponding instances of "the proposition that P", and instances of "Tp" for instances of "p is true". We make use of three rules concerning, respectively, the existence of propositions, the truth of a proposition, and the grounds of a proposition's being true:

P-*Existence*		$/ \exists x\, x = \langle A \rangle$
T-*Introduction*	$p = \langle A \rangle, A$	$/ Tp$
T-*Grounding*	$Tp, p = \langle A \rangle$	$/ A \prec Tp$

Finally, we need the standard logical rule of $\exists$-*Elimination*, which says that a conclusion C may be derived from $\exists x\, B(x)$ and some auxiliary assumptions $A_1, \ldots, A_n$ if it may be derived from $B(a)$ and $A_1, \ldots, A_n$, provided that a does not occur in any of $\exists x\, B(x)$ and $A_1, \ldots, A_n$.

The Puzzles

All of the puzzles to be discussed essentially involve an instance of what may be called *circular aboutness*. A simple example of this is a proposition that *quantifies* over a range of things including that proposition itself. Thus, the proposition that there are propositions exhibits circular aboutness (short: is circular). Another class of examples involves *singular self-reference*. Suppose my first utterance this morning was: "The first proposition I express today is interesting". Then the proposition I expressed in making this utterance is also circular. (In principle, the circles of aboutness can be longer, when one proposition is about another, which is in turn about the first, but we shall mainly be concerned with the simple cases.)

In all of the puzzles, the existence of a circular *aboutness* relationship is then exploited to derive, using the principles introduced earlier, the existence of a circular *grounding* relationship. This yields a violation of the principles of *Irreflexivity*, *Transitivity*, and/or *Asymmetry*, which jointly exclude the possibility of a grounding circle. In many cases, the circles are indeed so small as to include only two members, so that *Asymmetry* is directly violated without the help of *Transitivity*. As we shall see, there is even an example of a one-element circle of ground, directly violating *Irreflexivity*.

We begin with a detailed presentation of the puzzle mentioned in the introduction to this article, centering on the proposition that some proposition is true. Both this puzzle and the next are adapted with minor modifications from (Fine 2010: §5). The arguments will be presented in a Lemmon-style notation for derivations in which each line represents one step in the derivation, consisting in, first, the (possibly empty) list of assumptions upon which the derived sentence depends, second, the number of the line, third, the derived sentence, and fourth, the rule in whose application the step in question consists.

	(1)	$\exists y\ y = \langle \exists x\ Tx \rangle$	P-*Existence*
2	(2)	$\exists x\ Tx$	Assumption
3	(3)	$p = \langle \exists x\ Tx \rangle$	Assumption, for ∃-*Elimination*
2, 3	(4)	Tp	3, 2 T-*Introduction*
2, 3	(5)	$\exists x\ Tx \prec Tp$	3, 4 T-*Grounding*
2, 3	(6)	$Tp \prec \exists x\ Tx$	4 ∃-*Grounding*
2, 3	(7)	$\bot$	5, 6 *Asymmetry*
2	(8)	$\bot$	1, 3, 7 ∃-*Elimination*

Two comments. First, beyond the rules introduced in the previous section, the derivation depends solely on the highly plausible assumption (2) that there is at least one truth. In many contexts, even that assumption can be dispensed with, since most logical systems allow the premise-free derivation of some statement, from which (2) may then be derived using P-*Existence* and T-*Introduction*. Second, the proposition $p = \langle \exists x\ Tx \rangle$ whose existence is established in the first step of the derivation is circular, for p is itself within the range of the quantifier $\exists x$ used in expressing p. If $\exists x$ was restricted so that p lies outside its range, we would not be justified in applying ∃-*Grounding* to obtain $Tp \prec \exists x\ Tx$.

A contradiction can also be derived for suitable *universally* rather than existentially quantified circular propositions:

	(1)	$\exists y\ y = \langle \forall x\ (Tx \lor \neg Tx) \rangle$	P-*Existence*
2	(2)	$\forall x\ (Tx \lor \neg Tx)$	Assumption
3	(3)	$p = \langle \forall x\ (Tx \lor \neg Tx) \rangle$	Assumption, for ∃-*Elimination*
2, 3	(4)	Tp	3, 2 T-*Introduction*
2, 3	(5)	$\forall x\ (Tx \lor \neg Tx) \prec Tp$	3, 4 T-*Grounding*
2, 3	(6)	$Tp \prec Tp \lor \neg Tp$	4 ∨-*Grounding*
2	(7)	$Tp \lor \neg Tp \prec \forall x\ (Tx \lor \neg Tx)$	2 ∀-*Grounding*

2, 3	(8)	$Tp \prec \forall x\, (Tx \lor \neg Tx)$	6, 7 *Transitivity*
2, 3	(9)	$\bot$	5, 8 *Asymmetry*
2	(10)	$\bot$	1, 3, 9 ∃-*Elimination*

Beyond the rules of the previous section, this derivation of a contradiction depends only on the assumption that everything is either true or not true. Again, in many logics, we can derive this claim. The first step in the derivation consists in using P-*Existence* to establish the existence of a circular proposition, in this case $\langle \forall x\, (Tx \lor \neg Tx)\rangle$. Similarly as in the first puzzle, it is essential that the proposition be construed as itself within the range of the quantifier used in expressing it; otherwise the application of ∀-*Grounding* would not be justified. And just as before, the contradiction is obtained by establishing a circle of ground. Only in this case, the circle is slightly longer than before. To establish its individual links, we need both ∨-*Grounding* and ∀-*Grounding*, whereas previously we only needed ∃-*Grounding*, and to obtain a violation of *Asymmetry*, we first need to apply *Transitivity*.

Let me now state a puzzle that invokes a proposition employing singular reference to itself, adapted from Fine (2010: 117n15):

	(1)	$\exists y\; y = \langle 0 = 0 \lor Ty\rangle$	P-*Existence*
2	(2)	$p = \langle 0 = 0 \lor Tp\rangle$	Assumption, for ∃-*Elimination*
3	(3)	$0 = 0 \lor Tp$	Assumption
2, 3	(4)	Tp	2, 3 T-*Introduction*
2, 3	(5)	$Tp \prec 0 = 0 \lor Tp$	4 ∨-*Grounding*
2, 3	(6)	$0 = 0 \lor Tp \prec Tp$	4, 2 T-*Grounding*
2, 3	(7)	$\bot$	5, 6 *Asymmetry*
3	(8)	$\bot$	1, 2, 7 ∃-*Elimination*

Beyond the rules of the previous section, this derivation of a contradiction depends only on the classical logical truth that $0 = 0 \lor Tp$. The crucial first step uses P-*Existence* to establish the existence of a proposition that "says of itself" that either it is true or $0 = 0$. As in the first case, the circle of ground that we establish has only two members, so no appeal to *Transitivity* is needed to establish the contradiction. (For further discussion of this puzzle, see (Litland 2015: 488f).)

One can construct variations of these puzzles that do not require P-*Existence*, T-*Introduction*, and T-*Grounding*. First, as already mentioned, in place of propositions, one may appeal to sentences or facts. Talk of facts, sentences, or propositions serves a similar purpose in the arguments: it is used to transform a sentence like "$\exists x\, Tx$" into a singular term "$\langle \exists x\, Tx\rangle$", the denotation of which may then be assumed to be itself within the range of a quantifier in the sentence—or, in the case of self-reference, referred to by another singular term in the sentence.

Second, that same effect can also be achieved by more direct means, through the use of *higher-order* quantifiers—specifically, quantifiers binding variables in the position of *sentences* or *predicates*. Once such quantifiers are in play, it is plausible to suppose that they satisfy analogous grounding principles to the first-order ones. In particular, given the principle ∃-*Grounding* that every true instance of a first-order existential quantification grounds that quantification, it is natural also to accept its sentential and predicational higher-order analogues

∃-*Grounding*$_S$	B(A)	/ $B(A) \prec \exists P\, B(P)$
∃-*Grounding*$_P$	$B(F)$	/ $B(F) \prec \exists X\, B(X)$

according to which every true instance of a higher-order existential quantification grounds that quantification. Using ∃-*Grounding*$_S$, a particularly efficient construction of a logical puzzle of ground becomes possible (Krämer 2013):

1	(1)	$\exists P\, P$	Assumption
1	(2)	$\exists P\, P \prec \exists P\, P$	1 ∃-*Grounding*$_S$
1	(3)	$\bot$	2 *Irreflexivity*

The circle of aboutness here results directly from the fact that "$\exists P\, P$" is seen as a *true instance of itself*, justifying the application ∃-*Grounding*$_S$. (A hybrid of this argument and its earlier Finean counterpart is given in (Correia 2014: §6), using both quantification into sentence position and a version of T-*Grounding*.)

To construct a puzzle using ∃-*Grounding*$_P$, we require two additional principles that govern the abstraction of a predicate using the λ operator. Intuitively, "λ" may be read as "is such that", so that "λx(Fx)a" for example means: *a* is such that it is *F*. The principles we require are roughly analogous to T-*Introduction* and T-*Grounding*:

λ-*Introduction*	$B(a)$	/ $\lambda x(B(x))a$
λ-*Grounding*	$\lambda x(B(x))a$	/ $B(a) \prec \lambda x(B(x))a$

Using these two principles in conjunction with ∃-*Grounding*$_P$, we can derive a contradiction as follows (cf. Donaldson 2017: 794):

1	(1)	$\exists X\, Xa$	Assumption
1	(2)	$\lambda x(\exists X\, Xx)a$	1 λ-*Introduction*
1	(3)	$\exists X\, Xa \prec \lambda x(\exists X\, Xx)a$	2 λ-*Grounding*
1	(3)	$\lambda x(\exists X\, Xx)a \prec \exists X\, Xa$	2 ∃-*Grounding*$_P$

A circle of aboutness here results from the fact that in taking $\lambda x(\exists X\, Xx)a$ to be an instance of $\exists X\, Xa$, we are taking the existential quantifier to include in its range the very property ascribed to *a* in $\exists X\, Xa$. For both types of higher-order quantification, one can also construct versions of the puzzles using universal rather than existential quantification.

It is possible to formulate a large number of further variations on these puzzles. Most of them depend on essentially the same assumptions and therefore introduce nothing interestingly new. One way to try to obtain interestingly new puzzles is to look for alternative interpretations of the predicate "*T*" for which at least the specific applications of T-*Grounding* and T-*Introduction* in the given puzzles may be argued to yield true results. For example, it appears that when "*T*" is read as "Socrates knows" in the very first puzzle, we obtain another compelling version. ("Knows" should here be read so that knowing a proposition *p* is knowing that *P*, when $p = \langle P \rangle$.) For the assumption that $\exists x\, Tx$ then amounts to the very plausible claim that Socrates knows something. T-*Grounding* now expresses the plausible principle that whenever Socrates knows that *P*, then *P* ≺ Socrates knows that *P*. The rule of T-*Introduction* is of course invalid under this reading, since Socrates doesn't know everything. But all that is required for the argument is the claim that Socrates knows that Socrates knows something, which is independently plausible. By ∃-*Grounding*, we then obtain that (Socrates knows that Socrates knows something) ≺ Socrates knows something, and by T-*Grounding*, that Socrates knows something ≺ (Socrates knows that Socrates knows something). So once more, we have a violation of *Asymmetry*.[1,2]

Desiderata for a Solution

Having considered the range of logical puzzles of ground, it is worth considering in general terms what we should expect of a satisfactory solution to the puzzles. Although this may not be wholly uncontroversial, I suspect that most philosophers would be happy to accept, at least under some reasonable interpretation, the following four desiderata.

Generality: A solution should cover all the puzzles of ground, not just a few of them. Of course, partial solutions may still be of interest, but their plausibility ultimately depends on the possibility of extending them in such a way as to provide a general solution.

Uniformity: A solution should be appropriately uniform across all cases. This is harder to make precise than the previous desideratum, but there are some clear cases of appropriately uniform (would-be) solutions and objectionably disuniform ones. For instance, a general ban on circular aboutness would constitute a uniform solution. A solution that rejects the relevant application of P-*Existence* in the first puzzle but only the application of *Transitivity* in the second puzzle, on the other hand, would appear objectionably disuniform. It would be overly simplistic to demand that a uniform solution reject a single principle invoked in all puzzles. But it may perhaps be plausible to demand, when different premises are rejected, that there be a unified explanation both of their falsity and of their initial appeal.

Proportionality. Given that all the assumptions required to generate the puzzles seemed initially plausible, there is no way around the fact that something that looks intuitive will have to go. But a solution should at least preserve as much of our intuitive views as possible: if one solution retains more of our intuitive judgements than another, then this constitutes a comparative advantage of that solution. We might add that other things being equal, preference should be given to retaining those principles that are most central to our overall philosophical and scientific theorizing.

Independent Motivation. A solution should provide independent motivation for rejecting the assumptions that the solution gives up. Thus, if the rejection of a particular principle is proposed, then a reason should be offered for giving up that principle independently of its role in the puzzles of ground.

Proposed Solutions

We turn finally to the question of how the logical puzzles of ground may be resolved. It is impossible within the confines of this survey to do justice to the various ingenious proposals that have been made on how best to do this. Instead, I shall focus on trying to explain the basic philosophical ideas underlying these various proposals and to say which of our desiderata would seem to be the most challenging ones to meet when developing these ideas.

Banning Circular Aboutness. As already emphasized, all of the puzzles can be described as exploiting some form of circular aboutness to derive an instance of circular grounding. Similar forms of circular aboutness lie at the heart of the more familiar *Liar* paradox of truth: a proposition that says of itself that it is not true seems to be both true and not true. In light of this, it is at first glance quite tempting to conclude that the solution to the puzzles must lie in the rejection as incoherent of such circular aboutness. One would then reject the applications of P-*Existence* in the first three puzzles, and in the higher-order versions, deny that $\exists P\, P$ and $\lambda x(\exists X\, Xx)a$ are genuine *instances* of $\exists P\, P$ and $\exists X\, Xa$, respectively.

At second glance, however, this seems to be an overreaction. For many examples of circular aboutness seem entirely harmless, such as the proposition that all the propositions I express today are interesting. This may be true or, more likely, false, but it does not seem paradoxical. The situation is even clearer with respect to sentences. Thus, consider

(X) The unique sentence on this page labelled (X) contains more than five words.

which is both self-referential and unproblematically true.

Typing. A natural reaction to this difficulty is to try to further specify conditions under which circular aboutness creates problems. Now, in each of the puzzles that apply P-*Existence*, the sentence expressing the problematic proposition *p* is not only *about* itself but contains an application of the *truth-predicate* to (a range of objects including) itself. What is illegitimate, one might therefore suggest, is not circular aboutness *per se* but application of a truth-predicate to a proposition which is itself expressed using that same truth-predicate. This is the conclusion that (Tarski 1944) drew from the Liar paradox (though for sentences, not propositions).

It would seem implausible, though, to conclude that one simply cannot coherently ascribe truth or untruth to propositions that are themselves about the truth of some proposition. Instead, one may advocate a *typed* conception of truth, on which there are many truth-predicates T_1, T_2, . . . of different types or *levels* 1, 2, . . . The application of a truth-predicate T_m to a proposition *p* is then allowed only if *p* is expressed using only truth-predicates of levels $n < m$. In the context of the theory of ground, a version of this solution is presented in (Korbmacher 2017); (Donaldson 2017: 795f) describes a way in which it may be transposed to the predicational higher-order setting (the additional extension to sentential higher-order quantification is unproblematic).

The main difficulty facing this kind of approach—sometimes called the *predicative* approach in the literature—is that it seems to impose excessively severe and counterintuitive constraints on what may be said and thereby violates the desideratum of *Proportionality*. Suppose, for example, that I say to you that all of the claims you're making today are true, and you say that some claim I'm making today is false. Then there is no way to assign levels to the truth-predicates we use under which we both succeed in generalizing over the entirety of the other's claims today. But this seems implausible. Moreover, our claims need not be paradoxical at all. For instance, if you make some other false claim, then my claim seems simply false and your claim true. (However, if all our other claims today are true, then the utterances described do give rise to a Liar-like paradox. This illustrates the striking fact that empirical matters can help determine if a self-application of truth is paradoxical.)

Building on Kripke's Constructions. Considerations like these led Saul Kripke, in his groundbreaking "Outline of a Theory of Truth" (1975), to develop a more sophisticated approach that allows for some self-applications of truth-predicates to receive a determinate truth-value without giving rise to paradox. Building on Kripke's methods, (Fine 2010) has described (though not advocated) solutions to the puzzles of ground that likewise permit circular aboutness while avoiding in a principled way the corresponding circularities of ground; parts of (Litland 2015) are devoted to extending and improving on these solutions.

Informally put, the key idea of the Kripkean approach is to classify a sentence A containing *T* as true (false) iff A (¬A) can be established through a certain step-wise procedure, *starting from sentences not containing T*. As long as neither A nor ¬A has been established in this way—and in some cases, in particular when A is paradoxical, this never happens—A counts as neither true nor false. If A is of the form *T*("B"), then A (¬A) may be established if B (¬B) has been established. There are then a number of possible views on how A may be established when it is of another form, such as *T*("B") ∨ C, say. It would take us too far afield to discuss all, or indeed any, of them in detail. What follows is a sketch of what happens to (sentential versions of) our examples under perhaps the most promising, "Strong Kleene" approach. Consider the sentence "$\exists s\, T(s)$"—"there is a true sentence". Despite the fact that it self-applies *T*, it may be established via the following steps: $0 = 0 \Rightarrow T(\text{“}0 = 0\text{”}) \Rightarrow \exists s\, T(s)$. Similarly, the sentence s = "$0 = 0 \vee T(s)$", although it self-applies *T*, can easily be established: $0 = 0 \Rightarrow 0 = 0 \vee T(s)$.

How does all this lead to a solution to the puzzles of ground? Fine's basic idea is that we may take this step-wise construction to indicate how truths involving T are *grounded* by truths obtained at a previous stage in the construction. Roughly, the idea is that a truth A involving T is grounded by some truth B iff there is an appropriate way of establishing A that proceeds via B. The step-wise derivations would be taken to show, for example, that $T(\text{“}0 = 0\text{”}) \prec \exists s\, Ts$ and that $0 = 0 \prec (0 = 0 \vee Ts)$. Any putative circles of ground can in this way be avoided. For example, $\exists s\, T(s)$ cannot be established on the basis of $T\text{“}\exists s\, T(s)\text{”}$. For the only way the latter could be established is if we *already had* established $\exists s\, T(s)$. And likewise $0 = 0 \vee T(s)$ cannot be established on the basis of $T(s) = T(\text{“}0 = 0 \vee T(s)\text{”})$, for the only way to establish this would be if we *already had* established $0 = 0 \vee T(s)$.

The biggest challenge for this kind of approach, it might be thought, is the joint satisfaction of *Proportionality* and *Generality*. Some versions of the approach quite drastically weaken classical logic (the "Weak Kleene" approach in Fine), and some others yield a drastic revision of the impure logic of ground (the "Supervaluational" approach). The "Strong Kleene" approach lies between these extremes, weakening both classical logic and the impure logic of ground, but less drastically so. For example, $\exists$-*Grounding* and $\vee$-*Grounding* are slightly weakened so as to yield only that *some* true instance must ground a true existential generalization, and *some* true disjunct must ground a true disjunction, not that any such instance or disjunct does this. While these revisions may be tolerable, it is not easy to see how the approach may be extended to the fact-based or the higher-order versions of the puzzle at similarly limited cost. As Fine points out, in application to a fact-based version of the puzzles, we are led to deny the apparent truism that everything exists. Similarly, an extension to the higher-order setting may prevent us from endorsing *any* universal quantification $\forall P\, \mathrm{B}(P)$.

Allowing Circularity of Ground. A different option is to take the puzzles to show that there can be circles of ground after all. This option is favoured by (Correia 2014: 54f), (Rodriguez-Pereyra 2015), and (Woods 2018). This solution seems adequately *general* and *uniform*. It is less clear that there is *independent motivation* for allowing circles of ground, but a number of authors have argued for cases of symmetric or reflexive ground for independent reasons (see Chapter 17). The main issue, though, would seem to be *Proportionality*, since a *mere* rejection of *Irreflexivity* and *Asymmetry* would surely violate this constraint. For even if the puzzles have convinced us that circles of ground can sometimes arise, it would seem that this can happen only in special circumstances. Rather than merely being rejected, the principles should then be weakened so as to apply only in "normal" circumstances. The challenge is then to say which circumstances are special and which are normal.

(Woods 2018) addresses this task, maintaining that violations of *Irreflexivity* arise only in cases of *vacuous* grounding: cases in which the particular content of the truth that does the grounding does not matter, so that any other truth might replace it, and we should still have a true grounding claim. This claim, Woods argues, also helps to dissolve the intuitive discomfort with accepting instances of self-grounding. For once we have recognized that in certain cases, all we need to ground a given truth is some truth, *no matter which*, then we should not be surprised that even the truth to be grounded itself does the job. In addition, we seem to keep a fairly strong version of *Irreflexivity*. For in many cases, it will be relatively straightforward to establish that a putative instance of self-grounding would be nonvacuous and thus may still be ruled out by the envisaged weakening of *Irreflexivity*.

The main challenge, perhaps, for this approach is to show that in *all* versions of the puzzles, the putative violations of *Irreflexivity* are instances of vacuous grounding. This is easy for some of our examples, like $\exists P\, P \prec \exists P\, P$, but less so for others. As it stands, the solution does not straightforwardly apply to the predicational higher-order case establishing $\exists X\, Xa \prec \exists X\, Xa$, though perhaps it can be suitably extended. After all, there does seem to be some sort of vacuity involved here, too. Roughly, all it takes to ground $\exists X\, Xa$ is *any* truth *about a*—no wonder, then, that $\exists X\, Xa$ itself does the trick, too. Whether the strategy works in all cases is an open question.

Replacing Strict by Weak Grounding Principles. Fine (e.g., 2012a: 51ff, 2012b: 1–4) makes a distinction between the familiar, *strict* variety of grounding and a somewhat less familiar, *weak* variety. Every case of strict grounding is also taken to be a case of weak grounding but not the other way round. For unlike strict grounding, weak grounding is not supposed to be asymmetric or irreflexive. Indeed, it is taken to be *reflexive*, so that every truth weakly grounds itself. We may continue to focus on the case of partial grounding. In that case, strict and weak grounding may be taken to be related in a particularly simple way in that a case of weak partial grounding ($\preceq$) is always either strict or mutual, so $A \prec B$ if and only if ($A \preceq B$ and not $B \preceq A$). With the notion of weak grounding at hand, another approach to the puzzles, defended in (Lovett forthcoming), is to weaken the relevant *impure* grounding principles—specifically, T-, $\exists$-, $\forall$-, and $\vee$-*Grounding*—by replacing $\prec$ in the conclusion by $\preceq$. This blocks the derivation of an absurdity in the very last step, since $\preceq$ satisfies neither *Irreflexivity* nor *Asymmetry*.

This weakening of the impure rules can be motivated on independent if not uncontentious grounds. (Correia 2010) has argued that grounding should be considered a *worldly* phenomenon and that there is no plausible worldly difference to be made between, for example, a truth A and its self-disjunction $A \vee A$. If so, A and $A \vee A$ should be interchangeable in the context of ground. (For more on the worldliness or otherwise of grounding, see Chapter 15.) Indeed, this is borne out under Fine's truthmaker semantics for ground (2012a, 2017). But then $\vee$-*Grounding* immediately clashes with Irreflexivity, since from $A \prec A \vee A$ we may, by interchange of A and $A \vee A$, obtain $A \prec A$. On this view, it is therefore natural to weaken $\vee$-*Grounding* in the way proposed and to permit the inference to $A \prec A \vee B$ only given the additional assumption that $\neg(A \vee B \preceq A)$. That assumption is always false in the applications made in the puzzles of the $\vee$-*Grounding* rule. Similar considerations apply with respect to $\exists$- and $\forall$-*Grounding* and their higher-order counterparts. Finally, the proposed weakening of the T-*Grounding* rule may be motivated on the basis of a deflationary conception of truth, on which, again, there is no worldly difference between a truth A and the truth $T\langle A \rangle$.

This approach appears to provide both a *general* and a *uniform* solution: it offers a common motivation for weakening the various rules, namely that their plausibility depends on an excessively fine-grained conception of ground. Moreover, since this general thought puts pressure on the rules under attack irrespective of the puzzles of ground, the approach also seems to satisfy *Independent Motivation*. But it might be argued that the approach does not do as well with respect to *Proportionality*. For note that it is not enough to claim that there is *an* interesting, worldly notion of grounding on which the impure rules must be weakened, and so the puzzles are avoided. The claim must be that there is *no* coherent notion of grounding, worldly or otherwise, that satisfies all these principles. And it has seemed very plausible to many—including, indeed, Correia in later work (e.g., 2011, 2014)—that although there may be this worldly notion of ground, there is *also* another, representational notion which does satisfy the original, "strict" impure rules. Once this is granted, we still require a story about how to avoid the puzzles for this other notion of ground.

Alleging Equivocation on "Grounds". Another strategy for responding to the puzzles is to argue that they commit a fallacy of equivocation. The idea would be that "grounding", as currently used by philosophers, is ambiguous between a number of interpretations and that there is no single interpretation of the term under which enough of our principles of grounding turn out true that we can generate one of the puzzles. A response of this sort is defended in (Peels 2013).

Adapted to our setting, Peels' proposal may be stated as follows: T-*Grounding* and ∃-*Grounding* require different interpretations of '≺' to come out valid. In order that ∃-*Grounding* be valid, '≺' must be given an interpretation under which grounding is a matter of one fact's *being an instance of* another. Under that interpretation, however, T-*Grounding* is not valid, since the fact that A is not an instance of the fact that $T\langle A\rangle$. In order for T-*Grounding* to be valid, '≺' must be given an interpretation under which grounding is a matter of one fact's *being ontologically more basic than and comprised by* another. Under that interpretation, however, ∃-*Grounding* is not valid, since the fact that $T\langle\exists x\, Tx\rangle$ is not ontologically more basic than the fact that $\exists x\, Tx$—indeed, it seems less basic—even though it is an instance of the latter. Finally, *Asymmetry* and *Transitivity* hold only under a uniform interpretation of '≺', and they do not hold under the (uniform) interpretation of '≺' as expressing the disjunction of the two forms of grounding just distinguished.

Prima facie, the proposal scores well with respect to *Proportionality*, in that it retains a version of each of the initially plausible-seeming principles invoked in the puzzles, though only in a restricted form. There may also be *Independent Motivation* for recognizing a multiplicity of notions of grounding. For claims like this, albeit often in a grounding-critical or -skeptical context, have been defended by several authors for reasons independent of the puzzles (cf. esp. (Koslicki 2015), (Wilson 2014), and Chapters 11–13). The main challenge would seem to be the desideratum of *Generality* in combination with *Uniformity*: it is unclear if the proposal can be extended, in an appropriately uniform way, to all the puzzles.

It seems straightforward to extend it to puzzles using ∀-*Grounding* instead of ∃-*Grounding*, since it may again be argued that this principle targets an *instance*-directed interpretation of "≺". With some more work, it may perhaps also be extended to puzzles only invoking ∨-*Grounding*, since the relation between a disjunction and its disjuncts seems very similar to the relation between an existential quantification and its instances and arguably dissimilar to that between the fact that $T\langle A\rangle$ and the fact that A. In particular, it may be argued that in the puzzle cases the disjunct–disjunction relationship fails to be a relationship from the ontologically more basic to the less basic. For these are cases in which the disjunct is Tp, where p is the very proposition expressed by the entire disjunction $Tp \vee 0 = 0$. It is harder to see, however, how the proposal could be extended to the puzzles involving higher-order quantifications. Here the sentential version appears particularly troublesome, since it invokes only one grounding principle, the apparently instance-directed ∃-*Grounding*$_S$.

Maintaining the Metaphysical Transparency of Truth and Higher-Order Quantification. It is a familiar idea in the theory of truth that the notion of truth exhibits a distinctive kind of *transparency*. The basic thought is that applying "it is true that" to a sentence seems to have only a minimal effect on the content expressed. What is said by an instance of "It is true that *P*", it seems, is hardly more than what is said by the corresponding instance of "*P*". One might be tempted to infer from this that instances of "*P*" and "$T\langle P\rangle$" are ground-theoretically equivalent: they ground and are grounded by exactly the same things. This would lead one to reject T-*Grounding*, but it would not avoid the puzzles, since from ∃-*Grounding*, we still obtain that $T\langle\exists x\, Tx\rangle \prec \exists x\, Tx$. Since it is being maintained that instances of "*P*" and "$T\langle P\rangle$" ground the same things, we may still infer that $\exists x\, Tx \prec \exists x\, Tx$, in violation of *Irreflexivity*.

However, there is an alternative way of understanding the transparency of truth that does block the derivation of the puzzle (deRosset manuscript). The idea is that truth is *metaphysically* transparent in the sense that truth-ascriptions play no robust *(metaphysically) explanatory* role: all of their explanatory power is merely inherited from the proposition to which truth is being ascribed. In particular, the fact that ⟨*P*⟩ is true *does not ground anything*—for to ground something is to have a robust explanatory role. In explanatory contexts, that fact merely serves as a *placeholder* for its own ground, the fact that *P*, and thereby is able to explain whatever may be explained by its being the case that *P*.

It is clear how the proposal avoids those puzzles of ground which make use of truth-ascriptions. In order to yield a fully *general* and *uniform* solution, it needs to be supplemented by analogous claims of metaphysical transparency for any alternative devices that can be used in place of the truth-predicate to generate a puzzle. For instance, it would also have to be claimed that predications involving λ-terms and/or higher-order quantifications do not ground anything but can only derivatively play an explanatory role by standing in for their instances or λ-free counterparts. Prima facie, it is particularly difficult to see how the proposal could be extended in a uniform way to the knowledge-versions of the puzzle, like the one about Socrates' knowing that Socrates knows something. The overall plausibility of the proposal and the extent to which it satisfies *Independent Motivation* seems to depend in large part on the plausibility of claiming that truth-ascriptions, higher-order quantifications, and so on literally do not ground anything and on the extent to which this idea serves to capture, in an independently attractive way, the intuition that truth is somehow transparent.[3]

Related Topics

Logics [Chapter 14]
Granularity [Chapter 15]
Strict Partial Order [Chapter 17]
Semantics [Chapter 36]

Notes

1 (Whitcomb 2012) uses a similar style of reasoning to derive a contradiction not from the assumption that Socrates knows something but from the assumption that there is an omniscient being and proposes to blame the contradiction on the latter assumption, concluding that an all-knowing God could not exist. (Rasmussen, Cullison, and Howard-Snyder 2013) and (Peels 2013) respond.

2 Two more puzzles concerning ground should be mentioned only to be set aside, since they are of a very different kind to those described earlier and require separate discussion. First, it is possible to construct puzzles invoking ground that are more similar to the classical Liar paradox. Famously, given a self-referential "Liar"-proposition p such that $p = \langle \neg Tp \rangle$, we obtain a contradiction by observing that the assumption that p is true leads to the conclusion that p is not true, which in turn leads to the conclusion that p is true. But as shown by (Korbmacher 2015), given common assumptions about ground, we can easily construct a sentence equivalent to "$\neg Tp$". For example, given the principles that grounding is factive and that any truth helps ground its disjunction with any other proposition, the sentence "$\neg(p \prec (p \vee 0 = 1))$" is equivalent to "$\neg Tp$". As a result, the assumption that there exists a proposition $p = \langle \neg(p \prec (p \vee 0 = 1)) \rangle$ yields a contradiction in much the same way as the assumption that the Liar-proposition exists. Second, (Fritz manuscript) shows that plausible and widely accepted assumptions about ground seem to conflict with a version of Cantor's Theorem by requiring the existence of a distinct proposition for each plurality of propositions. Although some solutions to this problem may also provide the required materials for a solution to the above puzzles and vice versa (such as Fritz's own proposed solution), the puzzle is too different and too complex to be discussed as part of this entry. See Chapter 15 for some more discussion of Fritz's puzzle.

3 For helpful comments, feedback, and advice, I would like to thank Fabrice Correia, Louis deRosset, Kit Fine, Jon Erling Litland, Adam Lovett, and Mike Raven.

References

Correia, F. (2010) "Grounding and Truth-Functions," *Logique et Analyse* 53(211), 251–79.

———. (2011) "From Grounding to Truth-Making: Some Thoughts," in A. Reboul (ed.) *Philosophical Papers Dedicated to Kevin Mulligan*, Genève. www.unige.ch/lettres/philo/mulligan/festschrift/. Reprinted 2014 in A. Reboul (ed.) *Mind, Values and Metaphysics: Philosophical Essays in Honour of Kevin Mulligan*, Vol. I, Heidelberg, New York, Dordrecht and London: Springer.

———. (2014) "Logical Grounds," *The Review of Symbolic Logic* 7(1), 251–79.

deRosset, L. (manuscript) "The Metaphysical Transparency of Truth," unpublished manuscript.

Donaldson, T. (2017) "The (Metaphysical) Foundations of Arithmetic?" *Noûs* 51(4), 775–801.

Fine, K. (2010) "Some Puzzles of Ground," *Notre Dame Journal of Formal Logic* 51(1), 97–118.

———. (2012a) "Guide to Ground," in F. Correia and B. Schnieder (eds.) *Metaphysical Grounding*, Cambridge: Cambridge University Press.

———. (2012b) "The Pure Logic of Ground," *The Review of Symbolic Logic* 5(1), 1–25.

———. (2017) "A Theory of Truthmaker Content II: Subject-matter, Common Content, Remainder and Ground," *Journal of Philosophical Logic* 46, 675–702.

Fritz, P. (manuscript) "Ground and Grain," unpublished manuscript.

Korbmacher, J. (2015) "Yet Another Puzzle of Ground," *Kriterion: A Journal of Philosophy* 29(2), 1–10.

———. (2017) "Axiomatic Theories of Partial Ground II. Partial Ground and Typed Truth," *Journal of Philosophical Logic* 47(2), 193–226.

Koslicki, K. (2015) "The Coarse-Grainedness of Grounding," *Oxford Studies in Metaphysics* 9, 306–44.

Krämer, S. (2013) "A Simpler Puzzle of Ground," *Thought: A Journal of Philosophy* 2(2), 85–9.

Kripke, S. A. (1975) "Outline of a Theory of Truth," *Journal of Philosophy* 72(19), 690–716.

Litland, J. E. (2015) "Ground, Explanation, and the Limit of Internality," *The Philosophical Review* 124(4), 481–532.

Lovett, A. S. (forthcoming) "The Puzzles of Ground," forthcoming in *Philosophical Studies*.

Peels, R. (2013) "Is Omniscience Impossible?" *Religious Studies* 49, 481–90.

Rasmussen, J., Cullison, A., and Howard-Snyder, D. (2013) "On Whitcomb's Grounding Argument for Atheism," *Faith and Philosophy* 30(2), 198–204.

Rodriguez-Pereyra, G. (2015) "Grounding is not a Strict Order," *Journal of the American Philosophical Association* 1(3), 517–34.

Tarski, A. (1944) "The Semantic Conception of Truth," *Philosophy and Phenomenological Research* 4(3), 341–76.

Whitcomb, D. (2012) "Grounding and Omniscience," in Kvanvig, J. (ed.) *Oxford Studies in Philosophy of Religion*, Vol. IV, Oxford: Oxford University Press, 173–201.

Wilson, J. (2014) "No Work for a Theory of Grounding," *Inquiry* 57(5–6), 1–45.

Woods, J. (2018) "Emptying a Paradox of Ground," *Journal of Philosophical Logic* 47(4), 631–48.

THOUGHT
A JOURNAL OF PHILOSOPHY
THE NORTHERN INSTITUTE OF PHILOSOPHY

Thought ISSN 2161-2234

ORIGINAL ARTICLE

A Note on the Logic of Worldly Ground

Stephan Krämer& Stefan Roski

University of Hamburg

In his 2010 paper ‘Grounding and Truth-Functions’, Fabrice Correia has developed the first and so far only proposal for a logic of ground based on a *worldly* conception of facts. In this paper, we show that the logic allows the derivation of implausible grounding claims. We then generalize these results and draw some conclusions concerning the structural features of ground and its associated notion of relevance, which has so far not received the attention it deserves.

DOI:10.1002/tht3.158

1 Introduction

Kit Fine is a British philosopher. However, not all philosophers are British. For example, there are also American philosophers. Now, does this latter fact—that there are American philosophers—hold partly *because*, or *in virtue of* the fact that Fine is a philosopher? It seems not. Given that Fine is British, it seems that his being a philosopher does not contribute at all to bringing it about, or making it the case, that there are American philosophers. We show that an otherwise attractive-seeming logic for ground, more precisely a logic for what is often called a *worldly* conception of ground, yields contrary results, and discuss what conclusions should be drawn from this.

Section 2 introduces the notion of worldly ground and highlights some of the philosophically and technically attractive advantages it seems to enjoy in comparison with its rivals. Section 3 proves the seemingly untoward results in the most well-developed logic of worldly ground, proposed by Fabrice Correia (2010), and discusses the results in informal terms, highlighting some costs of accepting them. Section 4 presents alternative derivations of the results, and connects them to certain structural principles about ground and the kind of relevance that grounding involves. Section 5 explores that connection further. Section 6 explores how adherents of a worldly conception of ground may react in view of our results and concludes.

2 Worldly ground

Recent years have witnessed an increasing interest in the notion of ground, where ground is taken to be a kind of non-causal priority among facts. A standard way to motivate this notion is to point to certain uses of the sentential connective ‘because’ or to uses of the phrase ‘in virtue of’ and cognate phrases that are particularly widespread in philosophical

Correspondence to: E-mail: stephan.kraemer@uni-hamburg.de

discourse.[1] Claims to the effect (i) that someone is in a given mental state *in virtue of* being in a certain physical state, or (ii) that a given ball is red or round *because* it is red, are typical examples of statements of ground. We shall assume that such statements express that a relation of ground obtains between certain facts.[2] For instance, (ii) expresses that the fact that a given ball is red or round is grounded in the fact that it is red.

Theories of ground can be seen as attempts to account for the philosophical use of 'in virtue of' and the relation thereby expressed. Intuitions on pertinent uses of 'because' or 'in virtue of' are therefore not only helpful to motivate the notion of ground, they also provide a crucial, if defeasible, test of adequacy for proposed theories of ground. In particular, if a theory entails an inacceptable 'in virtue of'-claim, this is a *prima facie* reason to reject it. This, at least, is common practice among participants of the debate, and we will follow it.[3]

There is a standard distinction between *conceptualist* and *factualist*, or *worldly* conceptions of ground.[4] According to the factualist conception, grounding relates relatively coarse-grained 'worldly' facts, whereas on the conceptualist view, the grounding relation is in general highly sensitive to how the facts in question are conceptualized.[5] While this issue is relevant for a host of topics related to grounding, we will focus on how it affects the *logic* of ground, that is, the principles and rules that govern the interaction between ground and the logical connectives. Several conceptualists have proposed systems by the lights of which the fact that A grounds, for instance, the fact that $A \wedge A$, the fact that $A \vee A$, and the fact that $\neg\neg A$.[6] Factualists would deny this.[7] For on a worldly conception of facts, the fact that A is arguably the same as the facts that $A \wedge A$, $A \vee A$ and $\neg\neg A$ — albeit each time *represented* under a different guise. Given that grounding is irreflexive, factualists thus cannot allow for any grounding relations obtaining in these cases. This does not mean that they cannot hold on to the intuition that, for example, conjunctive facts are grounded in their conjuncts. This will be the case, however, only if the given conjunctive fact is really different from its conjuncts.

It is not clear that the conceptualist and the factualist conceptions of ground represent mutually exclusive options; they might simply capture different, though related phenomena, both worthy of investigation. To the extent to which they are seen as competing views, however, we want to highlight three points that seem to favor the factualist conception.

a. The factualist conception seems more natural, given a standard way to *motivate* both the viability and the importance of the notion of ground. For this is usually done by appealing to a picture of *reality* as a *layered structure* (cf. Bennett 2011; deRosset 2013). According to such a picture, the world is not a mere aggregate of facts but falls into several layers that are connected by various relations of priority. Grounding is then thought of as one such relation. On this kind of view, therefore, grounding emerges as a relation among constituents of the world that exist, and can be individuated, *independently* of our conceptual (or linguistic) representations of them — a relation that *carves reality at its joints*. Yet, the conceptualists' account of ground does not seem to pay this picture its proper due in that it introduces distinctions that are unduly sensitive to our conceptualizations of reality (cf. Correia (2010, p. 258f)). In other words, conceptualists conflate mere shadows of language with real features of the world.

b. A related worry is that even on the conceptualist view, there should be some non-trivial notion of 'saying the same thing' under which the following schema is valid ('[A]' abbreviates 'the fact that A'):

EQUIV (a) If A and B say the same thing and [A] grounds [C], then [B] grounds [C].
(b) If C and D say the same thing and [A] grounds [C], then [A] grounds [D].

For, surely, not *every* linguistic difference should make for a difference in ground-theoretic status.[8] But so far, conceptualists have not come up with a viable general criterion to account for that intuition.

c. Finally, conceptualists have not yet shown that the logics they propose are sufficiently well-behaved. For, although the proof-theoretic side of logics of conceptual ground is by now fairly well investigated (cf. esp. Fine 2012a, 2012b), there is still no *semantics* relative to which we could prove soundness and completeness for the deductive systems. Using an elegant and well-motivated variant of situation semantics (called *truthmaker semantics*), Fine (2012b) establishes soundness and completeness for a set of structural rules for various grounding operators. As noted in his (2012a, p. 74), however, rules for the interaction between grounding operators and the truth functors that are widely accepted among conceptualists turn out to be unsound under this semantics. Strikingly, this is precisely because the semantics cannot discriminate among the semantic values of, for example, A and $A \vee A$.

As we have already seen, the factualist conception takes the idea of grounding as a joint-carving, worldly relation, more seriously. Moreover, in his (2010) work, Correia has developed a logic of worldly ground that puts this idea on a firm formal footing. Crucially, Correia has offered a semantics with respect to which his logic is provably sound and complete. And while the semantics and some aspects of the logic may seem somewhat contrived, Fine has pointed out that a modified version of his truthmaker semantics allows for a very natural definition of a notion of ground that agrees with Correia's (cf. Fine ms-b, p. 11). Finally, on the Correia/Fine view, a principled account of the notion of saying the same thing relevant to EQUIV is possible; in particular, its logic is that of Angell's (1977) notion of *analytic equivalence*.

The prospects for worldly ground thus look good. It takes the idea of ground as a joint-carving, worldly relation seriously, and appears to have a well-behaved logic in the system proposed by Correia. However, we will show that these apparent advantages notwithstanding, all is not well for proponents of worldly ground.

3 The argument

The only properly developed proposal for a logic of worldly ground is that of Correia (2010). This section proves that Correia's logic entails a close cousin of the above implausible claim of ground (cf. Section 1).[9]

Central to Correia's approach is a notion of *factual equivalence* ($\approx$), the role of which is to provide a necessary and sufficient condition for two formulae being, under any interpretation, interchangeable in the scope of ground ($<$). Correia takes the logic of $\approx$ to be Angell's logic of analytic equivalence. At first glance, that logic seems well suited to the task. On the one hand, analytic equivalence is narrower even than equivalence in the logic of first-degree entailment. Correia's logic can thereby allow for pairs like A and $A \vee (A \wedge B)$ being factually inequivalent, and thus for the former to ground the latter (cf. Correia 2010, p. 263). On the other hand, the characteristic kinds of pairs over which factualists and conceptualists are divided, like A and $A \vee A$, come out factually equivalent. It is therefore reasonable to hope that analytic equivalence is sufficiently fine-grained, but not too fine-grained, to play the role of factual equivalence in the factualist's system.

For our present purposes, the main thing that matters about factual equivalence is that it satisfies the law of distribution of $\vee$ over $\wedge$:

(X) $A \vee (B \wedge C) \approx (A \vee B) \wedge (A \vee C)$

We further make use of the following rules for ground:[10,11]

($\wedge$-Introduction) $A; B; (A \wedge B) \not\approx A; (A \wedge B) \not\approx B \,/\, A, B < A \wedge B$
($\vee$-Introduction) $A; (A \vee B) \not\approx A \,/\, A < A \vee B$
$B; (A \vee B) \not\approx B \,/\, B < A \vee B$
(Factivity) $\Gamma < C \,/\, C$
$\Gamma, A < C \,/\, A$
(Cut) $\Gamma < B; B, \Delta < C \,/\, \Gamma, \Delta < C$

The $\wedge$-Introduction rule captures the intuitive thought that a true conjunction is grounded by its conjuncts, provided it is not factually equivalent to either. (If it is, the irreflexivity of ground prevents it from being grounded by the conjuncts.) The $\vee$-Introduction rules capture the intuitive thought that a true disjunction is grounded by any true disjunct, provided it is not factually equivalent with it. The Factivity rules encode the compelling principle that only truths are grounded and that only truths are grounds. The Cut rule, finally, captures the plausible view that grounding is transitive.

Consider any statements A, B, and C such that A and B are true, C is either true or false, and the following four assumptions concerning *lack* of factual equivalence hold:

(A1) $(A \vee C) \not\approx A$
(A2) $(A \vee B) \not\approx B$
(A3) $((A \vee B) \wedge (A \vee C)) \not\approx (A \vee B)$
(A4) $((A \vee B) \wedge (A \vee C)) \not\approx (A \vee C)$

(A1)–(A4) are weak assumptions that come out true for most choices of statements A, B, and C.[12]

The rules of $\vee$-Introduction yield

(1) $A < A \vee C$
(2) $B < A \vee B$

Since $A \vee C$ and $A \vee B$ are true, $\wedge$-Introduction and two applications of Cut yield

(3) $A, B < (A \vee B) \wedge (A \vee C)$

By (X), it follows that

(4) $A, B < A \vee (B \wedge C)$

We maintain that (4) has many instances that are implausible, even though the corresponding instances of the assumptions of our derivation are true.

Our first example is very close to the one with which we began this paper. It is a standard view that as far as grounding is concerned, an existential quantification like 'someone is an American philosopher' is just like the corresponding (potentially infinite) disjunction. If so, then we may paraphrase the quantification by 'someone other than Fine is an American philosopher ∨ Fine is an American philosopher'. The second disjunct is plausibly ground-theoretically equated with the conjunction 'Fine is a philosopher ∧ Fine is American'. It is easy to check that if we set A to 'someone other than Fine is an American philosopher', B to 'Fine is a philosopher' and C to 'Fine is American', all our premises come out true. So we obtain the conclusion that the fact that someone is an American philosopher is grounded by the facts that someone other than Fine is an American philosopher, and that Fine is a philosopher. In other words, this last fact *helps* ground the fact that someone is an American philosopher. But this is implausible given an understanding of ground that is mediated by our intuitive understanding of 'in virtue of'. For, we would not say that the fact that there are *American* philosophers obtains partially in virtue of some *non*-American (Fine) being a philosopher.

We note two ways of dramatizing the result. First, we can let C, and thus $B \wedge C$, be necessarily false. For instance, let A mean that something is a prime number, B that 4 is a number, and C that 4 is prime. It seems highly implausible to hold that the fact that something is a *prime* number is due in part to the fact that 4 is a number. Second, we may even set C to $\neg B$, obtaining that B helps ground $A \vee (B \wedge \neg B)$.[13] This seems most counter-intuitive, for it implies that for any truth A and any claim B, the fact that $A \vee (B \wedge \neg B)$ is due in part to either B or $\neg B$. Consider the fact that snow is white or the number of hairs on my head is both odd and even. Clearly, that this fact obtains is not due even in part to the number of hairs on my head being odd, or due to that number being even, whichever is in fact the case.

We conclude that Correia's logic of ground has implausible results. While the appeal to analytic equivalence permits us to abstract from mere differences in representation such as may seem to obtain between A and $A \vee A$, it does so at the cost of obliterating real distinctions tracked by our informal understanding of 'in virtue of' and cognate phrases.

4 Variations

It is instructive to note that we can also obtain similar results by appealing to resources that bear no obvious connection to the distributivity principle (X) appealed to before. The basis of these derivations is a structural principle of *convexity* for ground. It says that if a given truth C is grounded by some collection of facts as well as a *subset* of the same collection of facts, then it is grounded by any collection of facts that lies *between* these two, that is, which is both a subset of the first and a superset of the second collection:

CONVEXITY: $\Gamma < C; \Gamma, \Delta, E < C / \Gamma, \Delta < C$

Now if A, B, C are all true, we can (typically[14]) show both that $A < A \vee (B \wedge C)$ and that $A, B, C < A \vee (B \wedge C)$. By CONVEXITY, it then follows that $A, B < A \vee (B \wedge C)$. This is like the above case, only that here we had to assume that C is true.

We are tempted to suggest that typical instances of this grounding claim should be rejected. Roughly, it seems to us that some collection that includes B can only constitute a full ground of $A \vee (B \wedge C)$ *if it also includes C*. We admit, though, that intuitions about the appropriate use of 'in virtue of' et al on their own lend only limited support to this contention. Consider the fact that Quine is an American philosopher ∨ (Fine is a philosopher ∧ Fine is British). Does this fact obtain wholly in virtue of the facts that Quine is an American philosopher and that Fine is a philosopher? Or is it only the fact that Quine is an American on its own, which fully grounds the disjunctive fact? Intuition does not seem to decide the case.

However, in the presence of some additional plausible assumptions, CONVEXITY also implies the results previously found to be unacceptable. Firstly, according to a very popular view, grounding is an *internal* relation in the sense that if some facts Γ ground another C, they do so in every world in which all of Γ and C obtain.[15]

INTERNALITY: $\Gamma < C \rightarrow \Box((\bigwedge\Gamma \wedge C) \rightarrow \Gamma < C)$

Now assume that grounding is necessarily internal and that A and B are true. Assume further that $A \wedge B \wedge C$ is contingently false, so C is contingently false and compossible with A and B. Then in any world in which A, B, and C are all true, we have $A, B < A \vee (B \wedge C)$. But since grounding is internal in every world, and all of A, B, and $A \vee (B \wedge C)$ obtain in the actual world, it follows that in the actual world $A, B < A \vee (B \wedge C)$. So every instance of this schema which is derivable by the above method, *and* in which A, B, and C are compossible, can also be obtained from CONVEXITY and INTERNALITY.

Secondly, it is plausible that we should also recognize a *non-factive* counterpart of the notion of ground.[16] If so, it seems natural to think that a structural principle like CONVEXITY should hold for non-factive ground ($<_\circ$) if it holds for factive ground:

CONVEXITY*: $\Gamma <_\circ C; \Gamma, \Delta, E <_\circ C / \Gamma, \Delta <_\circ C$

Given the obvious connecting principle that if $\Gamma <_\circ C$ and $\bigwedge\Gamma \wedge C$, then $\Gamma < C$, we can then derive all the above unwelcome results.

Since Correia's logic for ground does not have modal operators, INTERNALITY is not derivable in it. Still, it is a very plausible seeming claim, and one that Correia has elsewhere committed himself to. CONVEXITY provably holds in Correia's logic.[17] As for CONVEXITY*, Correia has no operator for non-factive ground in his system, but it is clear what would correspond semantically to such an operator. If we were to introduce $<_\circ$ accordingly, CONVEXITY* would be derivable.[18] We conclude that the most plausible way of avoiding the results of Section 3 also involves a rejection of the principles of convexity.[19] As we see it, their problematic character also points to an important fact about the notion of *relevance* that guides central intuitions invoked by theorists of ground. We elaborate on this issue in the next section.

5 Relevance

It is commonly held that one of the distinctive features of the notion of ground is that it imposes a constraint of *relevance* on the grounds of a given fact; it is this feature of the notion which renders it *non-monotonic*, that is, such that adding arbitrary facts to a ground of a fact does not in general yield another ground of that fact.[20] In addition, the notion of *full* ground imposes a constraint of *sufficiency*: if some facts are to constitute a full ground of some further fact, then their obtaining must in some appropriate understanding be sufficient for the grounded fact's obtaining. Metaphorically speaking, these two constraints put bounds on the set of a fact's full grounds both from below and from above. It is bound 'from below' by the sufficiency constraint, which forces us to put enough into a given collection of facts that they are jointly sufficient for the fact to be grounded. It is bound 'from above' by the relevance constraint, which prevents us from enlarging the collection of facts that is to be the ground in arbitrary ways.[21]

It is plausible to suppose that the two constraints, on some suitable understanding, are jointly sufficient: if a collection of facts Γ is, in the appropriate senses, both sufficient and relevant for the obtaining of fact C, then $\Gamma < C$. At any rate, it seems very unclear what constraints might be satisfied in all typical examples for ground that cannot plausibly be subsumed under either the heading of sufficiency or that of relevance. But then it may seem as though ground *should* satisfy CONVEXITY. For suppose that (i) $\Gamma < C$ and (ii) $\Gamma, \Delta, E < C$. Then by (i) Γ, Δ is sufficient for C. Moreover, by (ii), a superset of Γ, Δ is relevant for C, so Γ, Δ must also be relevant. What our discussion shows is that this very tempting line of reasoning must be flawed somehow.

We would like to suggest that the flaw consists in an overly simplistic picture of relevance. We assumed, in effect, that if a given collection is relevant in the pertinent sense to a given fact, so is any subset of that collection. This follows immediately if the notion of relevance is understood *distributively*, so that the relevance of a given collection consists simply in the relevance of each member. We are inclined to hold that a notion of relevance that can do justice to our intuitive understanding of 'in virtue of' must be understood *collectively*, so that some facts can jointly be relevant to another without each member on its own, or indeed arbitrary subcollections, being thus relevant. In particular, we suggest that A, B, C are jointly relevant, in the appropriate sense, to $A \vee (B \wedge C)$, but A, B is not. The reason is that B is relevant only in combination with either C or a full ground of C. As long as A is neither, the collection A, B is not suitably relevant for $A \vee (B \wedge C)$.

6 Conclusion

We have shown that an initially plausible logic attempting to take the intuition of grounding as a *worldly* relation seriously entails unacceptable consequences. How can an adherent of a worldly conception of ground react? Perhaps the most natural option is modifying the logic. Our derivations relied on either the distributivity principle (X) for factual equivalence or the CONVEXITY principles. One thus might consider giving up these principles. However, since this would involve giving up on Angell's system as a guide for interchangeability in the scope of <, factualists would lose their respective principled account — one of the advantages Correia's logic seemed to offer. It is also unclear whether such a modified notion of factual equivalence allows for a well-behaved logical system — a second advantage the notion seemed to enjoy compared to conceptualist views. The question whether the notion of worldly ground has an attractive logic, and how it compares to conceptualist rivals, is still open.

Acknowledgments

We would like to thank Fabrice Correia, Kit Fine, as well as an anonymous referee for this journal, and our fellow members of the *Phlox* research group Michael Clark and Naomi Thompson for very helpful comments and discussion. Work on this paper was carried out while S.K. was working on his research project *The Logic and Metaphysics of Ground*, funded by the DFG, and S.R. was part of the *Kompetenzzentrum Nachhaltige Universität* of the University of Hamburg. We gratefully acknowledge the support.

Notes

1 This is not the place to give a comprehensive introduction to the debate. For this, consider Correia and Schnieder (2012a) and Trogdon (2013).

2 We make this assumption purely for ease of expression; nothing in our argument turns on whether ground is strictly a relation between facts (for discussion, see Correia and Schnieder 2012a, p. 10ff).

3 Compare, for example, Correia (2010, p. 263).

4 Compare, for example, Correia (2010, p. 256f) and Correia and Schnieder (2012a, p. 14f).

5 Whether one wants to still call those relata facts or perhaps rather true propositions is of no matter for present purposes. We continue to talk about facts for brevity's sake. Note that on both conceptions, grounding is hyperintensional, and does not in general allow even logically equivalent sentences to be interchanged salva veritate. Moreover, there are disagreements in both camps as to just *how* fine- or coarse-grained the relata are to be construed.

6 This holds in Fine's (2012a) *Impure Logic of Ground* and also given the principles proposed by Rosen (2009, p. 117ff). Related results can be obtained in Schnieder's Logic of 'because' (2011).

7 This holds at least for Correia (2010), on whose system we shall concentrate below. Perhaps not all factualists will give the same verdict on the case. Audi (2012, p. 700f) takes the relation to be even more coarse-grained.

8 Compare Correia (2010, p. 266), where he makes some suggestions on behalf of the conceptualist. Schnieder (2010) discusses a closely related point.

9 After we finished the paper, Correia informed us that he has independently discovered the same problem for his logic of factual equivalence, and develops an improvement in an as yet unpublished manuscript (Correia ms).

10 The rules are written in the format 'Premise; Premise; … /Conclusion'. The *semi-colon* is used to separate premises, and the comma to separate the sentences in the left-hand argument place of <, which can be occupied by any finite number of sentences. We have

changed Correia's notation to that employed in Fine's papers. By '$\not\approx$' we denote lack of factual equivalence [i.e., '$A \not\approx B$' abbreviates '$\neg (A \approx B)$'].

11 Note that Correia (2010, p. 269) in the end replaces the introduction rules by refinements in which the appeal to factual inequivalence is replaced by one to lack of conjunctive-disjunctive containment, that is, factual equivalence of a premise to some disjunction of conjunctions, one of which contains the conclusion as a conjunct. The change has no real effect on the uses we make of the rules.

12 Using the refined rules (cp. the previous footnote), we instead require assumptions stating lack of conjunctive-disjunctive containment. These are still true for most choices of A, B, and C. Roughly, as long as our A, B, and C have independent truthmakers, the assumptions all turn out true.

13 This slightly reduces the number of statements for which (A1)–(A4) are true. Again, it suffices to assign statements with independent truthmakers to A and B to ensure that the assumptions hold good.

14 Modulo the relevant conditions of factual inequivalence; this caveat will be left implicit in what follows.

15 Correia (2005, p. 61) and (2014a, p. 88), Bennett (2011, p. 32f), Audi (2012, p. 697), and deRosset (2013, p. 20) all accept principles that clearly entail INTERNALITY. Also Fine (2012a, p. 76) accepts it. Leuenberger (2014) presents counterexamples to INTERNALITY.

16 See, for example, Fine (2012a, p. 48f) and Correia (2014b, p. 36).

17 *Proof sketch*: Say that C is disjunctively contained in A ($A \geq^d C$) iff A is factually equivalent to some disjunction of which C is a disjunct. Application of Correia's Reduction Theorem (2010, p. 19) to the premises and conclusion of CONVEXITY reveals that it suffices to show that if the conjunction of the sentences in Γ and the conjunction of the sentences in Γ, Δ, and E is disjunctively contained in C, then so is the conjunction of the sentences in Γ and Δ. This can be shown essentially by applications of the distributivity of ∨ over ∧ and ∧ over ∨ as well as the fact that $A \approx A \vee A \approx A \wedge A$. We should note that Correia voices some dissatisfaction with the Reduction Theorem. However, against the background of Fine's truthmaker semantics, it seems perfectly well motivated.

18 The proof is a proper part of the proof of CONVEXITY sketched above.

19 There is an intriguing connection between CONVEXITY and (X) in Fine's truthmaker semantics for worldly ground (cf. Fine ms-b). In that semantics, both turn crucially on a condition of convexity that Fine imposes on ground-theoretic content: that if something is both a part of a truthmaker of a proposition and has a truthmaker of the same proposition as part, then that object is itself a truthmaker of the proposition. It would be very interesting to investigate the logics of factual equivalence and ground obtained by dropping the condition of convexity on Finean propositions.

20 See, for example, Fine (2012a, p. 56f; 2012b, p. 2), Dasgupta (2014, p. 4), Schnieder (2011, p. 450), and Correia (2010, p. 11f).

21 We owe the picture to Fine (ms-a).

References

Angell, Robert B. "Three Systems of First-Degree Entailment." *Journal of Symbolic Logic* **47** (1977): 147.

Audi, Paul. "Grounding: Toward a Theory of the *In-Virtue-Of* Relation." *Journal of Philosophy* **CIX** (2012): 685–711.

Bennett, Karen. "By Our Bootstraps." *Philosophical Perspectives* **25** (2011): 27–41.
Correia, Fabrice. *Existential Dependence and Cognate Notions*. Philosophia: München, 2005.
———. "Grounding and Truth-Functions." *Logique et Analyse* **53**.211 (2010): 251–79.
———. "From Grounding to Truth-Making: Some Thoughts," in *Mind, Values and Metaphysics: Philosophical Essays in Honor of Kevin Mulligan*, Vol. **I**, edited by A. Reboul. Heidelberg, New York, Dordrecht, London: Springer, 2014a, 85–98.
———. "Logical Grounds." *Review of Symbolic Logic* **7**.1 (2014b): 31–59.
———. "On the Logic of Factual Equivalence.'" (ms). Unpublished manuscript.
Correia, Fabrice and Benjamin Schnieder. "Grounding: An Opinionated Introduction," in , edited by Fabrice Correia and Benjamin Schnieder. 2012a, 1–36.
Correia, Fabrice and Benjamin Schnieder, (eds). *Metaphysical Grounding*. Cambridge: Cambridge University Press, 2012b.
Dasgupta, Shamik. "On the Plurality of Grounds." *Philosopher's Imprint* **14** (2014): 1–28.
Fine, Kit. "Guide to Ground," in , edited by Fabrice Correia and Benjamin Schnieder. 2012a, 37–80.
———. "The Pure Logic of Ground." *The Review of Symbolic Logic* **5**.1 (2012b): 1–25.
———. "A Theory of Truth-Conditional Content I." (ms-a).
———. "A Theory of Truth-Conditional Content II." (ms-b).
Leuenberger, Stephan. "Grounding and Necessity." *Inquiry* **57**.2 (2014): 151–174.
Rosen, Gideon. "Metaphysical Dependence: Grounding and Reduction," in *Modality: Metaphysics, Logic, and Epistemology*, edited by B. Hale and A. Hoffmann. Oxford: Oxford University Press, 2009, 109–136.
deRosset, Louis. "Grounding Explanations." *Philosopher's Imprint* **13** (2013): 1–26.
Schnieder, Benjamin. "A Puzzle about "Because"." *Logique et Analyse* **53**.211 (2010): 317–343.
———. "A Logic for "Because"." *The Review of Symbolic Logic* **4**.3 (2011): 445–465.
Trogdon, Kelly. "An Introduction to Grounding," in *Varieties of Dependence*, edited by M. Hoeltje, B. Schnieder, and A. Steinberg. München: Philosophia, 2013, 97–122.

DOI: 10.1111/phpr.12692

ORIGINAL ARTICLE

Philosophy and Phenomenological Research

Singular troubles with singleton socrates

Stephan Krämer

University of Hamburg

Correspondence
Stephan Krämer, University of Hamburg.
Email: stephan.kraemer@uni-hamburg.de

Funding information
Deutsche Forschungsgemeinschaft, Grant/Award Number: KR 4516/2-1

Abstract
I identify a problem for Kit Fine's truthmaker semantics for ground: it excludes a range of plausible structures of grounding hierarchies. Specifically, there is an attractive view about the grounds of the existence of individual singleton sets which is inconsistent under the truthmaker account. I then develop a modification of the truthmaker account which avoids this difficulty and show that it preserves most of the desirable features of the original account, such as the propositional logic of ground it yields.

1 | INTRODUCTION

Some facts obtain in virtue of other facts, which may then be said to *ground* the former. Many important metaphysical questions concern matters of ground: whether normative facts are grounded by non-normative ones, whether mental facts are grounded by physical ones, and so on. In order that debates on such questions can proceed in a fruitful and rigorous fashion, it is desirable that we have an appropriate formal framework available, within which the various competing positions can be articulated and their consequences drawn out. As part of this, it is desirable that we have an adequate semantics for statements of ground. Ideally, this will help clarify and structure debates on metaphysical ground in a similar way in which possible worlds semantics does this for debates about metaphysical necessity.

A promising candidate for this role is the *truthmaker semantics* for ground, as developed by Kit Fine in several recent publications (Fine, 2012a, 2012b, and 2017b). The basic idea of this approach is to model ground as a form of *entailment*, which is however characterized by an inclusion relation not between the sets of *worlds* in which premises and conclusion are true, but between the sets of *situations* or *states* that *make* premises and conclusion true. The relationship of truthmaking here may be seen as a (partial) semantic correlate of the relationship of ground that is expressed in the object language. Roughly, a state is taken to make a proposition true iff the state's obtaining would ground the truth of the proposition. Correspondingly, the truthmaking relation is taken to share several distinctive logico-structural features of grounding. In particular, the truthmaker must *necessitate* what it makes true, much as grounds are standardly taken to necessitate what they ground, it must be *relevant* to what it makes true, much as grounds are to what they ground, and like ground, the truthmaking relation is consequently *non-monotonic*.

In the first part of the paper, I present a problem for the truthmaker semantics for ground. The problem consists in the fact that there seem to be possible *grounding structures* which are ruled out by the logic of ground that we obtain under the truthmaker account. Specifically, as I will show, the account excludes the possibility that the grounding tree below a truth P could instantiate this structure:

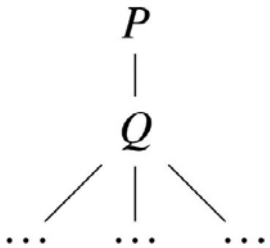

In this structure, there is a single proposition Q which grounds P, such that every other ground of P also grounds Q, and thus grounds P *via* Q, as it were. We might therefore describe it as a *single-conduit* grounding structure. Although it is not obvious that this structure is ever instantiated, there seems to be no independent motivation for taking it to be ruled out by the very *logic* of ground. Moreover, there are perfectly reasonable and attractive views on which there are instances of this structure. The example I shall focus on concerns the existence of singleton sets: one may plausibly take the grounds of the truth that {Socrates} exists to instantiate the single-conduit structure. I conclude that its exclusion of this structure is at least a highly problematic feature of the truthmaker account. So the question arises if there is an attractive way to avoid it.

In the second part of the paper, I show that there is, by developing a modification of the truthmaker account that accommodates the structures in question. The key idea is to recognize an additional *source* of grounding relationships. As I will explain, under Fine's account, grounding relationships turn out to arise in essentially two ways. Given a proposition with a certain set of truthmakers, we can obtain a ground of that truth either by *removing* some of its truthmakers, or by *decomposing* a truthmaker into its proper parts. I propose that we countenance a third way to obtain a ground, in which the truthmakers are not decomposed into proper parts, but replaced by more fundamental states that – as I shall say – *generate* them. We thereby recognize another partial semantic correlate of ground: in addition to the state-proposition relation of truthmaking, a state-state relation of generation. The resulting framework for ground, I argue, is strictly more powerful than the old one. It allows us to straightforwardly capture the previously problematical views, and it allows for every view that could be articulated in the old framework to still be articulated, and in effectively just the same way. Moreover, as I show in a formal appendix, the modified account I propose makes room for single-conduit structures without affecting the more basic parts of the logic of ground; in particular, it yields the same so-called *pure* and *propositional* logics of ground as the old one.[1]

2 | THE TRUTHMAKER FRAMEWORK

This section provides the background for the subsequent discussion by describing a formal framework for theorizing about ground, based on the accounts proposed by Fine (2012a, 2012b, 2017b).

[1]The term 'pure logic of ground' is borrowed from (Fine, 2012b). Fine's pure logic of ground comprises only relatively simple structural principles of ground, such as transitivity and irreflexivity principles. The propositional logic of ground additionally includes principles connecting ground to the boolean operations of conjunction, disjunction, and negation.

We assume as given some background language $\mathscr{L}$ in which we can express the propositions whose ground-theoretic relations we are interested in. To study these relations, we use what we may call the *language of ground* $\mathscr{L}_G$ *over* $\mathscr{L}$, consisting of all and only the expressions of the forms

- $\Gamma \leq C$
- $\Gamma < C$
- $A \preccurlyeq C$
- $A \prec C$

where A and C are formulas of $\mathscr{L}$, and Γ is any set of such formulas. The intended interpretation is that $<$ stands for strict full ground, $\leq$ for weak full ground, $\preccurlyeq$ for strict partial ground, and $\preccurlyeq$ for weak partial ground. We shall say more about what these notions are below.

We may then interpret the sentences of $\mathscr{L}$ by assigning propositions to them. In truthmaker semantics, propositions are identified not with sets of worlds, but with sets of states.[2] Like worlds, states are conceived of as specific, or determinate; in particular, whenever a state verifies a disjunction, it does so by verifying at least one of its disjuncts.[3,4] Unlike possible worlds, however, they may be incomplete, leaving open the truth-value of many propositions, and they may be impossible, like the state of a ball being red all over and green all over at the same time.[5]

Instead of the relation of a proposition being true at a world, in truthmaker semantics we appeal to a relation of a proposition being *exactly made true by*, or *exactly verified by*, a state. For a state to (exactly) verify a proposition in this sense, it is required that the state be *wholly relevant* to the truth of the proposition. So the state of snow's being white does not verify the proposition that $2 + 2 = 4$, because it is irrelevant to the truth of that proposition. And the state of it being sunny and warm does not verify the proposition that it is sunny or rainy, because it contains as an irrelevant part the state of it being warm, and is therefore not wholly relevant to the truth of the proposition.

As the talk of irrelevant parts implies, states are taken to be ordered by part-whole ($\sqsubseteq$). It is assumed furthermore that given any set of states $T = \{s_1, s_2, \ldots\}$ we may form their fusion $\bigsqcup T = s_1 \sqcup s_2 \sqcup \ldots$, which is taken to be the smallest state containing each of $s_1, s_2, \ldots$ as a part.[6]

A proposition, within truthmaker semantics, is thus a set of states. But we may not wish to count *every* set of states as a proposition. Following Fine, we will require propositions to be non-empty, and

[2]Strictly speaking, this is true only of so-called *unilateral* propositions. In many contexts, it is useful to work with a *bilateral* conception of propositions, under which a proposition is identified with a pair of two sets of states, one comprising the states verifying the proposition, the other comprising the states falsifying the proposition. Since the move to bilateral propositions does not affect any of the issues of concern in this paper, for simplicity's sake, it is better to stick to unilateral ones. For details on these issues, see (Fine, 2017a).

[3]I use the terms 'specific' and 'determinate' interchangeably. On my use, the elimination of disjuncts in a disjunction, and the move from a determinable property to one of its determinates, is correlated with an increase in specificity. The addition of a conjunct to a conjunction, or of differentia to a genus, however, is not. (This contrasts with the usage in (Rosen, 2010, §11), on which all these moves count as increasing specificity, but only the former count as increasing determinacy.)

[4]Strictly speaking, as we shall see below, a state may also verify a disjunction by being the fusion of verifiers of the disjuncts. This subtlety does not affect the point regarding specificity.

[5]Not all applications of truthmaker semantics require impossible states. Indeed, in (Fine, 2012b) and (Fine, 2012a), Fine does not envisage inconsistent or even just non-obtaining states. He does do so in (Fine, 2017b), and I believe that such states are clearly required in order to obtain an adequate account of ground. For the purposes of this paper, not much hangs on this.

[6]The operation of fusion is crucial in particular to the treatment of conjunction: a state is taken to verify a conjunction just in case it is the fusion of a verifier of the one conjunct and a verifier of the other conjunct.

to be closed under (non-empty) fusion, so that whenever some states verify a given proposition, so does their fusion.[7] Given an assignment of propositions to the sentences of $\mathscr{L}$, we then need to say which of the sentences of our language of ground $\mathscr{L}_G$ are true. To that end, we first define four relationships of ground over the propositions, corresponding to the four types of grounding statements in $\mathscr{L}_G$ (cf. Fine, 2012b, p. 9):

($\leq$)	$P_1, P_2, \ldots \leq Q$ iff $s_1 \sqcup s_2 \sqcup \ldots$ verifies Q whenever $s_1, s_2, \ldots$ verify $P_1, P_2, \ldots$, respectively
($\preccurlyeq$)	$P \preccurlyeq Q$ iff $\Gamma, P \leq Q$ for some set of propositions Γ
($<$)	$P_1, P_2, \ldots < Q$ iff $P_1, P_2, \ldots \leq Q$ and $Q \preccurlyeq P_i$ for no $i \in \{1, 2, \ldots\}$
($\prec$)	$P \prec Q$ iff $P \preccurlyeq Q$ and not $Q \preccurlyeq P$

The grounding statements in $\mathscr{L}_G$ are then considered true iff the corresponding grounding relation holds between the propositions assigned to the relevant sentences of $\mathscr{L}$.

Two comments on these definitions of grounding relations are in order. First of all, it should be noted that the target notions of ground are *non-factive*: there is no requirement that $P_1, P_2, \ldots, Q$, or the propositions in Γ be true in order for them to instantiate the various grounding relationships.[8] It would be easy enough to define corresponding factive notions, but for present purposes this would needlessly complicate things.

Second of all, it should be noted that the most basic notion of ground defined here, symbolized by $\leq$, besides being non-factive, is somewhat unfamiliar in a further way. In contrast to the intuitive understanding of ground, it is *reflexive*: for any proposition P, we have $P \leq P$. To mark this feature, Fine labels the notion *weak* full ground. In terms of it, and its natural partial counterpart (weak *partial* ground, symbolized by $\preccurlyeq$), he then defines the more intuitive notion of *strict* full ground by imposing a kind of irreversibility requirement, namely that the grounded truth must not partially ground any of the grounding truths. It is the second, strict understanding of ground that we shall mainly focus on below.

Although the precise definition of strict full ground is thus somewhat complicated and indirect, it yields a clear and fairly intuitive picture. In effect, the truthmaker account recognizes two means by which we can move from a proposition to a strict full ground. The first is by (proper) *decomposition* of verifiers. For instance, consider a proposition {u} verified only by the state u. If u may be decomposed into two proper parts s and t, then the corresponding propositions {s}, {t} will jointly strictly fully ground the proposition {u}. The second means is by (strict) *specification*, i.e. eliminating some of the verifiers from a proposition, thereby moving from a proposition for which there are various ways in which it may be true to one which may be true in only some of these. For instance, if u and v are distinct states, then {u} strictly fully grounds $\{u, v, u \sqcup v\}$.[9]

In the case of *singular* grounding, i.e. cases in which a truth is fully grounded by a single truth rather than a plurality of them, no proper decomposition of the grounded proposition's verifiers occurs. So here, specification is the only available means to proceed from a proposition to a strict

[7]Cf. (Fine, 2012b, p. 8). The requirement of closure under fusion is the reason why the fusion of verifiers of the disjuncts in a disjunction must also be counted a verifier. In (Fine, 2017b, pp. 686 and 700ff.), a requirement of convexity is also imposed, which demands that whenever a state u lies between two verifiers s and t – that is, $s \sqsubseteq u \sqsubseteq t$ – then u is also a verifier. For present purposes, it does not matter whether convexity is imposed. For some reasons not to impose convexity, see (Krämer & Roski, 2015).

[8]For further discussion of non-factive ground, see e.g. (Fine, 2012a, pp. 48f).

[9]Of course, combinations of the two are possible, as when {s}, {t} jointly strictly fully ground {u, v} under the assumptions in the main text.

ground. Correspondingly, for singular grounding the conditions for weak full ground and strict full ground can be simplified:

($\leq_s$)	$P \leq Q$ iff $P \subseteq Q$
($<_s$)	$P < Q$ iff $P \subset Q$

This completes my exposition of the truthmaker framework for ground as proposed by Fine. I now turn to my objection.

3 | THE OBJECTION

I will argue that the truthmaker semantics yields too restrictive an account of the possible grounding structures. In particular, as indicated in the introduction, it excludes what I have described as *single-conduit* grounding structures:

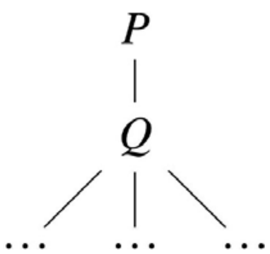

The lines here are to be taken to indicate relationships of strict full grounding, with the grounds lower than what they ground. So in the situation depicted, Q (strictly fully) grounds P, and any other ground of P grounds Q. Making use of the notion of weak full ground, we may summarize the key facts about the structure as follows:

(G.1)	$Q < P$
(G.2)	For all Γ, if $\Gamma < P$ then $\Gamma \leq Q$

Under the truthmaker account, (G.1) and (G.2) are jointly inconsistent. For assume $Q < P$. Then by ($<_s$), it follows that $Q \subset P$. So P must have a verifier which is not a verifier of Q. Call that verifier x, and let X be the proposition $\{x\}$, verified by x and only by x. Then $X \subset P$, and by ($<_s$) we have $X < P$. But since not $X \subseteq Q$, also not $X \leq Q$. So X is a counter-example to (G.2).

In more informal terms, the point is that any *singular* strict full ground must ground by *specification*, and any truth admitting of proper specification at all can be properly specified in at least two ways. So if a given proposition P has *any* singular strict full ground, it must have at least two:

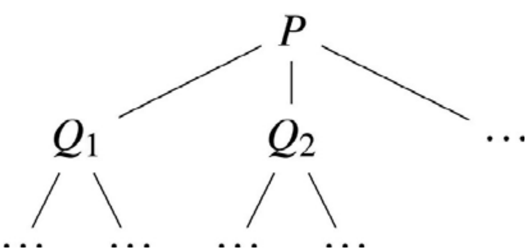

It is important to be clear what this result means. Truthmaker semantics provides an account of the *logic* of ground. If there is no model, within truthmaker semantics, in which a particular kind of grounding structure occurs, then this means that according to truthmaker semantics, the existence of such a structure is ruled out by the logic of ground. So, according to truthmaker semantics, the existence of single-conduit structures is ruled out by the logic of ground. This seems deeply problematic to me.

Before I further explain why, however, there is a subtlety that we should take note of. To characterize the single-conduit grounding structure, we need to *quantify* over all grounds of a given truth, as in (G.2). But within $\mathscr{L}_G$, we cannot formulate such quantifications. As a result, whether or not single-conduit grounding structures are allowed does not manifest itself in whether a particular set of sentences of $\mathscr{L}_G$ is consistent or not. Once we extend $\mathscr{L}_G$ by suitable quantificational devices, however, we can straightforwardly formalize (G.1) and (G.2), and the resulting sentences of our extended language will be jointly inconsistent under the truthmaker approach.

Why should we take this to be problematic? I think there are at least two good reasons. Firstly, there appears to be no *independent* motivation for taking single-conduit structures to be inconsistent. As far as I am aware, none of the principles that have been put forth in the extant literature on the logic of ground provide any reason for ruling out this structure. Indeed, discussion of the logic of ground has so far concerned itself almost exclusively with principles that can be stated within $\mathscr{L}_G$, or a language relevantly like it, which have no prospect of excluding single-conduit grounding. But absent positive reasons for rejecting a given grounding hypothesis as inconsistent with the very logic of ground, it seems the default view should be that the hypothesis is consistent relative to the logic of ground.

Secondly, there is also some direct evidence in favour of the consistency of single-conduit grounding. For there are plausible views which, if correct, provide *concrete instances* of single-conduit grounding. Perhaps the most compelling example concerns the grounds of the existence of a given singleton set, such as {Socrates}. For it seems to be a perfectly reasonable view that

(S.1)	Socrates exists $<$ {Socrates} exists
(S.2)	For all Γ, if $\Gamma <$ {Socrates} exists then $\Gamma \leq$ Socrates exists

For (S.1), note that one of the most frequently used examples of a plausible grounding claim is the claim that Socrates' existence grounds the existence of {Socrates} (cf. e.g. Correia & Schnieder, 2012, p. 14; Fine, 2015, p. 296; Bliss & Trogdon, 2016, §4).[10] In most cases, the authors leave unspecified whether they have in mind partial or full ground. But even this would at least seem to indicate that they do not consider its suitability as a paradigm example for a plausible grounding claim to depend on it being interpreted as a claim of partial ground. Moreover, (S.1) has also found explicit endorsement in the literature by Kelly Trogdon (2018, p. 1291).

As for (S.2), I am not aware that the claim has ever been discussed in print, let alone endorsed or rejected. But assuming (S.1), the only way for (S.2) to be false is for there to be a further grounding 'path' towards the existence of {Socrates}, in parallel to that via the existence of Socrates. Perhaps one could come up with potentially reasonable suggestions for such an additional grounding path. Perhaps one might have the view that while the existence of Socrates is sufficient to ground the existence of {Socrates}, there is another full ground of the existence of {Socrates} which includes not just the existence of Socrates, but also some sort of singleton-set-formation principle. But clearly one is not committed to countenancing such an additional grounding path purely because one assents to (S.2).

Aside from the grounds of singleton-existence claims, are there other plausible instances of single-conduit grounding?[11] I think so, though I am sceptical that anyone unconvinced by the example of {Socrates} will be swayed by any of the others. Let me mention just one more pertinent case, which concerns property-ascriptions. One may reasonable hold that true property-ascriptions – truths

[10]It is clear in each case that the example is not intended as a claim of weak ground.

[11]One may wonder if the same problem does not arise for sets with more members than one. In a sense, it does: just like the truthmaker account rules out that the truth that {Socrates} exists is strictly fully grounded by the truth that Socrates exists, it rules out that the truth that {Socrates, Plato} exists is strictly fully grounded by the truth that Socrates exists and Plato exists. But in this case, there is a related claim of strict full ground which *can* be true under the truthmaker account: that the truth that {Socrates, Plato} exists is *jointly* strictly fully grounded by the two truths that Socrates exists and that Plato exists. As a result, these examples are dialectically less compelling than the singleton case.

expressed by instances of 'a has the property of being F' – have the corresponding simple predication – expressed by an instance of 'a is F' – as their only immediate strict full ground, in the sense that every other ground has to pass through it.[12] If so, they, too, yield instances of single-conduit grounding.

It might be suggested that there are also examples of *logical* single-conduit grounding. For consider the following principles

$(L_\vee)$	$A < A \vee A$
$(L_\wedge)$	$A < A \wedge A$
(L)	$A < A$

It is a fairly common view that all instances of these schemata are true. Moreover, in all three cases, it also seems plausible that the ground A should be a sole conduit: all other grounds of $A \vee A$, $A \wedge A$, and A have to pass through A.[13] So the truthmaker account is incompatible with such a view. Indeed, as a number of authors have observed, under the usual truthmaker account, all of A, $A \vee A$, $A \wedge A$, and A have exactly the same verifiers, so $(L_\vee)$, $(L_\wedge)$, and (L) are not merely incompatible with the assumption of A as sole conduit, they are straightforwardly false.[14]

With respect to these examples, however, I think there is a plausible response available to the defender of the truthmaker account. Following Correia (2010), many grounding theorists distinguish between a *worldly* conception of ground on the one hand, and a *representational* conception of ground on the other. It is usually supposed that the former would be considerably less fine-grained than the latter, since it would be sensitive only to differences between truths that concern their relation to worldly items, rather than ones that concern purely representational differences. And cases such as $(L_\vee)$, $(L_\wedge)$, and (L) are commonly used to illustrate the difference between the two conceptions, the thought being that they are plausible only under a representational, but not under a worldly conception of ground (cf. Correia, 2010, pp. 267f), (Fine, 2017b, pp. 685f)). Since corresponding truths A and, say, $A \vee A$, plausibly represent the same worldly fact, under a worldly conception, one could not be held to ground the other without allowing us, absurdly, to draw the conclusion that each grounds itself.[15] With this in mind, the truthmaker semantics may then be put forward as adequate only relative to the worldly conception of ground – indeed, this is precisely what is done in (Fine, 2017b).

In light of this consideration, one might wonder if the same response could be given with respect to the other examples, concerning the existence of singleton sets and property-ascriptions. The idea would be that the truth that Socrates exists and the truth that {Socrates} exists represent the same worldly fact, and similarly the truth that Socrates is wise would be seen as representing the same worldly fact as the truth that Socrates has the property of being wise. However, it seems to me that this response is far less convincing in application to these non-logical cases. Firstly, in both cases the plausibility of the claim that the same worldly fact is represented turns on substantive and contentious issues in the philosophy of set theory and the metaphysics of properties, respectively. Other things being equal, we should not expect such issues to be decided by the correct

[12]A view like this is endorsed by (Fine, 2012a, p. 67).

[13]This is a straightforward consequence of the so-called *elimination rules* for strict ground proposed in (Fine, 2012a, p. 63ff).

[14]$(L_\vee)$, $(L_\wedge)$, and (L), or closely related principles, are endorsed, among others, in (Fine, 2012a), (Rosen, 2010), (Schnieder, 2011) and (Correia, 2014). Their failure under the truthmaker account is discussed, e.g., in (Krämer, 2018, p. 790), where I use it as motivation to develop a different kind of semantics for ground, as well as by Fine himself, e.g. in his (2017b, p. 685f) and (2012a, p. 74, n. 22).

[15]This is absurd because it would allow us to conclude for *every* truth that it strictly fully grounds itself.

logic of ground. Secondly, any considerations in support of identifying the relevant worldly facts, it seems to me, would have to about the specific subject matter under consideration, i.e. sets and properties. As such, while they may support doubts concerning these specific putative examples of single-conduit grounding, they do not seem to provide much support for the far stronger and more general claim that such grounding structures are excluded already *by virtue of the logical features* of ground. Crucially, the case against this claim developed above does not seem to depend on ground being thought of as representational rather than worldly: no one has ever proposed logical principles for worldly ground that would exclude single-conduit grounding, and the singleton- and property-examples are plausible, though contentious, under both conceptions of ground.

Admittedly, setting aside the question of single-conduit grounding, the truthmaker account of ground has much to be said for it. So it might be suggested that this fact itself gives us reason to accept the truthmaker account, and hence its implication that single-conduit grounding is incompatible with the logic of (worldly) grounding.[16] The plausibility of this suggestion depends, however, on whether this implication is an *essential* consequence of the desirable features of the truthmaker account. So the question is, *must* an alternative semantics that allows for single-conduit grounding be significantly less attractive in other respects, or less in line with a conception of ground as worldly? Or can we give a semantics that allows for single-conduit grounding and which (near enough) matches the truthmaker semantics in the other relevant respects? If we can, then the general virtues of truthmaker semantics do not give much support for the alleged inconsistency of single-conduit grounding. In the next section, I try to develop such an alternative semantics.[17]

4 | THE SOLUTION

In a nutshell, the solution I want to propose is to recognize a *third* way in which grounding connections can arise. Like decomposition, it may be described as involving the reduction of a verifier to something more basic. But whereas in the case of decomposition, the reduction of a state s is to a multiplicity of states $t_1, t_2, \ldots$ whose *fusion* is s, in the present case, the reduction of a state s is to a *single* state t. When s may in this way be reduced to the state t, I shall say that t *(strictly) generates* s. For want of another term, I shall call this additional source of grounding connections *(non-mereological) reduction*.

In more detail, the proposal is as follows. We begin by defining, on the basis of our intuitive notion of (strict full) ground, a binary relation of generation on the states. For any states s and t, we say that s *strictly generates* t ($s \Rightarrow t$) iff the proposition verified by s and only s strictly fully grounds the proposition verified by t and only t.[18] We say that s (weakly) generates t ($s \rightarrow t$) iff $s = t$ or s strictly generates

[16]Thanks to an anonymous referee for pressing me on this.

[17]It is worth noting that the situation seems to me quite different when we consider the putative *logical* instances of single-conduit grounding. There exist two worked-out proposals for a semantics of ground that validates ($L_\vee$), ($L_\wedge$), and (L), and renders them instances of single-conduit grounding, due to Correia (2017) and myself (Krämer 2018, 2019). Both are significantly more complicated than the truthmaker semantics, and more importantly, both are significantly less in line with a conception of ground as *worldly*. Correia's semantics is explicitly targeted at the representational conception of ground. My approach in these other papers is to see ground as sensitive to whether a worldly fact verifies a given proposition *by* verifying some other proposition, and thus to how a given proposition relates to other propositions. Truthmaker semantics, in contrast, renders ground sensitive only to how a proposition relates to its worldly verifiers. As I emphasize in my (2019, §4.5) this constitutes at least one good sense in which his approach and the truthmaker approach are on opposite sides of the worldly/representational divide.

[18]Given its intimate connection to grounding, why not just call generation 'grounding'? The reason I prefer to use a different term is that there are two state-level connections – fusion and generation – that bear the same sort of connection to grounding as a relation between propositions. Calling just one of them 'grounding' would suggest an asymmetry in their connection to propositional grounding which is not there. We might wish to describe both of them as relationships of grounding on the level of states, but since their behaviour is rather different in many ways, it seems best to reflect that difference in our terminology.

t. Given some assumptions about grounding, we may then establish parallel principles about generation. In particular, using widely accepted assumptions about grounding, we may establish that weak generation is a partial order, i.e. reflexive, transitive, and anti-symmetric.[19] We then adjust Fine's definition of weak full ground to reflect the idea that generation gives rise to grounding connections. According to the original definition, recall, a proposition P weakly fully grounds a proposition Q iff every verifier of P is *identical* to some verifier of Q. We now replace the appeal to identity by an appeal to (weak) generation. In the general case where we may have a plurality of grounds, the new definition then reads:

(≤*) $P_1, P_2, \ldots \leq Q$ iff $s_1 \sqcup s_2 \sqcup \ldots$ generates some verifier t of Q whenever $s_1, s_2, \ldots$ verify $P_1, P_2, \ldots$, respectively.

We do not make any changes to how the other notions of ground are defined in terms of weak full ground.

The resulting account can easily allow for single-conduit grounding, and thus for the joint truth of (S.1) and (S.2). To see this, let s be the state that Socrates exists, and let t be the state that {Socrates} exists. Let S be the proposition that Socrates exists, and let T be the proposition that {Socrates} exists. We shall assume that

(A.1) $S = \{s\}$

(A.2) $T = \{t\}$

Then in order to obtain (S.1), the claim that $S < T$, it suffices to make the assumptions that s strictly generates t, and that t is not a part of any state generating s:

(A.3) $s \Rightarrow t$

(A.4) For all states u: not $t \sqcup u \Rightarrow s$

(A.3) straightforwardly ensures that $S \leq T$. (A.4) implies that not $T \preccurlyeq S$, which by the definition of strict full ground yields that $S < T$.[20]

To also obtain (S.2), we need to make sure that any strict full ground of T passes through S. So firstly, we need to ensure that T cannot be grounded by decomposition. We do this by assuming that t is *prime*, in the sense that no fusion of only proper parts of t is identical to t:

(A.5) For all sets of states T: if $t = \bigsqcup T$ then $t \in T$

Finally, we need to make sure that all grounding of T by reduction goes via S. To that end, we assume that every state generating t does so via s:

(A.6) For all states u: if $u \Rightarrow t$ then $u \rightarrow s$

[19] That is, for all states s, t, u, we have $s \rightarrow s$ (reflexivity), $s \rightarrow u$ if $s \rightarrow t$ and $t \rightarrow u$ (transitivity), and $s = t$ if $s \rightarrow t$ and $t \rightarrow s$ (anti-symmetry). – We shall see in the appendix that some additional assumptions about → have to be made in order that we obtain the right logic of ground, but they enjoy a similar degree of independent plausibility.

[20] *Proof*: Suppose $T \preccurlyeq S$. Then by definition of ≼, there is a set of propositions Γ such that $\Gamma \cup \{T\} \leq S$. Then let X result from picking one verifier from each member of Γ. By definition ≤, it follows that $\bigsqcup(X \cup \{t\}) \rightarrow s$. But $\bigsqcup(X \cup \{t\}) = \bigsqcup X \sqcup t$, so it follows that $\bigsqcup X \sqcup t \rightarrow s$, contrary to (A.4).

Given the assumptions (A.1)–(A.6), we can derive the truth of both (S.1) and (S.2), and thus obtain that the grounding tree for the existence of {Socrates} instantiates the single-conduit structure.[21]

It should be noted, moreover, that assumption (A.1) is not essential to getting (S.1) and (S.2) to come out true. One might hold instead that there are many verifiers of the proposition that Socrates exists, corresponding perhaps to the many ways for there to be some simples arranged Socrates-wise. One then only needs to adjust (A.3), (A.4), and (A.6) accordingly, making sure these states are exactly the states that strictly generate t, and that t does not help generate any of them.[22]

It is also worth pointing out that the primeness assumption (A.5) is compatible with t having proper parts, and in particular with s being a proper part of t. It only rules out that s is a supplemented proper part of t, i.e. that there is a further proper part u of t such that $s \sqcup u = t$. But the standard assumptions in truthmaker semantics about the mereology of states allow for these kinds of unsupplemented proper parts.[23]

So the modification I propose does some good: it makes room for more plausible grounding structures than the original truthmaker semantics. It remains to show that it does not do more harm than good. Fortunately, it turns out that many, if not all, of the desirable features of the original account are preserved under the modification. The change in the logic of ground it brings about is rather *local*: allowing for single-conduit grounding does not force us to give up any of the logical principles for ground that are typically considered in the debate. In particular, as I show in the appendix, we can retain exactly the same pure and propositional logic of ground that the original truthmaker account yields. Moreover, the modified semantics seems to fit just as well as the original one with a conception fo ground as a *worldly* relation. For just like the original account, it renders ground sensitive only to how ground and groundee relate to their worldly verifiers, and how these relate to one another. The only change is that we are taking into account an additional relation of generation between the worldly verifiers themselves.

The modification does, of course, introduce some additional complexity into the theory. Instead of dealing just with the part-whole relation on the states, we now have to also deal with a second relation of generation, and the way the two relations interact. But as the appendix shows, the assumptions needed here are few, and they are fairly simple and straightforward counterparts to plausible assumptions about ground.[24] So I think the advantages that the modification offers are well worth the addi-

[21]*Proof*: Suppose $\Gamma < T$. Then $\Gamma \leq T$, and not $T \preccurlyeq Y$ for any $Y \in \Gamma$. We wish to show that $\Gamma \leq S$. So let $x = \bigsqcup X$ for any X obtained by picking one verifier from each member of Γ. Then what we need to show is simply that $x \rightarrow s$. Since $\Gamma \leq T$, $x \rightarrow t$. Now suppose for contradiction that $x = t$. Then $t = \bigsqcup X$, so by (A.5), $t \in X$. Hence $t \in Y$ for some $Y \in \Gamma$. But then $T \leq Y$, and hence $T \preccurlyeq Y$, contrary to our assumption. So since $x \rightarrow t$ and $x \neq t$, it follows that $x \Rightarrow t$. But then by (A.6), $x \rightarrow s$, as desired.

[22]If we assume the plausible seeming principle that no state generates any of its proper parts, we may actually derive (A.6) or its adjusted counterpart from the other assumptions. Thus, assume for contradiction that $t \sqcup u \rightarrow s$. By transitivity of $\rightarrow$ and (A.3), $t \sqcup u \rightarrow t$. Then by the principle just described, t cannot a proper part of $t \sqcup u$ and thus must be identical to $t \sqcup u$. But then $t \rightarrow s$, and hence by anti-symmetry of $\rightarrow$, $s = t$, in contradiction to (A.3). – It would be very interesting to examine what further principles might govern the interaction of $\rightarrow$ and $\sqsubseteq$, but this is something I shall have to leave for future work, although some related issues briefly come up in the formal appendix.

[23]Indeed, it is a natural hypothesis that s is a proper part of t, and perhaps even an unsupplemented proper part of t, whenever $s \Rightarrow t$. But as mentioned in the previous footnote, these issues call for a more extended discussion, which will have to wait for another occasion.

[24]It is perhaps worth nothing that one could also avoid working with two relations on states, and instead postulate just one relation of grounding between sets of states and states. The idea would be that the generation relation could be recovered as the special case of a singleton set grounding a state, and the fusion operation could be recovered by letting the fusion of a set of states be the least state, with respect to grounding, which is grounded by the given set of states. From the fusion operation, we could then recover parthood by letting $s \sqsubseteq t$ iff $s \sqcup t = t$. I am skeptical, however, that any real simplification can be achieved in this way. As far as I can see, the assumptions required about this kind of grounding relation on the states are considerably less simple and less natural than the ones needed for parthood and generation.

tional complexity. Moreover, the general picture of how grounding relations arise that the account offers is still a very simple and intuitive one: starting from derivative truths, we work towards more fundamental ones by *eliminating* a verifier, thereby obtaining a *more specific* description of reality, or by *reducing* a verifier, either to a generating state, or to a collection of states of which it is the fusion, thereby obtaining a description of *more basic elements* of reality.

It should be noted that there is a sense in which part of my proposal is a *generalization* of the truthmaker semantics, rather than a competitor to it. For we may distinguish between two components of the proposal. One is an adjusted definition of a truthmaker interpretation, or *model*, of the language of ground, and a definition of what it is for a sentence of that language to be true in a model:[25] in contrast to the original truthmaker models, ours include a relation of generation in addition to the part-whole relation on the states and the truth-conditions for grounding claims appeal to that additional relation. Now we may understand the notion of logical consequence as truth-preservation across all *admissible* models. There is then a separate question which of our models are admissible. In particular, there is the question whether any models that give rise to single-conduit structures are admissible. So the second component of the view I have proposed is a characterization of the admissible models under which we do obtain examples of single-conduit structures. But we can of course retain the first component of my proposal and consider other possible characterizations of the class of admissible models. For instance, if we consider admissible only models in which generation coincides with identity, we obtain exactly the same account of the logic of ground that the original truthmaker semantics gives rise to.

There is an instructive parallel to this situation in the theory of metaphysical modality. In the simplest form of possible world semantics for metaphysical modality, we do not appeal to any kind of accessibility relation on the worlds: necessity is simply understood as truth in all possible worlds. By introducing an accessibility relation, and taking necessity to be a matter of truth in all accessible worlds, we obtain a more general framework. It collapses into the old one under the assumption that every world is accessible from every other world. But it can also capture other possible views, under which the assumptions about accessibility are weaker. So without a claim concerning which models are admissible, we may see my proposal as standing to standard truthmaker semantics in the same way that possible world semantics with an accessibility relation stands to the simpler version without accessibility.

Finally, it might be objected that my modification does not preserve the *reductive* nature of the original account, since the novel element of the generation relation is explicitly defined in terms of a prior, intuitive concept of ground, which is taken as primitive. It is true that my proposed account cannot satisfy reductive ambitions. Along with most participants in the contemporary debate, I take the concept of ground to resist reductive analysis. But I maintain that even the original truthmaker semantics should not be regarded as offering any kind of conceptual reduction of ground (nor is there any indication that Fine meant to propose it in such a spirit). The most compelling reason for this is that the notion of exact truthmaking is so closely related to grounding that it seems implausible that our grasp of this notion is independent of our grasp of the notion of ground. Indeed, Fine himself suggests that truthmaking may be explained in terms of ground, rather than the other way round:

> *[I]ndeed, we might think of the notion of exact verification as being obtained through a process of ontological and semantic ascent from a claim of ground [to the effect that A_1, A_2, … ground C]. For we first convert the statements A_1, A_2, … into the corresponding*

[25]The precise definitions are given in the appendix.

facts f_1, f_2, … (that A_1, A_2, … obtain) and then take the sum f of the facts f_1, f_2, … to be an exact verifier for the truth of C. (Fine, 2017c, *§3)*[26]

But if the semantics cannot serve reductive ambitions, exactly what is it good for?[27] There are a number of important uses for a formal semantics which do not depend on its reductive potential. Some of them were already in evidence in this paper, others we have hinted at. A formal semantics provides the resources to define relations of *consequence* – truth-preservation across all (admissible) models – and *consistency* – truth in at least one (admissible) model. It thereby allows us to evaluate proposed deductive systems for the relevant languages for soundness and completeness. It may also allow us to connect the question of the consistency of a particular set of grounding claims with specific conditions on models. We can then see what impact it has on the logic if we allow or disallow such models, and thereby obtain further evidence for or against the consistency of the grounding claims in question. Thus, in this paper, we have connected the question of the consistency of single-conduit grounding structures, or sets of sentences characterizing them, to specific conditions on the behaviour of the generation and parthood relations on the states, and we have examined the impact on the logic of ground of allowing for models of the relevant kind.[28,29]

5 | CONCLUSION

My aim in this paper was to develop and defend a modification of the truthmaker semantics for ground as developed by Kit Fine. I first argued that Fine's account yields a problematically restrictive view of the possible grounding structures, by ruling out single-conduit structures: structures with unique, singular, immediate strict full grounds. The account thereby excludes as inconsistent with the very logic of ground an otherwise natural and attractive view of how the existence of a singleton set like {Socrates} is grounded in the existence of its member. I then described a modified version of the semantics which avoids this difficulty by countenancing an additional source of grounding relationships in the form of a relation of *generation* between truthmakers. At the same time, as I argued in the previous section, the modification preserves most of the desirable features of the original account. I conclude that the evidence available so far favours the view of single-conduit grounding as consistent, and hence my modification over the original truthmaker account. Moreover, even if we leave open this question of consistency, the proposed modification constitutes progress, since it provides a strictly more general formal framework within which we may articulate, study and compare competing

[26]The idea that Fine sketches here is developed in more detail in (Pleitz, ms).

[27]Thanks to an anonymous referee for urging me to elaborate on this. The question is of course as pressing for Fine as it is for me, given that the original truthmaker semantics also cannot satisfy reductive ambitions. Fine briefly addresses the point in his (2012b, p. 2), offering a similar account of the benefits of a formal semantics to the one I give.

[28]There are other examples in the literature in which the truthmaker semantics is put to this kind of role. For instance, Leuenberger (2019) discusses whether a truth can have a strict partial ground without having any strict full ground. He notes that within truthmaker semantics, this turns on whether we allow for states that have proper parts, but are prime in the sense defined above.

[29]Again, the modal analogy may be illuminating. By varying the constraints on the accessibility, a range of different candidate modal logics are obtained, and a number of interesting modal hypothesis are in this way connected to different such logics. For instance, Salmon (1989) has defended the view that some claims may be impossible, but possibly possible. This is incompatible with the modal logic S5, which is validated by the simple possible world semantics without accessibility. Within the semantics with an accessibility relation, the consistency of Salmon's examples turns on whether accessibility is allowed to be non-transitive, and we can study exactly what sort of modal logics we can obtain if we make this assumption.

accounts both of the general logic of ground, and of more specific questions concerning what grounds what.[30]

ORCID

Stephan Krämer https://orcid.org/0000-0003-1783-8742

REFERENCES

Bliss, R., & Trogdon, K. (2016). Metaphysical grounding. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition, https://plato.stanford.edu/archives/win2016/entries/grounding/

Correia, F. (2010). Grounding and truth-functions. *Logique et Analyse*, *53*, 251–279.

Correia, F. (2014). Logical grounds. *The Review of Symbolic Logic*, *7*, 31–59. https://doi.org/10.1017/S1755020313000300

Correia, F. (2017). An impure logic of representational grounding. *The Journal of Philosophical Logic*, *46*, 507–538. https://doi.org/10.1007/s10992-016-9409-7

Correia, F., & Schnieder, B. (2012). Grounding: an opinionated introduction. In F. Correia & B. Schnieder (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (pp. 1–36). Cambridge: Cambridge University Press.

Fine, K. (2012a). Guide to ground. In F. Correia & B. Schnieder (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (pp. 37–80). Cambridge: Cambridge University Press.

Fine, K. (2012b). The pure logic of ground. *The Review of Symbolic Logic*, *5*, 1–25. https://doi.org/10.1017/S1755020311000086

Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, *43*, 549–577. https://doi.org/10.1007/s10992-013-9281-7

Fine, K. (2015). Unified foundations for essence and ground. *Journal of the American Philosophical Association*, *1*, 296–311. https://doi.org/10.1017/apa.2014.26

Fine, K. (2017a). A theory of truthmaker content I: conjunction, disjunction and negation. *Journal of Philosophical Logic*, *46*, 625–674. https://doi.org/10.1007/s10992-016-9413-y

Fine, K. (2017b). A theory of truthmaker content II: subject-matter, common content, remainder and ground. *Journal of Philosophical Logic*, *46*, 675–702. https://doi.org/10.1007/s10992-016-9419-5

Fine, K. (2017c). Truthmaker semantics. In B. Hale, C. Wright, & A. Miller (Eds.), *A Companion to the Philosophy of Language* (pp. 556–577). Chichester: John Wiley & Sons Ltd.

Krämer, S. (2018). Towards a theory of ground-theoretic content. *Synthese*, *195*, 785–814. https://doi.org/10.1007/s11229-016-1242-6

Krämer, S. (2019). Ground-theoretic equivalence. *Synthese*, https://doi.org/10.1007/s11229-019-02154-4

Krämer, S., & Roski, S. (2015). A note on the logic of worldly ground. *Thought*, *4*, 59–68. https://doi.org/10.1002/tht3.158

Leuenberger, S. (2019). The fundamental: ungrounded or all-grounding? *Philosophical Studies*, Online First, https://doi.org/10.1007/s11098-019-01332-x

Pleitz, M. (ms). Making truthmakers. Unpublished manuscript.

Rosen, G. (2010). Metaphysical dependence: grounding and reduction. In B. Hale & A. Hoffmann (Eds.), *Modality: Metaphysics, Logic, and Epistemology* (pp. 109–136). Oxford: Oxford University Press.

Salmon, N. (1989). The logic of what might have been. *The Philosophical Review*, *98*, 3–34. https://doi.org/10.2307/2185369

[30]For helpful discussion and feedback, I would like to thank Singa Behrens, Martin Glazier, Stephan Leuenberger, Roberto Loss, Martin Pleitz, Stefan Roski, Louis deRosset, Josh Schechter, Moritz Schulz and Bruno Whittle. I am also grateful to an anonymous referee for this journal, whose comments led to substantial improvements to the paper. My work on this paper was funded by the Deutsche Forschungsgemeinschaft through an Emmy Noether grant for my project on *Relevance* (Grant KR 4516/2-1). I am very grateful for this support.

Schnieder, B. (2011). A logic for ‘because’. *The Review of Symbolic Logic*, *4*, 445–465. https://doi.org/10.1017/S1755020311000104

Trogdon, K. (2018). Grounding-mechanical explanation. *Philosophical Studies*, *175*, 1289–1309. https://doi.org/10.1007/s11098-017-0911-8

APPENDIX

In this appendix, I show that my proposed modification of the truthmaker semantics for ground does not lead to a change in either the pure or the propositional logic of ground. I do this by showing that for every interpretation of the language of ground under the old semantics there is an equivalent interpretation of the language under the new semantics, and vice versa.

First, we give the definition of a state-space, the basic structure in the original truthmaker semantics.[31]

Definition 1. A *state-space* is any pair (S, $\sqsubseteq$) such that

1. S is a non-empty set
2. $\sqsubseteq$ is a partial order on S such that every subset of S has a least upper bound[32] with respect to $\sqsubseteq$ in S

The least upper bound of any set T = {$t_1, t_2, \ldots$} states is their fusion $\bigsqcup T = t_1 \sqcup t_2 \sqcup \ldots$ Following Fine, we assume that even the empty set has a fusion, which we call the nullstate and denote by $\square$. It is easily verified that the nullstate is part of every state, and that $s \sqcup \square = s$ for all states s.

The basic structure of the modified semantics is what I call a *generation-space*:

Definition 2. A *generation-space* is any triple (S, $\sqsubseteq$, $\rightarrow$) such that

1. (S, $\sqsubseteq$) is a state-space
2. $\rightarrow$ is a partial order on S such that
   a. if $s_1 \rightarrow t_1$ and $s_2 \rightarrow t_2$ and … and $t_1 \sqcup t_2 \sqcup \ldots \rightarrow v$ then $s_1 \sqcup s_2 \sqcup \ldots \rightarrow v$
   b. if $s \rightarrow t \sqcup u$ then $s = s_1 \sqcup s_2$ for some s_1, s_2 with $s_1 \rightarrow t$ and $s_2 \rightarrow u$

Note the two new conditions (2.a) and (2.b) constraining the interaction between $\rightarrow$ and $\sqsubseteq$. (2.a) is a *cut* constraint, asserting a strong form of transitivity for generation, parallel to the cut rule in Fine’s pure logic of ground, which says that given $\Gamma_1 \leq A_1$, $\Gamma_2 \leq A_2$, …, and $A_1, A_2, \ldots \leq C$, we may infer that $\Gamma_1, \Gamma_2, \ldots \leq C$ (Fine, 2012b, p. 5). The second constraint (2.b) says that a state only gets to generate a fusion of two states by being a fusion of generators of the states being fused. Its counterpart on the level of propositions is the principle that a proposition only gets to weakly fully ground a conjunction by being the conjunction of weak full grounds of the conjuncts, a valid rule of the propositional logic of ground under the truthmaker account.

We note some useful facts about generation-spaces.

Proposition 1. If (S, $\sqsubseteq$) is a state-space, (S, $\sqsubseteq$, =) is a generation-space.

Proof: It suffices to show that the identity relation satisfies the conditions on $\rightarrow$ in a generation-space. Identity is obviously reflexive, transitive, and anti-symmetric, and hence a partial order.

[31]For the formal details of Fine’s truthmaker semantics, see the appendices of (Fine, 2017a,b). The latter paper also briefly discusses the application to ground. My presentation of the Finean account is based on these most recent papers rather than (Fine, 2012b). The differences between the different versions of Fine’s account are inessential for our purposes.

[32]A state $s \in S$ is an *upper bound* (w.r.t. $\sqsubseteq$) of $T \subseteq S$ iff $t \sqsubseteq s$ for all $t \in T$, and it is a *least* upper bound of T iff $s \sqsubseteq u$ for every upper bound u of T. It is routine to show that least upper bounds are unique if they exist.

Moreover, if $s_1 = t_1$ and $s_2 = t_2$ and . . . and $t_1 \sqcup t_2 \sqcup \ldots = v$ then clearly $s_1 \sqcup s_2 \sqcup \ldots = v$. Finally, suppose $s = t \sqcup u$. Then let $t = s_1$ and $u = s_2$ to show that there are states s_1, s_2 with $s_1 \rightarrow t$ and $s_2 \rightarrow u$. QED.

Proposition 2. Let $(S, \sqsubseteq, \rightarrow)$ be any generation-space, and let $s_1, t_1, s_2, t_2, \ldots$ be members of S. Then if $s_1 \rightarrow t_1$ and $s_2 \rightarrow t_2$ and …, also $s_1 \sqcup s_2 \sqcup \ldots \rightarrow t_1 \sqcup t_2 \sqcup \ldots$

Proof: From Reflexivity of $\rightarrow$ by an application of the constraint (2.a), setting $v = t_1 \sqcup t_2 \sqcup \ldots$ QED.

Definition 3. Let $(S, \sqsubseteq, \rightarrow)$ be any generation-space. For $T \subseteq S$, let the *closure under generation* T^G of T be the set $\{s \in S\colon s \rightarrow t \text{ for some } t \in T\}$.

Proposition 3. Let $(S, \sqsubseteq, \rightarrow)$ be any generation-space, and let T be a non-empty subset of S which is closed under non-empty fusion. Then T^G is also non-empty and closed under non-empty fusion.

Proof: Non-emptiness follows from the non-emptiness of T and the reflexivity of $\rightarrow$. For closure, suppose U is a non-empty subset of T^G. Then for each $u \in U$, there is a state $s_u \in T$ with $u \rightarrow s_u$. By closure of T, the fusion of all s_u is also in T. By proposition 2, the fusion of all $u \in U$ generates the fusion of all the s_u, and hence is also in T^G. QED.

Now given any language $\mathscr{L}$, let the *language of ground* $\mathscr{L}_G$ *over* $\mathscr{L}$ consist of all and only the expressions of the forms

- $\Gamma \leq C$
- $\Gamma < C$
- $A \preccurlyeq C$
- $A \prec C$

where A and C are formulas of $\mathscr{L}$., and Γ is any set of such formulas.

Consider first the so-called pure logic of ground, which studies only the structural features of ground, without attention to the internal makeup of the relata of ground. Let $\mathscr{L}^P$ be a non-empty set of atomic sentences. Given a state-space $(S, \sqsubseteq)$, say that $\mathscr{M}_O = (S, \sqsubseteq, I)$ is an *old model* of $\mathcal{L}^P_G$ iff I maps every member of $\mathscr{L}^P$ to a non-empty subset of S which is closed under (non-empty) fusion. Likewise, given a generation-space $(S, \sqsubseteq, \rightarrow)$, say that $\mathscr{M}_N = (S, \sqsubseteq, \rightarrow, I)$ is a *new model* of $\mathcal{L}^P_G$ iff I maps every member of $\mathscr{L}^P$ to a non-empty subset of S which is closed under fusion. We give the obvious truth-conditions for the four types of grounding statements in $\mathcal{L}^P_G$, using our adjusted clause for weak full ground for truth in a new model. For Φ a subset and φ a member of $\mathcal{L}^P_G$, say that Φ entails$_O$ (entails$_N$) φ iff φ is true in every old (new) model in which every member of Φ is true.

Lemma 4. For every old model there is a new model in which exactly the same sentences of $\mathcal{L}^P_G$ are true.

Proof: Let $\mathscr{M}_O = (S, \sqsubseteq, I)$ be an old model and let $\varphi \in \mathcal{L}^P_G$. Let $\mathscr{M}_N = (S, \sqsubseteq, =, I)$, By proposition 1, $\mathscr{M}_N$ is a new model. Since all forms of partial and strict ground are defined in the same way in terms of weak full ground for both old and new models, it suffices to consider the case in which φ is of the form $\Gamma \leq C$. Then if φ is true in $\mathscr{M}_O$, any fusion of verifiers of the members of Γ is a verifier of C. Since the generation relation of the new model is reflexive, it follows that any fusion of verifiers of the members of Γ generates a verifier of C, and hence that φ is true in $\mathscr{M}_N$. But if φ is not true in $\mathscr{M}_O$, then some fusion t of verifiers of the members of Γ is not a verifier of C. Since the generation relation of the new model is the identity relation, it follows that t does not generate a verifier of C, and hence that φ is not true in $\mathscr{M}_N$. QED.

Lemma 5. For every new model there is an old model in which exactly the same sentences of $\mathcal{L}^P_G$ are true.

Proof: Let $\mathscr{M}_N = (S, \sqsubseteq, \rightarrow, I)$ be a new model and let $\varphi \in \mathcal{L}^P_G$. Let J map every member A of $\mathscr{L}^P$ to $I(A)^G$, the closure under generation of I(A). From proposition 3, it follows that $\mathscr{M}_O = (S, \sqsubseteq, I)$ is an old model. Again, since all forms of partial and strict ground are defined in the same way in terms of weak full ground for both old and new models, it suffices to consider the case in which φ is of the form $\Gamma \leq C$. For explicitness, let us write φ as $A_1, A_2, \ldots \leq C$. If $s \in I(A)$, we say that s is a verifier$_I$ of A, and likewise for $_J$.

Suppose first that φ is true in $\mathscr{M}_N$, and let $s_1, s_2, \ldots$ verify$_J$ $A_1, A_2, \ldots$, respectively. We need to show that $s_1 \sqcup s_2 \sqcup \ldots$ verifies$_J$ C. By definition of J, each s_i generates a verifier$_I$ t_i of the corresponding A_i. Since φ is true in $\mathscr{M}_N$, the fusion of the t_i generates a verifier$_I$ of C. By the cut constraint (2.a), it follows that the fusion of the s_i generates a verifier$_I$ of C. By the definition of J, we can thus infer that the fusion of the s_i is a verifier$_J$ of C, and hence that φ is true in $\mathscr{M}_O$, as desired.

Suppose now that φ is not true in $\mathscr{M}_N$. Then let $s_1, s_2, \ldots$ be verifiers$_I$ of $A_1, A_2, \ldots$, respectively, such that the fusion of the s_i does not generate a verifier$_I$ of C. By definition of J, the fusion of the s_i is not a verifier$_J$ of C, while each s_i still verifies$_J$ the corresponding A_i. It follows that φ is not true in $\mathscr{M}_O$, as desired. QED.

Theorem 6. For Φ any subset and φ any member of $\mathcal{L}^P_G$, Φ entails$_O$ φ iff Φ entails$_N$ φ.

Proof: From the previous two lemmas. QED.

We now turn to the *impure*, propositional logic of ground. To avoid irrelevant distractions, we set aside negation, and focus purely on conjunction and disjunction.[33] First, we define the operations of conjunction and disjunction on propositions.

Definition 4. Let $(S, \sqsubseteq)$ be any state-space, and let P and Q be non-empty subsets of S closed under non-empty fusion. Then

- $P \wedge Q = \{s \sqcup t: s \in P \text{ and } t \in Q\}$
- $P \vee Q = (P \cup Q) \cup (P \wedge Q)$

Now let the language $\mathscr{L}^I$ be the closure of $\mathscr{L}^P$ under the connectives ∧ and ∨. Old and new models of $\mathcal{L}^I_G$ are just like their counterparts of $\mathcal{L}^P_G$, but with I extended to the complex formulas in $\mathscr{L}^I$ in the obvious way, letting $I(A \wedge B) = I(A) \wedge I(B)$ and $I(A \vee B) = I(A) \vee I(B)$.

As before, we define 'old' and 'new' entailment relations in terms of the corresponding classes of models. To extend theorem 6 to $\mathcal{L}^I_G$, it suffices to show that closure under generation distributes over conjunction and disjunction. More precisely:

Lemma 7. Let $(S, \sqsubseteq, \rightarrow)$ be any generation-space, and let P and Q be non-empty subsets of S closed under non-empty fusion. Then

1. $(P \wedge Q)^G = P^G \wedge Q^G$
2. $(P \vee Q)^G = P^G \vee Q^G$

Proof: For 1., suppose $s \in (P \wedge Q)^G$, so $s \rightarrow t \sqcup u$ for some $t, u \in P, Q$. By constraint (2.b) on →, there are states s_t, s_u with $s = s_t \sqcup s_u$ and $s_t \rightarrow t$ and $s_u \rightarrow u$. Then $s_t \in P^G$ and $s_u \in Q^G$, hence $s \in P^G \sqcup Q^G$. Conversely, suppose $s \in P^G \wedge Q^G$, so $s = t \sqcup u$ for some t, u generating verifiers of P, Q, respectively. By proposition 2, s generates the fusion x of these verifiers of P and Q. By the definition of conjunction, x verifies $P \wedge Q$, hence $s \in (P \wedge Q)^G$.

For 2., suppose $s \in (P \vee Q)^G$, so s generates a verifier of P, or a verifier of Q, or a verifier of $P \wedge Q$. In the first case, $s \in P^G$. In the second case, $s \in Q^G$. In the third case, by the reasoning before, $s \in P^G \wedge Q^G$. So by the definition of disjunction, in all three cases, $s \in P^G \vee Q^G$. Conversely, suppose $s \in P^G$

[33]The treatment of negation introduces technical complexities which are irrelevant to our present concerns. For some ways to deal with negation within truthmaker semantics, see (Fine, 2017a, pp. 629ff), (Fine, 2017a, pp. 634f, 658), and (Fine, 2014, pp. 554ff).

$\vee Q^G$. Then either $s \in P^G$, in which case s generates a verifier of P, and hence of $P \vee Q$, or $s \in Q^G$, in which case s generates a verifier of Q, and hence again of $P \vee Q$, or $s \in (P^G \wedge Q^G)$, in which case by the above reasoning, s generates verifier of $P \wedge Q$, and hence again of $P \vee Q$. So in all three cases, s $\in (P \vee Q)^G$. QED.

Theorem 8. For Φ any subset and φ any member of $\mathcal{L}^I_G$, Φ entails_O φ iff Φ entails_N φ.

Proof: From the obvious counterparts to lemmas 4 and 5. The proof of the first, that to every old model there is an equivalent new one, carries over without any changes. For the proof of the second, we construct the old model from the new one in the same way as before, letting J assign to each atomic members A of $\mathscr{L}^I$ the closure under generation of its interpretation I(A) in the new model. We then conclude from lemma 7 that for *every* formula A of $\mathscr{L}^I$, the interpretation J(A) in the old model equals $I(A)^G$. The result then follows by exactly the same reasoning as before. QED.

Synthese (2018) 195:785–814
https://doi.org/10.1007/s11229-016-1242-6

CrossMark

Towards a theory of ground-theoretic content

Stephan Krämer[1]

Received: 8 February 2016 / Accepted: 1 October 2016 / Published online: 12 October 2016

Abstract A lot of research has recently been done on the topic of ground, and in particular on the logic of ground. According to a broad consensus in that debate, ground is hyperintensional in the sense that even logically equivalent truths may differ with respect to what grounds them, and what they ground. This renders pressing the question of what we may take to be the ground-theoretic content of a true statement, i.e. that aspect of the statement's overall content to which ground is sensitive. I propose a novel answer to this question, namely that ground tracks *how*, rather than just *by what*, a statement is made true. I develop that answer in the form of a formal theory of ground-theoretic content and show how the resulting framework may be used to articulate plausible theories of ground, including in particular a popular account of the grounds of truth-functionally complex truths that has proved difficult to accommodate on alternative views of content.

Keywords Ground · Content · Logic of ground · Truthmaking

1 Introduction

Recently a lot of research has been devoted to *ground*—the relation, as Kit Fine has put it, 'of one truth holding *in virtue of* others' (Fine 2012c, p. 1)—and in particular to the broadly *logical* features of ground.[1] A distinctive feature of ground, according

[1] The pioneering contributions initiating this debate were Batchelor (2010), Correia (2010, 2014a), Fine (2010, 2012c, b), Rosen (2010) and Schnieder (2011). More recent work includes Correia (2014b), deRosset (2013), deRosset (2014), Krämer (2013), Krämer and Roski (2015), Krämer and Roski (2016), Litland (2013, 2016a) and Poggiolesi (2015).

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

1 Department of Philosophy, University of Hamburg, Von-Melle Park 6, 20146 Hamburg, Germany

to current consensus, is that it is *hyperintensional* in the sense that even logically equivalent truths may differ with respect to what grounds them and what they ground (cf. Correia and Schnieder 2012, p. 14). For instance, the truth that snow is white is taken to ground the truth that snow is white or snow is not white, but not the logically equivalent truth that grass is green or grass is not green. And the truth that snow is white or snow is not white in turn grounds the truth obtained by adding as a further disjunct the proposition that $2+2=5$, which is not grounded by the logically equivalent truth that grass is green or grass is not green.

Ground is accordingly sensitive to features of a truth that go beyond its logical profile, and in particular beyond the matter of what possible worlds the truth obtains in. One pressing question in the development of the theory of ground is therefore just what the features of a truth are that are tracked by ground. For once we have an answer to this question, we may form a notion of *ground-theoretic content* by abstracting from those features of a truth to which ground is blind. We can then go on to construct a mathematical representation of ground-theoretic content, which may serve as a common framework within which to formally articulate, study, and compare the various competing views of ground.

In this paper I propose and defend a novel answer to the question what ground is sensitive to. My proposal builds on, and modifies, the view implicit in the semantics of ground that Fine has developed in his influential papers 'The Pure Logic of Ground' and 'Guide to Ground' (Fine 2012b, c). Central to this view is the notion of a *fact* – roughly, a proper or improper part of the actual world – *verifying* a truth. The view then presents the relationship of ground as sensitive purely to mereological relationships between the facts that verify the relevant truths. Unfortunately, as Fine concedes, the view is limited in its capability to serve as a framework for the articulation of plausible theories of ground. For it cannot accommodate certain widely held principles about the interaction of ground with the truth-functional operations of conjunction, disjunction, and negation.

I argue that we can overcome this problem by means of a very natural modification of Fine's approach. The key element of my proposal is the notion of a *mode of verification*, which corresponds to a certain kind of answer to the question *how* a truth is verified by a fact. A disjunction $P \vee Q$, for example, may plausibly be verified either by verifying its first disjunct P, or by verifying its second disjunct Q, and if P and Q are distinct propositions, then these modes of verifying $P \vee Q$ will be distinct as well. I suggest that it is these features of a truth, the modes in which it is verified, that ground tracks. In particular, I take some truths $P_1, P_2, \ldots$ to ground a truth Q just in case Q is verified by verifying $P_1, P_2, \ldots$[2]

The plan for the paper is as follows. Section 2 clarifies and motivates the project of the paper. Section 3 briefly describes Fine's proposal and its limitations. In Sect. 4 I informally introduce the notion of a mode of verification. I go on to show how a plausible account of how conjunctions, disjunctions, and negations are verified combines with the proposed view of ground to validate the principles that Fine was unable to

[2] A word on notation. I shall be somewhat sloppy in my use of the letters 'P', 'Q', etc., in that I sometimes use them as schematic sentence letters, and sometimes as variables ranging over contents that may be assigned to sentences.

accommodate. This constitutes a partial vindication of my proposal. A long penultimate Sect. 5 describes a mathematical representation of ground-theoretic content to serve as a formal framework for developing the theory of ground. I then use the framework to define a relation of ground and operations of conjunction, disjunction, and negation. The emerging view is again shown to validate the principles linking ground and the truth-functions that are invalid on Fine's account. Finally, I show how various competing views of the structural principles of ground, as well as of ground-theoretic equivalence, may be implemented within my formal framework. Section 6 concludes.

2 Preliminaries

Following Fine, I have informally spoken of ground as a relation between truths. It is controversial, however, whether this is the best, most perspicuous, or most fundamental way of speaking about ground. There are two worries about this, one targeting the term 'relation', one concerning the term 'truths'. I discuss them in turn.

A number of authors prefer to use a *sentential operator* to express ground, rather than a relational predicate such as 'ground(s)'.[3] Using $<$ as the symbol for ground, they would therefore write 'the ball is round $<$ the ball is red or round' rather than 'the truth that the ball is round grounds the truth that the ball is red or round'. Speaking in this way, it may be argued, no more commits one to a relation of ground holding between some truths than the use of 'if and only if' commits one to a relation of equivalence holding between some propositions.

However, even on this view, it is highly desirable to have a theory of ground-theoretic content of the kind I am after. Let it be granted that there is a legitimate notion of the overall *content* of a sentence. Then we may ask with respect to this notion which of the features of a sentence's content the grounding operator $<$ is sensitive to in the sense that sameness with respect to these features of content guarantees that two sentences may replace one another within the scope of $<$ without changing the truth value.[4] Abstracting again from any other features of a sentence's content, we get our notion of ground-theoretic content. The ground-theoretic content of a sentence will then be exactly what is suited to play the role of a sentence's *semantic value* in formulating a semantics for ground. A formal framework within which to theorize about ground-

[3] The operator option is chosen, for example, by Fine (2012b), Correia (2010) and Schnieder (2011). The predicate option is preferred by Rosen (2010), as well as Schaffer (2009). (The latter is something of an outlier in the current debate, though, in that he takes ground to relate not just truths, but objects of *any* kind. His conception of ground will not be canvassed in this paper.)

[4] I assume here that ground is sensitive only to differences between sentences that concern content. This is not obvious prior to investigation; it may be that the best account of the distinctions drawn by ground sees (some of) them as purely *syntactic*. For the purposes of this paper, the assumption has the status of a working hypothesis. That is, I propose that we try and see if we can make sufficiently fine-grained distinctions pertaining to content to capture the distinctions drawn by ground. (One potentially problematic kind of case arises in connection with conceptual analyses. It might be suggested that if Alice is a vixen, say, then this is so *because* Alice is a female fox, where the 'because' indicates grounding, and that nevertheless 'Alice is a vixen' and 'Alice is a female fox' are exactly alike in content. For discussion, see Schnieder (2010). Thanks to an anonymous referee for highlighting the relevance of these cases.)

theoretic content thus provides a framework for the development of semantic theories of ground.

We turn to the second worry concerning the talk of ground as a relation between truths, targeting the use of the term ‘truth’. A truth, presumably, is some kind of (accurate) representation of the world, and so ground, if it is a relation between truths, would appear to be a meta-representational relation. A number of authors however prefer to see ground as a relation between *worldly* items rather than representational ones, for which the term ‘fact’ would then appear more appropriate than ‘truth’ (cf. Correia 2010, p. 258f). An alternative option, preferred by Correia in later work, is to simply distinguish two legitimate notions of ground, one worldly, one representational (cf. Correia 2014a, Sect. 5, 2014b, p. 36).[5] I myself do not find the worldly/representational distinction very helpful or clear. It should be noted, though, that nothing in the way I have introduced the notion of ground-theoretic content hinges on the use of the term ‘truth’ rather than ‘fact’, or on taking the relata of ground to be representational entities.[6] So even if the only legitimate notion of ground is somehow worldly, this by itself threatens neither interest nor feasibility of my project.

A clearer distinction than that between worldly and representational items obtains between two alternative kinds of views on the *logical* principles for ground which Correia takes to be characteristic of worldly and representational conceptions of ground. On one kind of view, the following *introduction* principles for ground are taken to hold without restriction:[7]

($<$I$\vee$) If P then $P < P \vee Q$ and $P < Q \vee P$
($<$I$\wedge$) If P, Q then $P, Q < P \wedge Q$
($<$I$\neg\wedge$) If $\neg P$ then $\neg P < \neg(P \wedge Q)$ and $\neg P < \neg(Q \wedge P)$
($<$I$\neg\vee$) If $\neg P, \neg Q$ then $\neg P, \neg Q < \neg(P \vee Q)$
($<$I$\neg\neg$) If P then $P < \neg\neg P$

Now on all of the extant logics of ground, *typical* instances of the first four of these principles are taken to hold. But opinions diverge with respect to some special kinds of instances. Consider in particular the case in which $P = Q$. If the principles hold even for this case, then for any truth or fact P, we have that P grounds $P \vee P$ as well as that P grounds $P \wedge P$. Given the standard assumption that ground is not generally reflexive, it follows that a distinction must be made between P on the one hand, and $P \vee P$ as well as $P \wedge P$ on the other hand. But such a distinction, Correia claims, is only plausible given a conception of the relata of ground as representational entities,

5 Note that the distinction, if it can be made, may also be transposed to the sentential-operator setting, where it turns into a distinction between worldly and representational conceptions of ground-theoretic content; cf. Correia (2010, p. 257).

6 It may be objected that I treat ground, in effect, as a relation between ground-theoretic contents, and since contents are representational entities, this commits me to a representational conception of ground. If this were so, the same reasoning would reveal Correia’s ostensively worldly conception of ground as representational, for he, too, treats ground in effect as a relation between what he calls the worldly contents of sentences. In his intended sense, then, a content may be worldly and thus non-representational. This is one of the reasons why I do not find this way of making the distinction very helpful.

7 This kind of view is endorsed, for example, by Fine (2012b), and in related forms by Rosen (2010), Schnieder (2011) and Correia (2014b).

and he takes the same view with respect to the distinction required by the last principle between a truth P and its double negation $\neg\neg P$. On a worldly conception of ground, he concludes, (<I$\neg\neg$) is to be rejected, and the other principles have to be restricted to rule out, at least, the case of $P = Q$.[8]

So let us set aside the worldly/representational distinction and instead focus directly on the introduction principles. Should we take them to hold in unrestricted form? Or should they be restricted in some way? Or does it simply depend on which of two equally legitimate notions of ground is at issue? Rather than try to decide the issue in advance of developing a theory of ground-theoretic content, I suggest that well-founded answers should be based in part on how the possible views can be implemented within such a theory. And so the question arises whether each view can be accommodated within an attractive theory of ground-theoretic content.

A partial answer to this question is contained in Fine's work, which provides us with an elegant theory of ground-theoretic content that naturally accommodates the view that the introduction principles should be restricted. For on that account and given Fine's definition of ground, certain relatively natural restrictions of the introduction principles turn out to hold. The question is then whether it is also possible to formulate an alternative account that equally naturally accommodates the view that the introduction principles should not be restricted; as Fine himself points out, his own account cannot serve in this role. My aim in what follows is therefore to develop such a theory of ground-theoretic content and thereby to establish a positive answer to this second question.

3 The truthmaker conception of content

Fine (2012b) formulates a semantics for ground, and implicit in that semantics is, firstly, an answer to the question to what features of a sentence's content ground is sensitive, and secondly, a theory of ground-theoretic content. This section briefly summarizes Fine's account and identifies its limitations hinted at above.

The semantics Fine proposes associates with each sentence A two sets of facts, namely the set $[A]^+$ of facts verifying the sentence and the set of facts $[A]^-$ falsifying the sentence. Facts are described by Fine in comparison with possible worlds: they are like possible worlds except in that they are all of them *actual*, and in that they are generally *incomplete* in the sense that they leave open the truth-value of some propositions. Moreover, it is assumed that facts may be *fused*, so that given any facts $s, t, u, \ldots$, we may form their fusion $\bigsqcup\{s, t, u, \ldots\} = s \sqcup t \sqcup u \sqcup \ldots$, which contains all of s, t, u as parts. For a sentence to be true is for it to be verified by some fact, and for it to be false is for it to be falsified by some fact. So if we assume that no sentence is both true and false and no sentence is neither true nor false, then for every sentence A, exactly one of $[A]^+$ and $[A]^-$ is empty.

A content P, on this view, may accordingly be identified with a pair $\langle P^+, P^-\rangle$ of two sets of facts, where any member of P^+ is a verifier of P, and any member of P^- is a falsifier of P. Fine assumes that any fusion of verifiers of a sentence also verifies

[8] Cf. Correia (2010, p. 267f)—it is plausible that a number of other cases should then also be disallowed; cf. ibid, p. 269.

the sentence and similarly for falsifiers. So both components of a content are required to be closed under fusion.

Fine then formulates semantic clauses for four distinct grounding operators; the corresponding relations on contents may be defined as follows:[9]

$$\begin{array}{ll} P_1,\ P_2, \ldots \leq Q & \text{iff } s_1 \sqcup s_2 \sqcup \ldots \in Q^+ \text{ whenever } s_1 \in P_1^+,\ s_2 \in P_2^+,\ \ldots \\ P \preceq Q & \text{iff } P,\ R_1,\ R_2, \ldots \leq Q \text{ for some } R_1,\ R_2, \ldots \\ P_1,\ P_2, \ldots < Q & \text{iff } P_1,\ P_2, \ldots \leq Q \text{ and } Q \npreceq P_1,\ Q \npreceq P_2,\ \ldots \\ P \prec Q & \text{iff } P \preceq Q \text{ and } Q \npreceq P \end{array}$$

Somewhat curiously, the most basic notion of ground of the four is not the familiar notion of ground but a *reflexive* notion that Fine calls *weak full ground* ($\leq$). In terms of this notion and its natural *partial* counterpart ($\preceq$), Fine then defines an irreflexive notion of *strict full ground* ($<$), as well as a notion of *strict partial ground* ($\prec$).

Weak full ground, and thereby ultimately every notion of ground, is defined essentially in terms of the notion of fusion and verification. Ground, on this picture, is thus taken to track mereological relationships between the facts verifying the relevant truths, and ground-theoretic content is accordingly taken to be a matter of which facts verify a given truth.

Given this view of ground-theoretic content, however, it does not appear possible to accommodate the unrestricted introduction principles. To see this, we need to ask how the operations of conjunction, disjunction, and negation may be defined on Fine's contents. The view adopted in Fine (2012b) corresponds to these definitions:

$$\begin{array}{l} (P \wedge Q)^+ = \{s \sqcup t : s \in P^+ \text{ and } t \in Q^+\} \\ (P \wedge Q)^- = P^- \cup Q^- \cup \{s \sqcup t : s \in P^- \text{ and } t \in Q^-\} \\ (P \vee Q)^+ = P^+ \cup Q^+ \cup \{s \sqcup t : s \in P^+ \text{ and } t \in Q^+\} \\ (P \vee Q)^- = \{s \sqcup t : s \in P^- \text{ and } t \in Q^-\} \\ (\neg P)^+ \quad = P^- \\ (\neg P)^- \quad = P^+ \end{array}$$

But this implies that for all P, $P = P \wedge P = P \vee P = \neg\neg P$, and accordingly we never have $P < P \wedge P$, $P < P \vee P$, or $P < \neg\neg P$. Moreover, even independently of these specific clauses, it is hard to see how the problem could be solved within the present framework. For there appears to be no plausible way to distinguish between a truth P on the one hand and the corresponding truths $P \vee P$, $P \wedge P$, or $\neg\neg P$ on the other hand, appealing only to the facts verifying them, and in such a way as to generally render the relevant grounding claims true. The most obvious way to allow for some sort of distinction between P and $P \vee P$ as well as $P \wedge P$ would be to lift the requirement that P be closed under fusion. But this would not help for the case of $\neg\neg P$, and even for $P \wedge P$ and $P \vee P$, it would of course help only in the special case where P is not closed under fusion. If we wish to accommodate the envisaged introduction principles in unrestricted form, we need to modify the Finean framework.[10]

[9] Although Fine does not say so, the definitions should be read as restricted to truths, otherwise $P_1,\ P_2,\ \ldots \leq Q$ will vacuously hold whenever one of $P_1,\ P_2, \ldots$ is false.

[10] Note that I am relying here on the assumption, noted in Footnote 4 above, that ground is sensitive only to differences in content. If this assumption were given up, and ground seen as sensitive to the linguistic guises of contents, then the above line of reasoning could be resisted. Thanks here to an anonymous referee.

4 Introducing modes of verification

To see how a suitable conception of ground-theoretic content might be obtained, we adopt for the moment an informal understanding of propositions and ask what distinctions we can then draw between a proposition P and the corresponding disjunction $P \vee P$ that might be tracked by the relation of ground.

Consider first the easier case of the disjunction $P \vee Q$ of the distinct propositions P that Alice is smart and Q that Bob is tall. Suppose the fact s of Alice having an IQ of 150 verifies P and the fact t that Bob is 6 foot tall verifies Q. Then both s and t separately verify $P \vee Q$. But it seems natural to think that there is a *difference* between the verification of $P \vee Q$ by s and the verification of $P \vee Q$ by t, which we can bring out by asking *how* the facts verify the proposition. The fact s verifies the proposition $P \vee Q$ by verifying the disjunct P that Alice is smart. The fact t, in contrast, verifies $P \vee Q$ by verifying its other disjunct Q that Bob is tall.

Thus attending to the ways, or *modes*, in which a proposition is verified by its verifiers also enables us to distinguish between P and $P \vee P$. For it is very natural to say that s verifies $P \vee P$ just as it verifies $P \vee Q$, by verifying P. At the same time, it is not at all tempting to say that s verifies P itself by verifying P. I would suggest, moreover, that it is this distinction which is tracked by the grounding relation. Roughly speaking, P grounds $P \vee P$ precisely because $P \vee P$ may be verified by verifying P, whereas P does not ground P, since P may not be verified by verifying P.

Before we move on to consider the case of conjunction, a brief comment on my use of ‘by’ is required to avoid misunderstanding. As used in ordinary language, a ‘by’-statement can be true even if it provides only a *partial* answer to the corresponding ‘how’-question. For instance, we may say that someone got into the house by breaking a window without thereby implying that breaking the window was on its own sufficient for the person to get into the house. Even if it was also required that the person climb through the window, the ‘by’-statement may express a truth by ordinary language standards. I wish to highlight, therefore, that the uses of ‘by’ I make in the formulation of my proposal, in contrast, *are* to be understood as imposing the kind of *fullness* or *sufficiency* condition that is lacking in the ordinary reading of ‘by’.[11]

We turn now to the case of P and $P \wedge P$. Again, we first consider the easier case of the conjunction $P \wedge Q$ of the distinct propositions P and Q. Since s verifies P and t verifies Q, their fusion $s \sqcup t$ verifies $P \wedge Q$. But *how* does $s \sqcup t$ verify $P \wedge Q$? A natural first thought is: by verifying P and Q. However, on the perhaps most natural, *distributive* reading of that suggestion, it implies that $s \sqcup t$ verifies *each* of these two propositions. And this is typically not the case, for on the Finean view, a verifier of a conjunction does not in general verify each, or even just one of the conjuncts.[12] The difficulty can be resolved by employing a *non-distributive* interpretation of the claim

[11] Cf. Schnieder (2008, p. 665f). As Schnieder points out elsewhere (cf. his 2011, p. 450f) similar cautionary remarks are in order when locutions such as ‘because’ are used to convey relationships of full ground.

[12] This illustrates that Fine’s notion of verification is non-monotonic in the sense that a fact may fail to verify a proposition even though it has a part which verifies the proposition. For this reason, Fine sometimes describes this notion of verification, which he also calls *exact* verification, as imposing a requirement of *holistic relevance*: for a fact to verify a given proposition, it must not contain any part which is irrelevant to the truth of the proposition (cf. Fine 2012a, p. 234; 2014, p. 551f et passim).

that the propositions P and Q are verified by a fact like $s \sqcup t$.[13] The idea is that in the relevant sense, the propositions $P_1, P_2, \ldots$ are verified by a fact iff the fact is the fusion of some facts $s_1, s_2, \ldots$ verifying $P_1, P_2, \ldots$ respectively. Since s, t verify P and Q, respectively, their fusion $s \sqcup t$ verifies, in the non-distributive sense, P and Q. Moreover, I claim that it is *by* verifying P and Q that $s \sqcup t$ verifies $P \wedge Q$.[14]

So by taking into account the modes in which a proposition is verified, we are also able to distinguish between P and $P \wedge P$. For while s verifies $P \wedge P$ by verifying P, it is not the case that s verifies P by verifying P. And as before, P may be taken to ground $P \wedge P$ on the strength of the fact that $P \wedge P$ may be verified by verifying P, whereas P does not ground P, since P may not be verified by verifying P. More generally, the account of ground I propose may in a first approximation be stated thus:

($<$) $Q_1, Q_2, \ldots < P$ iff
- (i) $P, Q_1, Q_2, \ldots$ are true, and
- (ii) every verifier of $Q_1, Q_2, \ldots$ verifies P by verifying $Q_1, Q_2, \ldots$

Consider now finally the negative proposition $\neg Q$ that it is not the case that Bob is tall. How should we take it to be verified? Like Fine, I adopt a *bilateral* conception of content to account for negation. On this conception, a content encodes information concerning both how it is verified and how it is falsified. So in giving an account of $\neg Q$, we may appeal to both what verifies its negatum Q, and what falsifies it. Again with Fine, I take it that $\neg Q$ is verified by exactly those facts that falsify Q, and that $\neg Q$ is falsified by exactly those facts that verify Q. For instance, the fact t' that Bob is 5 foot tall would plausibly falsify the proposition Q that Bob is tall, and hence verify $\neg Q$, whereas the fact t that Bob is 6 foot tall verifies Q, and therefore falsifies $\neg Q$.

We use again the $+/-$-notation to talk about bilateral contents and their components. That is, for P a bilateral content, P^+ is the positive component of P, and P^- is the negative component, so that $P = \langle P^+, P^- \rangle$. The notion of verification is applied with harmless ambiguity to both bilateral contents and their (unilateral) positive and negative components. Verification of a bilateral content P is defined as verification of P^+, and falsification of P as verification of P^-.

In giving an account of negation, I now need to say *how* a negation is verified by its verifiers, and falsified by its falsifiers. For each of these questions, two views are prima facie possible. I call them the *by-view* and the *identity-view*, respectively. So we have to consider four views with respect to $\neg P$:

(BV) The by-view of verification: $\neg P$ is verified by falsifying P.
(BF) The by-view of falsification: $\neg P$ is falsified by verifying P.
(IV) The identity-view of verification: $\neg P$ is verified how P is falsified.

[13] The non-distributive notion of verification has also been put to use in Litland (2016b) in developing a logic of a many-many notion of grounding, on which what is grounded is irreducibly a collection of truths or facts.

[14] The case of conjunction makes especially clear why the ‘full’ interpretation of ‘by’ must be assumed. For suppose that every fact s that verifies P also verifies Q and therefore $P \wedge Q$. By ordinary standards, it might then be true to say that every verifier s of P verifies $P \wedge Q$ by (among other things) verifying P. But it would not therefore be true to say that $P < P \wedge Q$. We avoid this result if we read ‘by’ as requiring fullness. For on such a reading, it is false that in general every verifier of P verifies $P \wedge Q$ by verifying P, since verifying P is usually only part of what is required for verifying $P \wedge Q$.

(IF) The identity-view of falsification: $\neg P$ is falsified how P is verified.

For example, according to (BV), the fact t' would verify the proposition $\neg Q$ that Bob is not tall by falsifying its negatum Q, the proposition that Bob is tall. More generally, any fact verifying Q^-, by verifying Q^-, verifies $(\neg Q)^+$. And according to (BF), the state t that Bob is 6 foot tall falsifies $\neg Q$ by verifying Q. More generally, any verifier of Q^+, by verifying Q^+, verifies $(\neg Q)^-$. According to (IV), in contrast, $\neg Q$ is verified in exactly the ways Q is falsified, so that $(\neg Q)^+ = Q^-$. And on (IF), $\neg Q$ is falsified in exactly the ways Q is verified, so $(\neg Q)^- = Q^+$.

In principle, any view of verification can be consistently combined with any view of falsification, so we obtain a total of four possible views of how the bilateral content of a negation is determined: (BV + BF), (BV + IF), (IV + BF), and (IV + IF). Some of the views do not fit our desired principles for ground, though. Firstly, if we are to obtain that $P < \neg\neg P$, then it must hold that the double negation $\neg\neg P$ of a given content P may be verified *by* verifying P. But this rules out (IV + IF), since that view implies $P = \neg\neg P$. (IV + BF) and (BV + IF), in contrast, directly imply this result. For assume that s verifies P^+. Suppose we accept (BV + IF). By (IF), $(\neg P)^- = P^+$, so s verifies $(\neg P)^-$ in whatever ways it verifies P^+. By (BV), it follows that s verifies $(\neg\neg P)^+$ by verifying $(\neg P)^- = P^+$. But that just means that s verifies $\neg\neg P$ by verifying P, as desired. Now consider (IV + BF). By (BF), s verifies $(\neg P)^-$ by verifying P^+. By (IV), $(\neg\neg P)^+ = (\neg P)^-$, so it follows that s verifies $(\neg\neg P)^+$ by verifying P^+, and thus s verifies $\neg\neg P$ by verifying P, as desired.

The situation is more complicated in the case of (BV+BF). Assume s verifies P^+. Then by (BF), s verifies $(\neg P)^-$ by verifying P^+. Moreover, by (BV), s verifies $(\neg\neg P)^+$ by verifying $(\neg P)^-$. Now, nothing we have said so far allows us to conclude that s verifies $(\neg\neg P)^+$ by verifying P^+. Nevertheless, it seems plausible that ‘by’, at least in its pertinent ‘full’ use, is *transitive* in the sense that if s ϕs by ψing, and ψs by χing, then s ϕs by χing. If so, then we may still infer that s verifies $(\neg\neg P)^+$ by verifying P^+. So it seems that all three remaining views will yield the desired mode of verification for the double negation $\neg\neg P$, and correspondingly that $P < \neg\neg P$.

The differences between the views emerge more clearly with respect to the other kinds of negative propositions, namely negations of conjunctions and disjunctions. Recall the above introduction principles for ground governing these:

($<$I$\neg\wedge$) If $\neg P$ then $\neg P < \neg(P \wedge Q)$ and $\neg P < \neg(Q \wedge P)$
($<$I$\neg\vee$) If $\neg P, \neg Q$, then $\neg P, \neg Q < \neg(P \vee Q)$

If we are to accommodate them, we must have that $\neg(P \wedge Q)$ may be verified by verifying $\neg P$ and by verifying $\neg Q$, and that $\neg(P \vee Q)$ may be verified by verifying $\neg P$, $\neg Q$. It turns out that this demand favours (IV+BF), the combination of the identity-view of the verification of a negation with the by-view of its falsification.

To see this, we first need to look at how conjunctions and disjunctions are falsified. Note that on the truthmaker conception, there is a strong analogy between the falsification of conjunctions and the verification of disjunctions, as well as between the falsification of disjunctions and the verification of conjunctions. Specifically, the negative content of a conjunction relates to the negative contents of its conjuncts like the positive content of a disjunction relates to the positive contents of its disjuncts. And similarly, the negative content of a disjunction relates to the negative contents of its disjuncts like the positive content of a conjunction relates to the positive contents of its conjuncts. It is natural to continue the analogy from the case of *what* verifies or falsifies to *how* it does so. In particular, just as $(P \vee Q)^+$ is verified via P^+ and via Q^+, we shall take $(P \wedge Q)^-$ to be verified via P^- and via Q^-, and just as $(P \wedge Q)^+$ is verified via $P^+,\ Q^+$, we shall take $(P \vee Q)^-$ to be verified via $P^-,\ Q^-$.

Now consider the claim that every verifier of $\neg P$, by verifying $\neg P$, verifies $\neg(P \wedge Q)$. First, we show that (IV + BF) implies this claim. Assume that s verifies $\neg P$. Then s verifies P^-, and so verifies $(P \wedge Q)^-$ by verifying P^-. By (IV), $(\neg(P \wedge Q))^+ = (P \wedge Q)^-$, so it follows that s verifies $(\neg(P \wedge Q))^+$ by verifying P^-. Moreover, we have $P^- = (\neg P)^+$, so it follows that s verifies $(\neg(P \wedge Q))^+$ by verifying $(\neg P)^+$, as desired. But now consider (BV + IF). If s verifies $(\neg P)^+$, then s verifies P^-, so s verifies $(P \wedge Q)^-$ via P^-. Given (BV), we have that s verifies $(\neg(P \wedge Q))^+$ via $(P \wedge Q)^-$. Assuming transitivity, we could infer that s verifies $(\neg(P \wedge Q))^+$ via P^-, but without the identity of P^- and $(\neg P)^+$, there is no way we can then obtain the conclusion that s verifies $(\neg(P \wedge Q))^+$ via $(\neg P)^+$. Moreover, as this case does not involve the falsification of any negation, the situation is exactly the same for (BV + BF).

To validate all of the desired introduction principles for ground, we therefore have to accept (IV + BF) as our account of negation. This account is striking in embodying an *asymmetric* view of negation. It distinguishes between the falsification of a negation and the verification of the negatum, but *not* between the verification of a negation and the falsification of the negatum. In other words, the view is that to falsify a content *is* to verify its negation, whereas it is not the case that to verify a content is to falsify its negation. Rather, falsifying the negation is something achieved through, but distinct from, the verification of the negatum.[15]

5 A mode-ified truthmaker theory of content

On the view I have proposed, the features of a truth to which ground is sensitive are features concerning what facts verify the truth and *how*, i.e. in what *modes* they do so. We therefore have to encode in (a mathematical representation of) a proposition both

[15] That this kind of asymmetric account should be accepted is a surprising result; independently of the connection to the introduction principles, it might have seemed more natural to endorse one of the symmetric accounts. So the question arises whether there might be independent philosophical reasons for endorsing the asymmetric account. Although the matter calls for a much more extended discussion than I can offer here, it may be worth mentioning one possible source of independent motivation. I have in mind the kind of asymmetric account of truth and falsity that is endorsed, for example, by Williamson (1994, p. 188), which can be captured by the following principles:

(T) If a proposition says that P, then it is true iff P.
(F) If a proposition says that P, then it is false iff $\neg P$.

This account immediately ties both the truth of a proposition saying that $\neg P$ and the falsity of a proposition saying that P to the same thing: it being the case that $\neg P$. But it does not in the same way tie the falsity of a proposition saying that $\neg P$ and the truth of a proposition saying that P to the same thing. Rather, the first is tied, in the first instance, to it being the case that $\neg\neg P$, and the second to it being the case that P. This asymmetry is at least strongly reminiscent of the identification of the falsification of P with the verification of $\neg P$, in the absence of the identification of the verification of P with the falsification of $\neg P$. As such, it may perhaps provide an independent basis for the latter.

what facts verify or falsify it as well as in what modes they do so. The obvious idea is to replace the sets of facts in Fine's account by sets of *pairs* $\langle s, m\rangle$ of a fact s and a mode m. The presence of $\langle s, m\rangle$ in P^+ (P^-) would then be taken to represent that s verifies (falsifies) P in mode m. The limiting case of a state verifying a proposition *directly*, as it were, i.e. not by verifying any propositions, may be represented by means of a special mode m_0 of directness. Any indirect mode m would be identified by the propositions $P_1, P_2, \ldots$ in the verification of which it consists. The truths $P_1, P_2, \ldots$ would be taken to ground a truth Q just in case for every fact s verifying $P_1, P_2, \ldots$, when m is the mode corresponding to $P_1, P_2, \ldots$, $\langle s, m\rangle$ is a member of Q.

I shall however deviate from this proposal in two ways. Firstly, it turns out that given our account of ground, we can work with a significantly simpler construction than the one just described: we may replace Fine's sets of facts simply with sets of modes, counting facts as special, *direct* modes of verification.[16] The presence of a fact s in P^+ represents that s verifies P directly, and the presence of an (indirect) mode m in P^+ represents that P is verified in mode m, by every fact verifying the propositions $Q_1, Q_2, \ldots$ corresponding to m. We may then take the propositions $P_1, P_2, \ldots$ to ground a proposition Q just in case there is a mode of verification corresponding to $P_1, P_2, \ldots$, and it is a member of Q. It is straightforward to show that the relation of ground holds according to the latter, simpler picture just in case it holds between the corresponding contents on the former, more complicated picture.

The second deviation is motivated by a desire to accommodate a *non-factive* understanding of ground. Roughly speaking, some propositions $P_1, P_2, \ldots$ non-factively ground a proposition Q iff they satisfy the conditions to ground Q, bar perhaps the condition of being true. This means in particular that non-factive ground satisfies unconditional versions of the introduction principles stated above, so that for example $P < P \vee Q$ holds irrespective of whether P or Q are true.[17] To characterize non-factive ground, we move to a similarly non-factive conception of indirect modes of verification.[18] Mainly, this means that we do not demand that the propositions $P_1, P_2, \ldots$ be true if they are to correspond to a mode of verification. We may then include in $P \vee Q$ the mode of verifying via P, irrespective of whether P is true. We can then say that P non-factively grounds $P \vee Q$ on the strength of the fact that $P \vee Q$ contains the mode

[16] The possibility of this simplification was suggested to me by Kit Fine.

[17] On the idea of a non-factive notion of ground, cf. e.g. Fine (2012b, p. 48ff) and Correia(2014b, p. 36).

[18] A similar move may be considered for the direct modes. Specifically, we may follow (cf. Fine 2014, p. 557f, 2016, p. 8), and replace the appeal to the notion of a fact by an appeal to a broader notion of a state, which is like that of a fact except in that a state need not be actual, and indeed need not even be possible. There is then no obstacle to assuming every proposition to have at least one verifier and at least one falsifier. However, since ground, on my account, is determined purely by the presence or absence of indirect modes in a given proposition, this extension of the conception of verifiers is not required for our construction to work as intended. To the extent that non-actual and impossible modes are less problematic than non-actual and impossible states, it is an advantage of my framework that it can accommodate non-factive ground using less worrisome resources than are required on Fine's approach. (It should be noted, though, that if we do not allow non-actual and impossible states, then the presence of a mode m in a proposition P does not in general represent exactly that every verifier of the propositions corresponding to m thereby verifies P. For the latter condition will be vacuously satisfied for any non-actual mode. The presence of m in P should then simply be understood to represent that verifying the propositions corresponding to m is a way to verify P—though it may be logically impossible to verify P in this way.)

of verifying via P. Given a suitable selection of modes as *actual*, we may then count a proposition true just in case one of its modes of verification is actual, and factive ground may be defined in the natural way in terms of truth and non-factive ground. For simplicity, I shall henceforth focus exclusively on non-factive ground.

Our first task will be to describe the basic behaviour of modes; this is done in Sect. 5.1. Section 5.2 defines notions of ground on our contents, and in Sect. 5.3 I then define suitable notions of conjunction, disjunction, and negation, which are shown to relate to ground in the desired way. Section 5.4 identifies the constraints on our ground-theoretic contents corresponding to the structural features ground is sometimes taken to possess. In Sect. 5.5 I turn to the matter of the identity conditions on modes, and explain how they bear on the question of equivalence in ground-theoretic content.

5.1 Modes

We shall begin by describing the basic behaviour of modes, leaving open their exact nature and identity.

Firstly, any mode is either *direct*, which is to say that it is a fact, or it is *indirect*, which is to say that there is some list of propositions $P_1, P_2, \ldots$ such that m is the mode of verifying *via* $P_1, P_2, \ldots$ We will also call direct modes *fundamental*, and indirect ones *derivative*. Given our informal account of the previous section, it is clear that *some* lists of propositions determine a mode, i.e. that there are some derivative modes. For instance, the verifiers of the conjunction $P \wedge Q$ of truths P, Q verify the conjunction by verifying P, Q. So there is a mode of verification corresponding to the list of propositions P, Q.

There appears to be no motivation, intuitive or theoretical, for allowing distinct modes to correspond to the same list of propositions. The mathematical structure corresponding to a list is a sequence, so we may take there to be a function mapping certain sequences of propositions to modes. We shall write this function V, alluding to the fact that the mode in question will be the mode of verifing *via* the relevant list of propositions. So when $\langle P_1, P_2, \ldots\rangle$ is a suitable sequence of propositions, $V\langle P_1, P_2, \ldots\rangle$ is the mode of verifying via $P_1, P_2, \ldots$

The basic structure that we will work with is a *mode-space*, which is a pair $\langle M, V\rangle$ of a non-empty set of modes M, and a via-function V. A maximally liberal conception of propositions is obtained as follows. We call any subset P of M a *unilateral proposition*, which is verified in exactly the modes which are its members. And we call any pair $\mathbf{P} = \langle \mathbf{P}^+, \mathbf{P}^-\rangle$ of unilateral propositions a *bilateral proposition*, which is verified in exactly the modes that are members of $\mathbf{P}^+$, and falsified in exactly the modes that are members of $\mathbf{P}^-$. (Here and in what follows, I mark the unilateral/bilateral distinction typographically by using boldface variables for bilateral propositions.) For now, we may focus on unilateral propositions.

The function V is assumed to be a mapping from some set of sequences of propositions into M. It is not required that V be defined for every sequence of propositions, or even that it be defined for every singleton sequence $\langle P\rangle$ with P a proposition. Informally, that V is defined for some sequence of propositions $\langle P_1, P_2, \ldots\rangle$ means that there is such a thing as doing something by verifying $P_1, P_2, \ldots$ Correspondingly, if V is undefined for $\langle P_1, P_2, \ldots\rangle$, this means that there is no such thing as doing something by verifying $P_1, P_2, \ldots$ One reason for not requiring that V be defined for all sequences of propositions is that such a requirement would be inconsistent with a view of ground as irreflexive. To see this, note that M is itself a proposition, so $V\langle P\rangle$ should have to be defined for $P = M$. But then $V\langle P\rangle$ would have to a member of P, which is to say that verifying P is a way to verify P. Given the proposed account of ground in terms of modes of verification, it would follow that P grounds itself.

We shall however take for granted that two natural closure principles hold for the range of sequences on which V is defined. Firstly, if V is defined for a given sequence $\langle P_1, P_2, \ldots\rangle$ then it is also defined for any non-empty subsequence. Secondly, if V is defined for each of the sequences $\gamma_1, \gamma_2, \ldots$, then V is also defined for their *concatenation* $\gamma_1{}^\frown\gamma_2{}^\frown\ldots$[19] For simplicity, we require that V be undefined for the empty sequence $\langle\rangle$.[20]

Note that V is not required to be one-to-one, but is allowed to map different sequences to the same mode. We shall assume that the property of determining the same mode is preserved under concatenation of sequences. That is, if $V(\gamma_1) = V(\delta_1),\ V(\gamma_2) = V(\delta_2),\ \ldots$, then $V(\gamma_1{}^\frown\gamma_2{}^\frown\ldots) = V(\delta_1{}^\frown\delta_2{}^\frown\ldots)$. If $m_1 = V(\gamma_1),\ m_2 = V(\gamma_2), \ldots$, we call $V(\gamma_1{}^\frown\gamma_2{}^\frown\ldots)$ *a fusion of* $\langle m_1, m_2, \ldots\rangle$. Given the previous constraints, any sequence of indirect modes $\langle m_1, m_2, \ldots\rangle$ has a unique fusion which we denote by $\bigsqcup\langle m_1, m_2, \ldots\rangle$ or $m_1 \sqcup m_2 \sqcup \ldots$

Whenever $m = V\langle P_1, P_2, \ldots\rangle$, we call the *set* $\{P_1, P_2, \ldots\}$ a *ground-set* of m, since $\{P_1, P_2, \ldots\}$ will ground a proposition Q if $m \in Q$. A mode-space will be called *constrained* iff: if V maps two sequences $\langle P_1, P_2, \ldots\rangle$ and $\langle Q_1, Q_2, \ldots\rangle$ of propositions to the same mode, then the corresponding sets $\{P_1, P_2, \ldots\}$ and $\{Q_1, Q_2, \ldots\}$ are identical. In that case, any indirect mode m has a unique ground-set, which we denote by $|m|$. We shall usually assume that we are working in a constrained mode-space.[21]

Having described what modes are like, it is natural to ask what modes *are*. In view of their intimate relationship with sequences of propositions, an obvious suggestion is that perhaps modes may simply be *identified* with these sequences. With respect to a large variety of mode-spaces, such an identification may indeed be carried through. There are, however, interesting sorts of mode-spaces for which this is not possible. The basic point is that sequences of propositions are set-theoretic constructions from propositions, and propositions themselves are set-theoretic constructions from modes. As a result, if modes are identified with sequences of propositions, the relation of ground

[19] For present purposes, we may restrict attention to concatenations of an at most countable sequence of sequences.

[20] There may be purposes for which a putative 'nullmode' corresponding to $\langle\rangle$ may be useful. If a mode is counted actual iff all propositions in the corresponding sequence are true, then the nullmode would automatically be actual and hence every proposition containing it trivially true. It would thereby have a similar profile to the *nullfact* (or nullstate) in Fine's framework, which is the fusion of the empty set of facts (or states), and part of every fact (state). The most obvious application of the nullmode would be to capture Fine's idea that some truths may be *zero-grounded*, where this is supposed to be distinct from being ungrounded; cf. Fine (2012b, p. 47f).

[21] It may be worth pointing out that for cardinality reasons, in a constrained mode-space, V will be undefined for most sequences. Thanks here to an anonymous referee.

will automatically inherit certain features of the membership-relation, in particular its irreflexivity, its asymmetry, and its well-foundedness. So for any mode-spaces that yield relationships of ground that violate these principles, the reduction of modes to sequences, or other set-theoretic constructions, of propositions, cannot be carried out. Whether a different useful kind of reduction is then possible, perhaps using non-well-founded sets, is a question I shall leave for another occasion. At any rate, in devising a framework for theorizing about ground, I believe we are well-advised first to try and work out what sort of constraints different plausible theories of ground impose on the modes and their structure. Once we understand this, we may return to the metaphysical question of what sorts of things modes may be taken to be on the various views. For the time being, we shall therefore adopt towards modes the stance that Fine adopts towards facts or states, and simply take them as given.

5.2 Ground

We first define notions of (non-factive) ground for unilateral contents. Given our informal account of strict full ground above, we may define $<$ and its partial cousin $\prec$ as follows. Let Γ be a non-empty set of (unilateral) contents and P and Q (unilateral) contents. Then:

($<$) $\Gamma < P$ iff for some mode $m \in P$, Γ is a ground-set of m
($\prec$) $P \prec Q$ iff $\Delta,\ P < Q$ for some set of propositions Δ

Note that in contrast to Fine's account above, strict ground is given a direct definition in terms of modes of verification, and not defined via a reflexive notion of weak ground. Nevertheless, we shall later have use for such a notion, and for a partial version of it, which we define as follows:

($\leq$) $\Gamma \leq P$ iff $\Gamma < P$ or $\Gamma = \{P\}$ or ($\{P\} \subset \Gamma$ and $\Gamma \backslash \{P\} < P$)
($\preceq$) $P \preceq Q$ iff $P = Q$ or $P \prec Q$

The following principles are straightforward consequences of these definitions:[22]

Identity($\leq$) $P \leq P$
Subsumption($<$/$\prec$) If $\Gamma,\ P < Q$ then $P \prec Q$
Subsumption($<$/$\leq$) If $\Gamma < Q$ then $\Gamma \leq Q$
Subsumption($\leq$/$\preceq$) If $\Gamma,\ P \leq Q$ then $P \preceq Q$
Subsumption($\prec$/$\preceq$) If $P \prec Q$ then $P \preceq Q$

The only non-trivial case is Sub($\leq$/$\preceq$). But suppose $\Gamma,\ P \leq Q$. There are three cases. (i) $\Gamma,\ P < Q$. Then $P \prec Q$ and hence $P \preceq Q$. (ii) $\Gamma \cup \{P\} = \{Q\}$. Then $P = Q$ and hence $P \preceq Q$. (iii) $\{Q\} \subset \Gamma \cup \{P\}$ and $(\Gamma \cup \{P\}) \backslash \{Q\} < Q$. Then either $P = Q$ and hence again $P \preceq Q$, or $P \in (\Gamma \cup \{P\}) \backslash \{Q\}$, and hence $P \prec Q$, and so $P \preceq Q$.

We might also have defined weak partial ground as the partial version of weak full ground rather than, as we have done above, as a weak version of strict partial ground.

[22] I borrow the labels from the corresponding inference rules in Fine (2012c)'s pure logic of ground. – Here and in what follows, I adopt the familiar convention of writing $\Gamma \cup \{Q\}$ as $\Gamma,\ Q$ as well as $\Gamma \cup \Delta$ as $\Gamma,\ \Delta$, and similarly in other cases.

For given the other definitions, the condition that $P \prec Q$ or $P = Q$ is equivalent to the condition that for some Δ, we have $\Delta,\ P \leq Q$. The right-to-left direction was just established. For the left-to-right direction, suppose first that $P \prec Q$. Then for some Δ, we have $\Delta,\ P < Q$ and hence $\Delta,\ P < Q$. Suppose then that $P = Q$. Then $P \leq Q$, and hence again, for some Δ, we have $\Delta,\ P \leq Q$.

The extension of the four notions of ground to bilateral contents is done in the simplest possible way. For a non-empty set $\boldsymbol{\Gamma}$ of bilateral contents, let $\boldsymbol{\Gamma}^+ = \{\mathbf{P}^+ : \mathbf{P} \in \Gamma\}$. Then we define $\boldsymbol{\Gamma} < \mathbf{P}$ by $\boldsymbol{\Gamma}^+ < \mathbf{P}^+$, and likewise in the other cases.

5.3 The truth-functional operations

In this section I define truth-functional operations on our ground-theoretic contents. They are shown to combine with the above account of ground to yield an attractive account of the grounds of truth-functionally complex contents, which includes the unrestricted introduction principles for ground. We first define conjunction and disjunction for unilateral contents. We then extend these operations to the case of bilateral contents and define an operation of negation. We then establish general necessary and sufficient conditions for some propositions to ground the various kinds of truth-functionally complex propositions.

Recall the informal account given in the previous section of how disjunctions, conjunctions, and negations are verified. With respect to disjunction, we said that any verifier of P verifies $P \vee Q$ by verifying P, and any verifier of Q verifies $P \vee Q$ by verifying Q. And with respect to conjunction, we said that any fusion $s \sqcup t$ of verifiers s of P and t of Q verifies $P \wedge Q$ by verifying $P,\ Q$. It follows that given any disjunction $P \vee Q$, V must be defined for $\langle P \rangle$, and $\langle Q \rangle$, and $P \vee Q$ must include both $V\langle P \rangle$ and $V\langle Q \rangle$. Likewise for any conjunction $P \wedge Q$, V must be defined for $\langle P,\ Q \rangle$, and $P \wedge Q$ must include $V\langle P,\ Q \rangle$.

Plausibly, both conjunctions and disjunctions may also be verified in other ways. Thus, suppose that P may be verified by verifying some propositions $R_1,\ R_2,\ \ldots$ Then it is natural to hold that $P \wedge Q$ may also be verified by verifying $R_1,\ R_2, \ldots,\ Q$. Similarly, if Q may be verified by verifying $R_1,\ R_2, \ldots$, it is natural to say that $P \wedge Q$ may also be verified by verifying $P,\ R_1,\ R_2,\ \ldots$ And finally, if P and Q may be verified via $R_1,\ R_2, \ldots$ and $S_1,\ S_2, \ldots$, respectively, $P \wedge Q$ may be verified via $R_1,\ R_2, \ldots,\ S_1,\ S_2, \ldots$

In the case of the disjunction $P \vee Q$, it seems similarly plausible that it may be verified not only via P and via Q, but also via $R_1,\ R_2, \ldots$ whenever either of P and Q may be so verified. In addition, we shall suppose that every mode of verifying the conjunction $P \wedge Q$ is also automatically a mode of verifying $P \vee Q$.[23]

We may usefully state this account of conjunction and disjunction more formally in terms of three auxiliary operations on unilateral propositions. At certain key places, we need to ensure that V is defined for $\langle P \rangle$. We shall say that a proposition P is

[23] This parallels Fine's account of the truthmakers of disjunctions which include not only the verifiers of the disjuncts, but also any fusions of these.

raisable exactly when this is the case. Then when $P,\ Q$ are propositions including only derivative modes, and when R is a raisable proposition:

(Df. $\sqcup$) $P \sqcup Q :=\quad \{m \sqcup n : m \in P \text{ and } n \in Q\}$
(Df. $+$) $P + Q :=\quad (P \cup Q) \cup (P \sqcup Q)$
(Df. $\Uparrow$) $\uparrow R :=\quad \{V\langle R\rangle\} + \{m \in R : m \text{ is derivative}\}$

Note that there is a close correspondence between the operations $\sqcup$ and $+$ on unilateral contents and Fine's operations of conjunction and disjunction, respectively. For the way $P \sqcup Q$ is obtained from P and Q is exactly analogous to how, one Fine's view, the positive content of a conjunction is obtained from that of its conjuncts. Likewise the way $P+Q$ is obtained from P and Q is analogous to how, one Fine's view, the positive content of a disjunction is obtained from that of its disjuncts. On the present account, however, conjunction and disjunction are defined in terms of these operations and the third one, which I call *raising*. Its effect is to produce a proposition $\uparrow P$ just like its argument P, except that it may be verified via P, and via $P,\ R_1,\ R_2,\ \ldots$ whenever P may be verified via $R_1,\ R_2, \ldots$[24] The following definitions of unilateral conjunction and disjunction accord with our above account of what modes should be included in $P \vee Q$ and $P \wedge Q$. For raisable propositions $P,\ Q$:

(Df. $\wedge$U) $P \wedge Q :=\quad \uparrow P \sqcup \uparrow Q$
(Df. $\vee$U) $P \vee Q :=\quad \uparrow P + \uparrow Q$

Given our informal discussion of negation in the previous section, it is immediate how conjunction, disjunction, and negation are now to be defined on bilateral contents. For pairs of raisable unilateral propositions $\mathbf{P},\ \mathbf{Q}$:

(Df. $\neg$B) $\neg\mathbf{P} :=\ \langle \mathbf{P}^-,\ \uparrow\mathbf{P}^+\rangle$
(Df. $\wedge$B) $\mathbf{P} \wedge \mathbf{Q} :=\ \langle \mathbf{P}^+ \wedge \mathbf{Q}^+,\ \mathbf{P}^- \vee \mathbf{Q}^-\rangle$
(Df. $\vee$B) $\mathbf{P} \vee \mathbf{Q} :=\ \langle \mathbf{P}^+ \vee \mathbf{Q}^+,\ \mathbf{P}^- \wedge \mathbf{Q}^-\rangle$

Note how the definition of $\neg$ reflects the *asymmetric* view of negation.

We may then establish substantive necessary and sufficient conditions for a set of propositions $\boldsymbol{\Gamma}$ to ground any kind of truth-functionally complex proposition. To state them concisely, we introduce some abbreviations. Let $\boldsymbol{\Gamma} \leq \{\mathbf{P}_1,\ \mathbf{P}_2, \ldots\}$ abbreviate:

24 Note that absent any assumptions to the effect that P cannot be verified in part by verifying P, there is no guarantee that $\uparrow P \neq P$. – One might wonder whether an argument is not needed for the claim that there always exists such a proposition as $\uparrow P$. Formally, the assumption that P is raisable, in conjunction with the closure of the set of modes under fusion, makes sure that a suitable proposition always exists. But we may then ask for a defence of this assumption; what justifies disregarding propositions that are not raisable? The simplest answer is perhaps this. For any legitimate proposition $\mathbf{P}$, it should be possible to form its double negation $\neg\neg\mathbf{P}$. Given the proposed account of ground and negation, $(\neg\neg\mathbf{P})^+$ relates to $\mathbf{P}^+$ exactly so that $\uparrow(\mathbf{P}^+) = (\neg\neg\mathbf{P})^+$ (similarly for $(\neg\neg\mathbf{P})^-$). So whenever $\uparrow P$ does not exist, P cannot occur within a legitimate bilateral content and may for that reason be discarded. We can perhaps also argue for the raisability of legitimate propositions on independent grounds. For it seems that if P is a legitimate (unilateral) proposition, then there is such a thing as (non-factively) verifying P – if no sense can be made of the idea of P being verified, there is something incoherent about P. But if there is such a thing as verifying P, then it seems we can ask what can be done *by* verifying P. Now this question is about a way, or mode of doing something, namely the mode of doing it by verifying P. So this mode should be taken to exist. But since this mode is just the mode $V\langle P\rangle$, it follows that P is raisable.

for some sets $\boldsymbol{\Gamma}_1, \boldsymbol{\Gamma}_2, \ldots$ with $\boldsymbol{\Gamma} = \boldsymbol{\Gamma}_1 \cup \boldsymbol{\Gamma}_2 \cup \ldots$, $\boldsymbol{\Gamma}_1 \leq \mathbf{P}_1$ and $\boldsymbol{\Gamma}_2 \leq \mathbf{P}_2$ and ... Then for raisable propositions $\mathbf{P}$ and $\mathbf{Q}$ and any set of raisable propositions $\boldsymbol{\Gamma}$:[25]

$(< \wedge)\ \boldsymbol{\Gamma} < \mathbf{P} \wedge \mathbf{Q}$ iff $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
$(< \vee)\ \boldsymbol{\Gamma} < \mathbf{P} \vee \mathbf{Q}$ iff $\boldsymbol{\Gamma} \leq \mathbf{P}$ or $\boldsymbol{\Gamma} \leq \mathbf{Q}$ or $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
$(< \neg\neg)\ \boldsymbol{\Gamma} < \neg\neg\mathbf{P}$ iff $\boldsymbol{\Gamma} \leq \mathbf{P}$
$(< \neg\wedge)\ \boldsymbol{\Gamma} < \neg(\mathbf{P} \wedge \mathbf{Q})$ iff $\boldsymbol{\Gamma} \leq \neg\mathbf{P}$ or $\boldsymbol{\Gamma} \leq \neg\mathbf{Q}$ or $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$
$(< \neg\vee)\ \boldsymbol{\Gamma} < \neg(\mathbf{P} \vee \mathbf{Q})$ iff $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$

By the reflexivity of $\leq$, the introduction principles for ground can be obtained from the right-to-left directions of these biconditionals. For example, since $\mathbf{P} \leq \mathbf{P}$ and $\mathbf{Q} \leq \mathbf{Q}$, it follows that $\mathbf{P}, \mathbf{Q} \leq \{\mathbf{P}, \mathbf{Q}\}$, and hence by the right-to-left direction of $(< \wedge)$, $\mathbf{P}, \mathbf{Q} < \mathbf{P} \wedge \mathbf{Q}$. The left-to-right directions of the biconditionals in turn correspond exactly to the elimination rules for ground proposed by Fine (2012b, pp. 63ff).[26]

A number of variations on the above definitions may be considered. For example, we might not allow (unilateral) disjunctions $P \vee Q$ to be verified via P, Q, but only via P and via Q, as well as perhaps in modes in which the latter are verified. The above results would then still hold once we drop the third disjunct from the right-hand sides of $(< \vee)$ and $(< \neg\wedge)$. It would also be very interesting to see exactly what logical principles would hold on a symmetric account of negation.

What about views, like that proposed in Correia (2010), which impose substantive restrictions even on, say, the principle that $P < P \vee Q$? Can they also be accommodated within the present framework? There is no general reason why this should not be possible. As long as the restrictions imposed can somehow be captured in terms of the modes in which the component propositions are verified, we may simply exclude the offending modes from the complex propositions as defined above. We might, for example, stipulate that $P \vee Q$ is verified via $V\langle P \rangle$ and $V\langle Q \rangle$ only if $P \neq Q$. This would ensure that $P \not< P \vee P$.[27]

Note, however, that these sorts of moves cannot be motivated in the way in which Correia argues for his favoured restrictions on the introduction principles. For as we saw earlier, Correia appeals to a *prior* standard for the individuation of ground-theoretic content on which P and $P \vee P$ are identified. But when ground-theoretic

[25] A proof of this result is given in the Appendix.

[26] The elimination rules are perhaps more controversial than the introduction rules; for instance, it might be suggested that $P \vee \neg P$ is not only grounded by the weak grounds of its true disjunct, but also by the laws of logic, on some suitable construal of that phrase (this idea is also mentioned, but not endorsed in Schnieder 2011, p. 457f). So it may be worth noting that the elimination rules may be invalidated in a natural way without the introduction rules thereby also becoming invalid. All we need to do is to drop the requirement that the mode-space be constrained. (Whether and how the suggestion that the laws of logic ground $P \vee \neg P$ could be implemented within the overall framework developed here is harder to answer, and depends strongly on how exactly that view is spelt out.) Thanks here to an anonymous referee.

[27] It might also be possible to achieve the same result not by changing the definition of disjunction, but by revising the definition of grounding, giving up on the tight connection that every mode of verification corresponds to an instance of grounding. – Note that the wish to reject $P < P \vee P$ is not the only possible motivation for wanting to restrict the disjunction principle that $P < P \vee Q$; some of the possible responses to the puzzles of ground presented in Fine (2010) also involve such a restriction. With respect to these views, the same comments apply: as long as the relevant restrictions can be captured within our framework, the views can be accommodated. Thanks to an anonymous referee for bringing up the matter of the puzzles of ground.

content is conceived as on the present view as capturing in what modes a content is verified, then the latter question needs to be answered independently of assumptions concerning the identity and distinctness of the contents involved. I take this to provide some additional support for Correia's contention that the view of ground as satisfying the unrestricted introduction principles targets a *different notion* of ground, sensitive to different features of a sentence's overall content, than the view on which only a restricted version of the principles hold (cf. Correia 2010, pp. 256ff, 2014a, Sects. 3, 5, 2014b, p. 36).

5.4 Structural properties of ground

In this section we discuss some structural properties that ground may be taken to possess, and how they may be captured within my framework. There is in principle an unsurveyable number of such properties; my choice in which of them to discuss has been guided in part by which look particularly natural from the perspective of my framework, and by which of them are endorsed in the currently most well-developed view on the matter, which is that of Fine (2012c).

It is often held that ground is irreflexive in the sense that nothing helps ground itself.[28] Given our definition of ground, this amounts to the claim that there is no proposition P such that P is a member of a ground-set of some mode $m \in P$.[29] Nothing we have said so far guarantees this. Specifically, given a proposition P, nothing we have said rules out that V maps $\langle P \rangle$ to a member of P, in which case $P < P$.

Call a proposition P *irreflexive* iff P is not an element of any sequence γ with $V(\gamma) \in P$. If $V\langle P_1, P_2, \ldots \rangle \notin \bigcup\{P_1, P_2, \ldots\}$ for any propositions $P_1, P_2, \ldots$, then all propositions are irreflexive. For irreflexive propositions, and *only* for them, the following principles hold:[30]

Irreflexivity($<$) $\Gamma, P \nless P$
Irreflexivity($\prec$) $P \nprec P$

Irreflexive propositions also satisfy the following principle.[31]

Reverse Subsumption($\leq$/$<$) If $\Gamma \leq P$ and $Q \prec P$ for all $Q \in \Gamma$, then $\Gamma < P$

This principle corresponds to a basic rule of inference in Fine's logic; it is closely related to the characterization of strict full ground as 'irreversible' weak full ground,

[28] Since this is the standard view, more interesting than listing sources for it is to give some sources where the principle has been called into question. One context in which self-grounding has been considered a possibility is that of the paradoxes of ground described in Fine (2010) and Krämer (2013); see in particular (Correia 2014b, Sect. 7). Different kinds of doubts about irreflexivity are raised in Jenkins (2011).

[29] Throughout this section, by 'content' and 'proposition' I shall mean unilateral content.

[30] For the first principle, assume $\Gamma, P < P$. Then Γ, P is a ground-set of some derivative mode $m \in P$, so for some sequence of propositions γ with $V(\gamma) = m$, Γ, P is the set underlying γ. But then P is an element of γ, contrary to our assumption. The second principle follows by the definition of $\prec$. The other directions are equally straightforward.

[31] Assume $\Gamma \leq P$ and $Q \prec P$ for all $Q \in \Gamma$. Note that $P \notin \Gamma$, for otherwise $P \prec P$, contradicting irreflexivity. But then neither $\Gamma = \{P\}$ nor $\{P\} \subset \Gamma$, and therefore $\Gamma < P$.

i.e. the equivalence of $\Gamma < P$ to $\Gamma \leq P$ and $P \npreceq Q$ for any $Q \in \Gamma$. The right-to-left direction of this equivalence is ensured already by our definition of $\leq$. For suppose $\Gamma \leq P$ and $P \npreceq Q$ whenever $Q \in \Gamma$. Unless $\Gamma < P$, then from the definition of $\leq$ it follows immediately that $P \in \Gamma$. But since $P \preceq P$, this contradicts the assumption that $P \npreceq Q$ whenever $Q \in \Gamma$.

The left-to-right direction of the equivalence holds whenever another commonly accepted principle for ground holds, namely that ground is asymmetric in the sense that if P helps ground Q, then Q does not help ground P. We may call a proposition P *asymmetric* iff whenever $V\langle Q_1, Q_2, \ldots\rangle \in P$ and $V\langle R_1, R_2, \ldots\rangle \in \bigcup\{Q_1, Q_2, \ldots\}$, then P is not among $R_1, R_2, \ldots$ It is straightforward to show that for exactly the asymmetric propositions P:

Asymmetry($<$) If $\Gamma, Q < P$, then $\Delta, P \nless Q$
Asymmetry($\prec$) If $Q \prec P$, then $P \nprec Q$

But now suppose $\Gamma < P$. Then, firstly, $\Gamma \leq P$. Secondly, the relationship of weak ground is irreversible. For suppose $P \preceq Q$ for some $Q \in \Gamma$. Then either $P \prec Q$ or $P = Q$. But since $Q \prec P$, Asymmetry($\prec$) rules out both possibilities. So strict full ground implies irreversible weak full ground for irreflexive and asymmetric propositions.

I wish to briefly mention two further principles of interest related to those just considered.

Redundancy($\leq$) If $\Gamma, P \leq P$ then $\Gamma \leq P$
Reverse Subsumption($\prec/=$) If $P \preceq Q$ and $Q \preceq P$, then $P = Q$

The first one follows from irreflexivity, the second from asymmetry.[32]

The principles are interesting in part because they do *not* hold in the logic of Fine (2012c). To see this, let fact s be a proper part of fact t, i.e. $t = t \sqcup s$ and $t \neq s$. Then consider $P = \{s, t\}$ and $Q = \{t\}$. Evidently, $P \neq Q$. But $Q \subseteq P$, so on Fine's account, $Q \leq P$, and hence $Q \preceq P$. Moreover, since $s \sqcup t = t \sqcup t = t$, every fusion of verifiers of P and Q is a verifier of Q, so we have $P, Q \leq Q$, and hence $P \preceq Q$. So P and Q are a counter-example to Reverse Subsumption($\prec/=$), and indeed to the weaker claim that mutual weak *full* ground implies identity. Moreover, they also give rise to a counter-example to Redundancy($\leq$), for as we saw, $P, Q \leq Q$, whereas $P \nleq Q$. The differences between Fine's approach to weak ground and ours thus manifest themselves in significant differences concerning the structural properties of weak ground.

It is also often maintained that ground is transitive.[33] A number of different principles may be subsumed under this heading. Relatively weak examples are that if $P \prec Q$

[32] For the first principle, assume $\Gamma, P \leq P$ for non-empty Γ. By irreflexivity, $\Gamma, P \nless P$. There remain two cases. (i) $\Gamma \cup \{P\} = \{P\}$. Then $\Gamma = \{P\}$, and hence $\Gamma \leq P$. (ii) $(\Gamma \cup \{P\}) \backslash \{P\} \leq P$. If $P \in \Gamma$, then $\Gamma \cup \{P\} = \Gamma$ and hence by assumption $\Gamma \leq P$. If $P \notin \Gamma$, then $(\Gamma \cup \{P\}) \backslash \{P\} = \Gamma$ and hence again, $\Gamma \leq P$. The second principle is immediate given asymmetry and the definition of $\preceq$.

[33] Again, it may be more interesting to mention dissenters. The most influential arguments against transitivity are perhaps those given in Schaffer (2012, p. 126ff). For replies, see Litland (2013), Raven (2013) and Krämer and Roski (2016).

and $Q \prec R$, then $P \prec R$, or that if $\Gamma < P$, and $P < Q$, then $\Gamma < Q$. A stronger form of transitivity may be stated as follows:

Transitivity($<$) If $\Gamma < P$ and $\Delta, P < Q$, then $\Gamma, \Delta < Q$

This principle holds, for example, on the views of Fine (2012c, p. 5, 22) and Correia (2010, p. 262).

Let us then consider what might be natural constraints of transitivity to impose on our contents. Here is one obvious idea, corresponding to the second weak interpretation of transitivity above. Suppose Q contains $V\langle P\rangle$, and P contains $V(\gamma)$. Then it might naturally be required that $V(\gamma)$ should also be contained in Q. This would ensure that if $\Gamma < P$, and $P < Q$, then $\Gamma < Q$.

But consider now the case in which Q does not contain $V\langle P\rangle$, but contains some mode $V(\delta)$ with P an element of δ. A natural suggestion is that transitivity requires that Q also contain $V(\delta^{\gamma/\langle P\rangle})$ where $\delta^{\gamma/\langle P\rangle}$ results from systematically replacing $\langle P\rangle$ in δ by γ.[34] This, however, is not sufficient for Transitivity($<$). For let Γ be the set corresponding to γ, so $\Gamma < P$. Now suppose Q contains $V\langle P, R\rangle$ with $R \neq P$, so $P, R < Q$. Now let $\Delta = \{P, R\}$, and note that since $\Delta = \Delta, P$, it follows that $\Delta, P < Q$. So by Transitivity($<$), $\Gamma, \Delta = \Gamma, P, R < Q$. However, the envisaged transitivity constraint on contents yields only that Q contains $V(\gamma {}^\frown \langle R\rangle)$. The ground-set of this mode is Γ, R which is distinct from Γ, P, R unless $P \in \Gamma$. So there is no guarantee that $\Gamma, P, R < Q$.

The transitivity principle for ground validated by this constraint may instead be stated thus:

Transitivity($<$)* If $\Gamma < P$ and $\Delta < Q$, then $\Delta^{\Gamma/P} < Q$

where $\Delta^{\Gamma/P}$ is $(\Delta \setminus \{P\}) \cup \Gamma$ if $P \in \Delta$ and Δ otherwise. If, on the other hand, we wish to ensure the stronger transitivity principle, we may impose the following constraint: If Q contains a mode with ground-set Δ, P, then Q contains a mode with ground-set Γ, Δ whenever P contains a mode with ground-set Γ.[35] Put in terms of corresponding via-sequences, we might say that if Q contains $V(\delta)$, P is an element of δ, and P contains $V(\gamma)$, then Q contains $V(\delta')$ whenever δ' may be obtained from δ by systematically replacing $\langle P\rangle$ in δ by either γ or $\langle P\rangle {}^\frown \gamma$. However, from an intuitive point of view, this seems to have a weaker claim than the previous constraint to amount simply to a principle of transitivity concerning modes of verification.

The final structural principle I wish to consider here is the following principle of amalgamation, which is derivable in Fine's pure logic of ground (cf. Fine 2012c, p. 7):

Amalgamation($<$) If $\Gamma_1 < P$ and $\Gamma_2 < P$ and ..., then $\Gamma_1 \cup \Gamma_2 \cup \ldots < P$

Like the previous principles, it does not hold for arbitrary contents. We may say that a content is *closed* iff, whenever it contains modes with ground-sets $\Gamma_1, \Gamma_2, \ldots$, it also contains some mode with ground-set $\Gamma_1 \cup \Gamma_2 \cup \ldots$ Then Amalgamation($<$) holds exactly for closed propositions P.

[34] We might perhaps also consider allowing the replacement of only some occurrences of $\langle P\rangle$, which would create some additional complications because what modes can be obtained by transitivity from a mode can then not be read off from the corresponding ground-set.

[35] It is immediate that Transitivity($<$) will hold exactly for the propositions satisfying this constraint.

Call a proposition *normal* iff it is closed, irreflexive, and satisfies the strong transitivity constraint. We may then report two welcome results about normal propositions.[36] Firstly, normal propositions satisfy all the principles corresponding to the rules of the pure logic of ground advocated in Fine (2012c, p. 5).[37] Most of them have already been stated, the remaining ones are principles of transitivity involving weak grounding relationships:

Transitivity($\leq$/$\leq$)	If $\Gamma \leq P$ and $\Delta, P \leq Q$ then $\Gamma, \Delta \leq P$
Transitivity($\preceq$/$\prec$)	If $P \preceq Q$ and $Q \prec R$ then $P \prec R$
Transitivity($\prec$/$\preceq$)	If $P \prec Q$ and $Q \preceq R$ then $P \prec R$
Transitivity($\prec$/$\prec$)	If $P \preceq Q$ and $Q \preceq R$ then $P \preceq R$

Secondly, the property of normality is preserved under our truth-functional operations. Our understanding of conjunction, disjunction, and negation therefore coheres with a view of ground as satisfying the normality properties in the sense that if we start with a range of normal propositions, application of our truth-functional operations to them will never take us to non-normal propositions.[38]

5.5 Ground-theoretic equivalence and the individuation of modes

We have so far said very little about the conditions under which contents obtained by application of truth-functional operations are identical. Closely related, but even more important for the purposes of the theory of ground is the question under what conditions they are *ground-theoretically equivalent*, where this term is understood in accordance with the following definition:[39]

($\approx$) **P** is ground-theoretically equivalent to **Q** ($\mathbf{P} \approx \mathbf{Q}$) iff
(i) for all $\boldsymbol{\Gamma}$: $\boldsymbol{\Gamma} < \mathbf{P}$ iff $\boldsymbol{\Gamma} < \mathbf{Q}$, and
(ii) for all $\boldsymbol{\Delta}$ and **R**: $\boldsymbol{\Delta}, \mathbf{P} < \mathbf{R}$ iff $\boldsymbol{\Delta}, \mathbf{Q} < \mathbf{R}$.

It is easily seen that $\approx$ is indeed an equivalence relation.

Ground-theoretic equivalence is closely related to identity on bilateral and unilateral contents. Evidently, the identity of bilateral contents implies their (ground-theoretic) equivalence. Indeed, because we have defined ground on bilateral contents without regard to negative content, two propositions will be equivalent already if they have the same *positive* content. This gives us a first substantive result about equivalence. For the DeMorgan laws are easily shown to hold for identity of positive content, and they therefore also hold for $\approx$:

---

36 The proofs are in the appendix.

37 As we have seen, however, they satisfy some additional principles as well.

38 Further structural principles about ground might of course be considered, and implemented in the form of suitable constraints on propositions. The most significant ones among them may be the various versions of the claim that ground is well-founded. For an illuminating discussion of the various possible interpretations of the claim, see Dixon (2016). See Litland (2016a) for an argument for an instance of non-well-founded grounding.

39 I borrow the term from Fine (2012b, pp. 63, 67) who does not explicitly define it but seems to use it in at least roughly the same sense.

(DeMorgan 1) $\neg(\mathbf{P} \wedge \mathbf{Q}) \approx \neg\mathbf{P} \vee \neg\mathbf{Q}$
(DeMorgan 2) $\neg(\mathbf{P} \vee \mathbf{Q}) \approx \neg\mathbf{P} \wedge \neg\mathbf{Q}$

Assuming irreflexivity, the relation between identity of positive content and $\approx$ turns out to be even tighter: $\approx$ is *equivalent* to identity of positive content.[40]

($\approx$-Equivalence) $\mathbf{P} \approx \mathbf{Q}$ iff $\mathbf{P}^+ = \mathbf{Q}^+$

Using this fact, it is easy to show that $\approx$ is preserved under $\wedge$, $\vee$, and $\neg\neg$:

($\approx \wedge$) If $\mathbf{P} \approx \mathbf{Q}$ then $\mathbf{P} \wedge \mathbf{R} \approx \mathbf{Q} \wedge \mathbf{R}$ and $\mathbf{R} \wedge \mathbf{P} \approx \mathbf{R} \wedge \mathbf{Q}$
($\approx \vee$) If $\mathbf{P} \approx \mathbf{Q}$ then $\mathbf{P} \vee \mathbf{R} \approx \mathbf{Q} \vee \mathbf{R}$ and $\mathbf{R} \vee \mathbf{P} \approx \mathbf{R} \vee \mathbf{Q}$
($\approx \neg\neg$) If $\mathbf{P} \approx \mathbf{Q}$ then $\neg\neg\mathbf{P} \approx \neg\neg\mathbf{Q}$

However, it can be shown that $\approx$ is *not* preserved under $\neg$. The DeMorgan equivalents constitute a relatively simple counter-example. To get an idea why, note that although $\neg(\mathbf{P} \wedge \mathbf{Q}) \approx \neg\mathbf{P} \vee \neg\mathbf{Q}$, negating the right-hand side yields $\neg(\neg\mathbf{P} \vee \neg\mathbf{Q})$, which by the other DeMorgan law is equivalent to $\neg\neg\mathbf{P} \wedge \neg\neg\mathbf{Q}$. This, however, is not in general equivalent to $\neg\neg(\mathbf{P} \wedge \mathbf{Q})$, that is, the negation of the left-hand side of the original equivalence. For while the former is always grounded by $\neg\neg\mathbf{P}$, $\neg\neg\mathbf{Q}$, this is not in general true of $\neg\neg(\mathbf{P} \wedge \mathbf{Q})$.

We should also like to know under what conditions the various kinds of truth-functionally complex propositions are equivalent to one another. The answer to this depends on the general conditions under which V is taken to map two sequences to the same mode. So far, we have committed to the thesis that given any sequence of propositions, there is at most one mode of verifying via that sequence. So any mode which is a mode of verifying via some sequence of propositions is uniquely identified by that sequence. This puts an upper bound on the fineness of grain with which we may individuate modes: they are at most as finely individuated as the corresponding sequences of propositions.

We have also assumed that two sequences correspond to the same mode *only if* the sequences correspond to the same *set*, i.e. if the same propositions belong to both. This puts a lower bound on the fineness of grain with which we may individuate modes: they are at least as finely individuated as the *sets* determined by their corresponding sequences. There is at least one natural intermediate option, which is to abstract from the order of the propositions in a sequence corresponding to a mode, but not from repetitions. On this view, modes are exactly as finely individuated as the *multi-sets* determined by their corresponding sequences.

We shall say that the via-function V is *sequential* iff: $V(\gamma) = V(\delta)$ just in case $\gamma = \delta$; *semi-extensional* iff: $V(\gamma) = V(\delta)$ just in case γ and δ determine the same multi-set; and *extensional* iff: $V(\gamma) = V(\delta)$ just in case γ and δ determine the same set. It is beyond the scope of this paper to determine, in terms of the relationships between the component propositions, the exact conditions under which truth-functionally complex propositions will be equivalent under each of these views. I shall here confine myself to some observations concerning the most distinctive features of the three approaches.

[40] For assume $\mathbf{P} \approx \mathbf{Q}$. Note that $\mathbf{P} < \neg\neg\mathbf{P}$, so $\mathbf{Q} < \neg\neg\mathbf{P}$, so $\mathbf{Q} \leq \mathbf{P}$, and hence either $\mathbf{Q} < \mathbf{P}$ or $\mathbf{Q}^+ = \mathbf{P}^+$. So if $\mathbf{Q}^+ \neq \mathbf{P}^+$, it follows that $\mathbf{Q} < \mathbf{P}$. By $\mathbf{P} \approx \mathbf{Q}$, we may infer $\mathbf{P} < \mathbf{P}$, contrary to the assumption of irreflexivity.

If V is sequential, then it will not in general be the case that $\mathbf{P} \wedge \mathbf{Q} \approx \mathbf{Q} \wedge \mathbf{P}$, or that $\mathbf{P} \vee \mathbf{Q} \approx \mathbf{Q} \vee \mathbf{P}$. For suppose $\mathbf{P}$ and $\mathbf{Q}$ are fundamental propositions, i.e. propositions whose positive components contain no derivative modes, and suppose $\mathbf{P}^+ \neq \mathbf{Q}^+$. For the case of conjunction, note that $\mathbf{P}^+ \wedge \mathbf{Q}^+$ and $\mathbf{Q}^+ \wedge \mathbf{P}^+$ will then each include exactly one derivative mode, namely $V\langle \mathbf{P}^+, \mathbf{Q}^+\rangle$ and $V\langle \mathbf{Q}^+, \mathbf{P}^+\rangle$, respectively. Since $\mathbf{P} \neq \mathbf{Q}$, $\langle \mathbf{P}, \mathbf{Q}\rangle \neq \langle \mathbf{Q}, \mathbf{P}\rangle$. By sequentiality of V, $V\langle \mathbf{P}^+, \mathbf{Q}^+\rangle \neq V\langle \mathbf{Q}^+, \mathbf{P}^+\rangle$, and hence $\mathbf{P}^+ \wedge \mathbf{Q}^+ \neq \mathbf{Q}^+ \wedge \mathbf{P}^+$, so by ($\approx$-Equivalence), $\mathbf{P} \wedge \mathbf{Q} \not\approx \mathbf{Q} \wedge \mathbf{P}$. Similar considerations establish the point for disjunction. However, if V is semi-extensional or extensional, these equivalences will hold. We consider only conjunction. Every mode in $\mathbf{P}^+ \wedge \mathbf{Q}^+$ may be written $V(\gamma \frown \delta)$ where $V(\gamma) \in \mathbf{P}^+$ and $V(\delta) \in \mathbf{Q}^+$. But then $\mathbf{Q}^+ \wedge \mathbf{P}^+$ includes $V(\delta \frown \gamma)$, and since $\gamma \frown \delta$ and $\delta \frown \gamma$ determine the same multi-set, and therefore the same set, the modes are identical.

A distinctive feature of the extensional approach is that it yields some general equivalences between the values of *different* truth-functional operations. In particular, for any closed $\mathbf{P}$, we have $\mathbf{P} \wedge \mathbf{P} \approx \mathbf{P} \vee \mathbf{P} \approx \neg\neg\mathbf{P}$. Indeed, the contents in question will be fully identical.[41] Now call a bilateral proposition conjunctive, disjunctive, or negative, according as it is the value of the operation of conjunction, disjunction, or negation on bilateral contents. Our observation then shows that on the extensional approach, there is no clear divide among the propositions that are conjunctive, disjunctive, or negative, but that these categories overlap. This is significant because it calls into question at least one interpretation of Correia's claim that the acceptance of the unrestricted introduction principles for ground commits us to a conceptual, or representational conception of the relata of ground. For one good sense we could give to the notion of a representational conception of content is that it applies just in case contents are individuated in terms of the concepts invoked in expressing them. But this would appear to imply, at the very least, that there is a mutually exclusive division into conjunctive and disjunctive contents.

Another striking feature of the extensional approach is that, assuming a modest form of the transitivity of ground, it allows us to characterize binary ground, weak or strict, purely in terms of disjunction, equivalence, and inequivalence. For $\mathbf{P} < \mathbf{Q}$ is equivalent to $\mathbf{P} \leq \mathbf{Q}$ and $\mathbf{P} \not\approx \mathbf{Q}$, and the condition that $\mathbf{P} \leq \mathbf{Q}$ now turns out equivalent to the condition that $\mathbf{P} \vee \mathbf{Q} \approx \mathbf{Q} \vee \mathbf{Q}$. To see this, note that the extensional approach renders equivalent the propositions $\mathbf{P}$, $\mathbf{Q}$ iff they have the same strict full grounds, assuming they have some strict full grounds at all. Then given transitivity, any strict full ground of $\mathbf{P}$ is a strict full ground of $\mathbf{Q}$ if $\mathbf{P} \leq \mathbf{Q}$, and hence the strict full grounds of $\mathbf{P} \vee \mathbf{Q}$ are exactly the same as the strict full grounds of $\mathbf{Q} \vee \mathbf{Q}$. This is significant in part because the very same equivalence holds on the Finean truthmaker account of ground above, and in the logic of Correia (2010). Only on the latter views, the equivalence can be further simplified, owing to the identification of a proposition with its self-disjunction. So we see that on the extensional approach

[41] The crucial observation is that if modes are extensional, the operation of fusion on the modes is idempotent, i.e. $m \sqcup m = m$. This renders the operations on closed contents of $\sqcup$ and $+$ idempotent, from which the above identities follow straightforwardly.

to modes, the view of ground that emerges stays remarkably close to the so-called worldly accounts.

6 Conclusion

In this paper, I have proposed a novel answer to the question what features of a sentence's content the notion of ground tracks: it tracks in what *modes* a content is verified. Roughly speaking, some truths Γ (strictly fully) ground a truth Q just in case verifying Γ is a mode of verifying Q. Based on this idea, I have presented an elementary formal theory of ground-theoretic content, on which a content encodes the information in what modes it may be verified or falsified. This theory, I maintain, is very well suited to serve as a framework within which a number of alternative views of ground may be articulated and examined. By way of defending this claim, I have presented a very natural and attractive account of the truth-functional operations and their interaction with ground, and I have shown how various possible views of the structural features of ground can be accommodated within my framework. Finally, I have described three views of the identity conditions of modes and some of their implications for the logic of ground, highlighting some surprising and intriguing features of the most coarse-grained of the views.

A natural next step is now to try and determine the exact pure and propositional logics of ground corresponding to a range of competing views of ground that may be implemented in the framework. In addition, various kinds of extensions of the framework may be attempted, for example by treating quantification, modal operators, and perhaps iterated ground, i.e. the grounds of truths of the form $\Gamma < P$. Already in its present form, however, the framework represents a significant step forward on the way towards a satisfactory and comprehensive theory of ground, and of ground-theoretic content.

Acknowledgements The research for this paper was funded by the Deutsche Forschungsgemeinschaft (Grant KR 4516/1-1). I gratefully acknowledge the support. Earlier versions of the material were presented at a research colloquium at the University of Hamburg, at a Senior Seminar at the Université Aix-Marseille, at the Workshop *On Ground and Consequence* in Gothenburg, at the conference *Truth and Grounds* in Ascona, and at the Mathematical Logic seminar of the Department of Mathematics at the University of Hamburg. I thank all members of the audiences for their comments and criticisms, and in particular Jon Litland as well as my fellow *phlox* members Stefan Roski and Benjamin Schnieder for their extensive and especially useful feedback. Special thanks are also due to Nick Haverkamp for very valuable comments on some of the more technical material, to Louis deRosset for a helpful email exchange on the topic, and to my wife for advice on presentational matters, as well as for a great deal of patience. The biggest debt of gratitude I owe to Kit Fine for many comments, suggestions, and discussions of the topic and of the material on which this paper is based. Finally, two anonymous referees for this journal have provided comments that have led to a number of improvements, and I am very grateful for their efforts.

Appendix

We begin with the definition of a mode-space.

Definition 1 (*Mode-spaces*) A *mode-space* is a pair $\langle M, V \rangle$ such that

1. M is non-empty

2. V is a non-empty, partial function taking non-empty sequences of non-empty subsets of M into members of M
3. the domain of V is closed under non-empty subsequences and countable concatenations of sequences
4. $V(\gamma_1 {}^\frown \gamma_2 {}^\frown \ldots) = V(\delta_1^\frown \delta_2^\frown \ldots)$ whenever $V(\gamma_1) = V(\delta_1)$, $V(\gamma_2) = V(\delta_2), \ldots$

Given a fixed mode-space $\langle M, V\rangle$, we call a *mode* any member of M. Any mode which is the value of V for some argument is called *derivative*, every other mode is called *fundamental*. We write M^D (M^F) for the set of derivative (fundamental) modes. Any subset of M will be called a *content*, and their set will be denoted by $\mathscr{C}$. The contents containing only derivative modes will themselves be called derivative, and the other contents will be called fundamental. We write $\mathscr{C}^D$ ($\mathscr{C}^F$) for the set of derivative (fundamental) contents. Any sequence for which V is defined is called a *via-sequence*. Since the domain of V is closed under non-empty subsequences, for any content P that is an element of some via-sequence, there is also the via-sequence $\langle P\rangle$ corresponding to the mode of verifying via P. We shall call any such content *raisable* and denote their set by $\mathscr{R}$. We call *a ground-set* of a derivative mode m any set of contents that underlies a content-sequence γ with $V(\gamma) = m$.

We say that a derivative mode m is *a fusion of* the sequence of derivative modes $\langle m_1, m_2, \ldots\rangle$ iff there are via-sequences $\gamma_1, \gamma_2, \ldots$ such that $V(\gamma_1) = m_1$, $V(\gamma_2) = m_2, \ldots$, and $m = V(\gamma_1 {}^\frown \gamma_2 {}^\frown \ldots)$.[42] By the third constraint on mode-spaces, fusions will always exist, and by the fourth constraint, they will be unique. For the sequence of derivative modes $\langle m_1, m_2, \ldots\rangle$, we write its fusion as $\bigsqcup\langle m_1, m_2, \ldots\rangle$ and also as $m_1 \sqcup m_2 \sqcup \ldots$ We note a first central lemma (the proof is elementary):

Lemma 1 *If $\Gamma_1, \Gamma_2, \ldots$ are ground-sets of $m_1, m_2, \ldots$, respectively, then $\Gamma_1 \cup \Gamma_2 \cup \ldots$ is a ground-set of $m_1 \sqcup m_2 \sqcup \ldots$*

For convenience, we repeat the definitions of the grounding relationships and of the truth-functions.

Definition 2 Let $\Gamma \subseteq \mathscr{C}$ and $P \in \mathscr{C}$. Then

$$\begin{aligned}
\Gamma < P &:\leftrightarrow \text{for some mode } m \in P,\ \Gamma \text{ is a ground-set of } m\\
Q \prec P &:\leftrightarrow \text{for some } m \in P,\ Q \text{ is a member of a ground-set of } m\\
\Gamma \leq P &:\leftrightarrow \Gamma = \{P\} \text{ or } \Gamma < P \text{ or } \Gamma\backslash\{P\} < P\\
Q \preceq P &:\leftrightarrow Q = P \text{ or } Q \prec P
\end{aligned}$$

As before, we write $\Gamma < \{P_1, P_2, \ldots\}$ to abbreviate that for some $\Gamma_1, \Gamma_2, \ldots, \Gamma = \Gamma_1 \cup \Gamma_2 \cup \ldots$ and $\Gamma_1 < P_1, \Gamma_2 < P_2, \ldots$, and similarly for the case of $\Gamma \leq \{P_1, P_2, \ldots\}$.

Definition 3 For $P, Q \in \mathscr{C}^D$:

$$\begin{aligned}
P \sqcup Q &:= \{m \sqcup n : m \in P \text{ and } n \in Q\}\\
P + Q &:= (P \cup Q) \cup (P \sqcup Q)
\end{aligned}$$

[42] Note that somewhat unusally, our operation of fusion applies to sequences of modes and is thereby potentially sensitive to order and repetition of the modes being fused.

Definition 4 For $P, Q \in \mathscr{R}$:

$$\begin{array}{lcl} \uparrow P & := & \{V\langle P\rangle\} \text{ if } P \cap M^D \text{ is empty, } \{V\langle P\rangle\} + (P \cap M^D) \text{ otherwise} \\ P \wedge Q & := & \uparrow P \sqcup \uparrow Q \\ P \vee Q & := & \uparrow P + \uparrow Q \end{array}$$

We shall mainly be interested in mode-spaces in which the set of raisable contents is closed under these operations.

Definition 5 A mode-space is called *complete* iff $P \wedge Q \in \mathscr{R}$, $P \vee Q \in \mathscr{R}$, and $\uparrow P \in \mathscr{R}$ whenever $P, Q \in \mathscr{R}$.

Lemma 2 (Introduction Lemma) *In any complete mode-space, for $\Gamma \subseteq \mathscr{R}$ and $P, Q \in \mathscr{R}$:*

1. *If $\Gamma \leq P$, then $\Gamma <\uparrow P$.*
2. *If $\Gamma < \{P, Q\}$, then $\Gamma < P \sqcup Q$.*
3. *If $\Gamma < P$ or $\Gamma < Q$ or $\Gamma < \{P, Q\}$, then $\Gamma < P + Q$*
4. *If $\Gamma \leq \{P, Q\}$, then $\Gamma < P \wedge Q$*
5. *If $\Gamma \leq P$ or $\Gamma \leq Q$ or $\Gamma < \{P, Q\}$, then $\Gamma < P \vee Q$.*

Proof For 1: Suppose $\Gamma \leq P$. Then either (i) $\Gamma = \{P\}$, or (ii) $\Gamma < P$, or (iii) $\Gamma \setminus \{P\} < P$. Suppose (i). By definition of $\updownarrow$ $V\langle P\rangle \in\uparrow P$. Since $\{P\}$ is the set underlying $\langle P\rangle$, it is a ground-set of $V\langle P\rangle$, so $\Gamma < P$. Suppose (ii). Then Γ is a ground-set of some mode $m \in P$. But then by definition of $\updownarrow$ it follows that $m \in \uparrow P$, and hence $\Gamma < P$. Suppose (iii). Then for some $m \in P$, $\Gamma \setminus \{P\}$ is a ground-set of m. But then by definition of $\updownarrow$ it follows that $V\langle P\rangle \sqcup m \in\uparrow P$. By Lemma 1, $\Gamma \setminus \{P\} \cup \{P\} = \Gamma$ is a ground-set of $V\langle P\rangle \sqcup m$, hence $\Gamma < P$.

For 2: Suppose $\Gamma < \{P, Q\}$. Let $\Gamma_P < P$ and $\Gamma_Q < Q$ with $\Gamma_P \cup \Gamma_Q = \Gamma$. Then Γ_P is a ground-set of some mode $m_P \in P$ and Γ_Q is a ground-set of some mode $m_Q \in Q$. By definition of $\sqcup$, $m_P \sqcup m_Q \in P \sqcup Q$, and by Lemma 1, $\Gamma_P \cup \Gamma_Q = \Gamma$ is a ground-set of $m_P \sqcup m_Q$, hence $\Gamma < P \sqcup Q$.

For 3: Suppose $\Gamma < P$. Then there is a mode $m \in P$ of which Γ is a ground-set. By definition of $+$, the same mode is included in $P + Q$, hence $\Gamma < P + Q$. Suppose $\Gamma < Q$. Then by the same reasoning, $\Gamma < P + Q$. Finally, suppose $\Gamma < \{P, Q\}$. Then by part 2., $\Gamma < P \sqcup Q$, and hence by definition of $+$, $\Gamma < P + Q$.

For 4: Suppose $\Gamma \leq \{P, Q\}$. Let $\Gamma_P \leq P$ and $\Gamma_Q \leq Q$ with $\Gamma_P \cup \Gamma_Q = \Gamma$. By part 1., $\Gamma_P <\uparrow P$ and $\Gamma_Q <\uparrow Q$, hence $\Gamma < \{\uparrow P, \uparrow Q\}$, and by part 2., $\Gamma < P \wedge Q$.

For 5: If $\Gamma \leq P$ or $\Gamma \leq Q$, then by part 2., $\Gamma <\uparrow P$ or $\Gamma <\uparrow Q$. If $\Gamma \leq \{P, Q\}$, then by the reasoning in part 4., $\Gamma < \{\uparrow P, \uparrow Q\}$. So by part 3., $\Gamma < P \vee Q$. □

Definition 6 A mode-space is called *constrained* iff $V(\gamma) = V(\delta)$ only if the same ground-set corresponds to γ and δ.

In a constrained mode-space, every derivative mode m corresponds to a unique ground-set, which we denote by $|m|$.

Lemma 3 (Elimination Lemma) *In any constrained and complete mode-space, for $\Gamma \subseteq \mathscr{R}$ and $P, Q \in \mathscr{R}$:*

1. *If* $\Gamma < \uparrow P$*, then* $\Gamma \leq P$
2. *If* $\Gamma < P \sqcup Q$*, then* $\Gamma < \{P, Q\}$
3. *If* $\Gamma < P + Q$*, then* $\Gamma < P$ *or* $\Gamma < Q$ *or* $\Gamma < \{P, Q\}$
4. *If* $\Gamma < P \wedge Q$*, then* $\Gamma \leq \{P, Q\}$
5. *If* $\Gamma < P \vee Q$*, then* $\Gamma \leq P$ *or* $\Gamma \leq Q$ *or* $\Gamma \leq \{P, Q\}$

Proof For 1: Suppose $\Gamma < \uparrow P$. Let $m \in \uparrow P$ with $\Gamma = |m|$. By definition of $\updownarrow$ there are three cases. (i) $m \in P \cap M^D$. Then $\Gamma < P$ and hence $\Gamma \leq P$. (ii) $m = V\langle P\rangle$. Then $\Gamma = |V\langle P\rangle| = \{P\}$, so again $\Gamma \leq P$. (iii) $m = V\langle P\rangle \sqcup n$ for some $n \in P \cap M^D$. Since $n \in P$, $|n| < P$. By Lemma 1, $\Gamma = |n| \cup \{P\}$. Then either $|n| = \Gamma$ or $|n| = \Gamma \backslash \{P\}$, so either $\Gamma < P$ or $\Gamma \backslash \{P\} < P$, and hence $\Gamma \leq P$.

For 2: Suppose $\Gamma < P \sqcup Q$. Let $|m| = \Gamma$ and $m \in P \sqcup Q$. Then $m = m_P \sqcup m_Q$ for some $m_P \in P$ and $m_Q \in Q$. So $|m_P| < P$ and $|m_Q| < Q$. But $\Gamma = |m_P| \cup |m_Q|$, and hence $\Gamma < \{P, Q\}$.

For 3: Suppose $\Gamma < P + Q$. Then it is immediate from the definition of $+$ that $\Gamma < P$ or $\Gamma < Q$ or $\Gamma < P \sqcup Q$, which by part 2. implies $\Gamma < \{P, Q\}$.

For 4: Suppose $\Gamma < P \wedge Q$. By part 2., $\Gamma < \{\uparrow P, \uparrow Q\}$. So let $\Gamma_P < \uparrow P$ and $\Gamma_Q < \uparrow Q$ with $\Gamma = \Gamma_P \cup \Gamma_Q$. By part 1., $\Gamma_P \leq P$ and $\Gamma_Q \leq Q$, hence $\Gamma \leq \{P, Q\}$.

For 5: Suppose $\Gamma < P \vee Q$. By part 3., there are three cases. (i) $\Gamma < \uparrow P$. Then by part 1., $\Gamma \leq P$. (ii) $\Gamma < \uparrow Q$. Then by part 1. again, $\Gamma \leq Q$. (iii) $\Gamma \leq \{\uparrow P, \uparrow Q\}$. Then by the reasoning in part 4., $\Gamma \leq \{P, Q\}$. □

We move on to the case of bilateral contents. We define the truth-functional operations on bilateral contents as well as the notion of ground-theoretic equivalence ($\approx$).

Definition 7 For $\mathbf{P}, \mathbf{Q} \in \mathscr{R} \times \mathscr{R}$:

$$\begin{array}{lll} \neg\mathbf{P} & := & \langle \mathbf{P}^-, \uparrow\mathbf{P}^+ \rangle \\ \mathbf{P} \wedge \mathbf{Q} & := & \langle \mathbf{P}^+ \wedge \mathbf{Q}^+, \mathbf{P}^- \vee \mathbf{Q}^- \rangle \\ \mathbf{P} \vee \mathbf{Q} & := & \langle \mathbf{P}^+ \vee \mathbf{Q}^+, \mathbf{P}^- \wedge \mathbf{Q}^- \rangle \end{array}$$

Definition 8 For $\mathbf{P}, \mathbf{Q} \in \mathscr{R} \times \mathscr{R}$: $\mathbf{P} \approx \mathbf{Q} :\leftrightarrow$ for all $\boldsymbol{\Gamma}$: $\boldsymbol{\Gamma} < \mathbf{P}$ iff $\boldsymbol{\Gamma} < \mathbf{Q}$, and for all $\boldsymbol{\Delta}$ and $\mathbf{R}$: $\boldsymbol{\Delta}, \mathbf{P} < \mathbf{R}$ iff $\boldsymbol{\Delta}, \mathbf{Q} < \mathbf{R}$.

Recall that ground between bilateral contents is defined simply as ground between the positive components. As a result, bilateral contents will be *ground-theoretically equivalent* (written $\approx$) provided their positive components are the same.

Lemma 4 (DeMorgan) *For* $\mathbf{P}, \mathbf{Q} \in \mathscr{R} \times \mathscr{R}$*:*

1. $\neg(\mathbf{P} \wedge \mathbf{Q}) \approx \neg\mathbf{P} \vee \neg\mathbf{Q}$
2. $\neg(\mathbf{P} \vee \mathbf{Q}) \approx \neg\mathbf{P} \wedge \neg\mathbf{Q}$

Proof By application of the definitions.

For 1: $(\neg(\mathbf{P} \wedge \mathbf{Q}))^+ = (\mathbf{P} \wedge \mathbf{Q})^- = \mathbf{P}^- \vee \mathbf{Q}^- = (\neg\mathbf{P})^+ \vee (\neg\mathbf{Q})^+ = (\neg\mathbf{P} \vee \neg\mathbf{Q})^+$

For 2: $(\neg(\mathbf{P} \vee \mathbf{Q}))^+ = (\mathbf{P} \vee \mathbf{Q})^- = \mathbf{P}^- \wedge \mathbf{Q}^- = (\neg\mathbf{P})^+ \wedge (\neg\mathbf{Q})^+ = (\neg\mathbf{P} \wedge \neg\mathbf{Q})^+$
□

As an immediate consequence of the definition of ground on bilateral contents and the previous lemmata, we obtain

Theorem 5 (Truth-functions and ground, introduction) *In any complete mode-space, for* $\boldsymbol{\Gamma} \subseteq \mathscr{R} \times \mathscr{R}$ *and* $\mathbf{P}, \mathbf{Q} \in \mathscr{R} \times \mathscr{R}$*:*

1. *If* $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$*, then* $\boldsymbol{\Gamma} < \mathbf{P} \wedge \mathbf{Q}$
2. *If* $\boldsymbol{\Gamma} \leq \mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \mathbf{Q}$ *or* $\boldsymbol{\Gamma} < \{\mathbf{P}, \mathbf{Q}\}$*, then* $\boldsymbol{\Gamma} < \mathbf{P} \vee \mathbf{Q}$.
3. *If* $\boldsymbol{\Gamma} \leq \mathbf{P}$*, then* $\boldsymbol{\Gamma} < \neg\neg\mathbf{P}$
4. *If* $\boldsymbol{\Gamma} \leq \neg\mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \neg\mathbf{Q}$ *or* $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$*, then* $\boldsymbol{\Gamma} < \neg(\mathbf{P} \wedge \mathbf{Q})$
5. *If* $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$*, then* $\boldsymbol{\Gamma} < \neg(\mathbf{P} \vee \mathbf{Q})$

Theorem 6 (Truth-functions and ground, elimination) *In any complete and constrained mode-space, for* $\boldsymbol{\Gamma} \subseteq \mathscr{R} \times \mathscr{R}$ *and* $\mathbf{P}, \mathbf{Q} \in \mathscr{R} \times \mathscr{R}$*:*

1. *If* $\boldsymbol{\Gamma} < \mathbf{P} \wedge \mathbf{Q}$*, then* $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
2. *If* $\boldsymbol{\Gamma} < \mathbf{P} \vee \mathbf{Q}$*, then* $\boldsymbol{\Gamma} \leq \mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \mathbf{Q}$ *or* $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
3. *If* $\boldsymbol{\Gamma} < \neg\neg\mathbf{P}$*, then* $\boldsymbol{\Gamma} \leq \mathbf{P}$
4. *If* $\boldsymbol{\Gamma} < \neg(\mathbf{P} \wedge \mathbf{Q})$*, then* $\boldsymbol{\Gamma} \leq \neg\mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \neg\mathbf{Q}$ *or* $\Gamma \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$
5. *If* $\boldsymbol{\Gamma} < \neg(\mathbf{P} \vee \mathbf{Q})$*, then* $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$

We turn now to the structural properties of ground.

Definition 9 A proposition P is

- *irreflexive* iff: P does not occur in γ whenever $V(\gamma) \in P$,
- *closed* iff: P contains a mode with ground-set $\Gamma_1 \cup \Gamma_2 \cup \ldots$ whenever P contains modes with ground-sets $\Gamma_1, \Gamma_2, \ldots$,
- *transitive* iff: P includes a mode with ground-set Γ, Δ whenever P includes a mode with ground-set Δ, Q and Q includes a mode with ground-set Γ.
- *normal* iff: irreflexive, closed, and transitive.

Theorem 7 (Structural Principles for Normal Propositions) *In any constrained mode-space, for normal propositions* $P, P_1, P_2, \ldots, Q, R$ *and sets of normal propositions* $\Gamma, \Gamma_1, \Gamma_2, \ldots, \Delta$*:*

1. $P \not< P$
2. *If* $\Gamma_1 < P$, $\Gamma_2 < P$, ..., *then* $\Gamma_1, \Gamma_2, \ldots < P$.
3. *If* $\Gamma \leq P$ *and* $Q \prec P$ *for all* $Q \in \Gamma$*, then* $\Gamma < P$.
4. *If* $\Gamma, P \leq P$*, then* $\Gamma \leq P$.
5. *If* $\Gamma_1 \leq P, \Gamma_2 \leq P, \ldots$*, then* $\Gamma_1 \cup \Gamma_2 \cup \ldots \leq P$.
6. *If* $\Gamma < P$ *and* $\Delta, P < Q$*, then* $\Gamma, \Delta < Q$.
7. *If* $\Gamma_1 \leq P_1$, $\Gamma_2 \leq P_2$, ..., *and* $P_1, P_2, \ldots \leq Q$ *then* $\Gamma_1, \Gamma_2, \ldots \leq Q$
8. *If* $P \preceq Q$ *and* $Q \prec R$ *then* $P \prec R$
9. *If* $P \prec Q$ *and* $Q \preceq R$ *then* $P \prec R$
10. *If* $P \preceq Q$ *and* $Q \preceq R$ *then* $P \preceq R$

Proof 1–4., and 6. were established in Sect. 5.4 above.

For 5., suppose $\Gamma_1 \leq P, \Gamma_2 \leq P, \ldots$ Consider all the Γ_i which are distinct from $\{P\}$. By definition of $\leq$ and 1., $\Gamma_i \setminus \{P\} < P$ for any such Γ_i. If there are any such Γ_i, let Γ' be the result of removing P from their union. By 2, $\Gamma' < P$, so $\Gamma_1 \cup \Gamma_2 \cup \ldots = \Gamma' \cup \{P\} \leq P$. If there are no such Γ_i, then $\Gamma_1 \cup \Gamma_2 \cup \ldots = \{P\} \leq P$.

For 7., suppose $\Gamma_1 \leq P_1, \Gamma_2 \leq P_2, \ldots$, and $P_1, P_2, \ldots \leq Q$. We confine ourselves to showing that $\Gamma_1 \cup \{P_2, \ldots\} \leq Q$ follows. By applying the same reasoning repeatedly and making use of the reflexivity of $\leq$, we may establish the desired conclusion. Now either (a) $\Gamma_1 = \{P_1\}$, or (b) $\Gamma_1 < P_1$, or (c) $\{P_1\} \subset \Gamma_1$ and $\Gamma_1 \setminus \{P_1\} < P_1$. If (a), then our intended result $\Gamma_1 \cup \{P_2, \ldots\} \leq Q$ follows immediately. Suppose that (b). Then if (b1) $P_1 = Q$, we have $\Gamma_1 < Q$ and hence $\Gamma_1 \leq Q$. Moreover, by 4., we have $\{P_2, \ldots\} \leq Q$. So by 5., $\Gamma_1 \cup \{P_2, \ldots\} \leq Q$ follows. But if (b2) $P_1 \neq Q$, then $P_1 \in \{P_1, P_2, \ldots\} \setminus \{Q\}$ and $\{P_1, P_2, \ldots\} \setminus \{Q\} < Q$, so by 6., $\Gamma_1 \cup \{P_2, \ldots\} \setminus \{Q\} < Q$, hence $\Gamma_1 \cup \{P_2, \ldots\} \leq Q$. Suppose finally that (c). Then if (c1) $P_1 = Q$, $\Gamma_1 \leq Q$, and so $\Gamma_1 \cup \{P_2, \ldots\} \leq Q$ follows as in case (b1). But if (c2) $P_1 \neq Q$, then by similar reasoning as in case (b2), $\Gamma_1 \setminus \{P_1\} \cup \{P_1, P_2, \ldots\} = \Gamma_1 \cup \{P_2, \ldots\} \leq Q$.

For 8., suppose $P \preceq Q$ and $Q \prec R$. If $P = Q$, then $P \prec R$ follows immediately. So suppose $P \neq Q$, and hence $P \prec Q$. Let $m \in Q$ be such that $P \in |m|$, and let $n \in R$ be such that $Q \in |n|$. Then $|m| \cup \{P\} < Q$ and $|n| \cup \{Q\} < R$, so by 6, $|n| \cup |m| \cup \{P\} < R$, so $P \prec R$.

For 9., suppose $P \prec Q$ and $Q \preceq R$. If $Q = R$, then $P \prec R$ follows immediately. So suppose $Q \neq R$, and hence $Q \prec R$. Then by the same reasoning as before, $P \prec R$.

For 10., suppose $P \preceq Q$ and $Q \preceq R$. If either $P \prec Q$ or $Q \prec R$, then it follows by the previous results that $P \prec R$, and hence $P \preceq R$. If neither $P \prec Q$ nor $Q \prec R$, then $P = Q$ and $Q = R$, hence $P = R$, and thus again $P \preceq R$. □

Together with the Subsumption principles and the Identity principle $P \leq P$ for weak ground established in Sect. 5.2, parts 1, 3, 7–10 correspond to the basic rules of the logic proposed in Fine (2012c), which is thereby shown to be sound with respect to the class of normal propositions in a constrained mode-space.

We call a bilateral content irreflexive, closed, transitive, or normal, just in case its positive component has the relevant property. We wish then to show that any truth-functional combinations of normal propositions are themselves normal. By the definitions of the truth-functional operations, it suffices to show that for unilateral contents, normality is preserved under $\wedge$, $\vee$, and $\updownarrow$

Theorem 8 *For $P, Q \in \mathscr{R}$, in some constrained mode-space: If P and Q are normal, then so are $P \wedge Q$, $P \vee Q$, and $\uparrow P$.*

Proof By reducing failures for $P \wedge Q$, $P \vee Q$, $\uparrow P$ of the characteristic principles of irreflexivity, amalgamation, and transitivity (1, 2 and 6 in Theorem 7) to failures of the corresponding principles for P or Q. We give the proof for $P \wedge Q$; the other cases may be established by parallel means.

For irreflexivity, suppose that $P \wedge Q \prec P \wedge Q$. Then for some Γ: $\Gamma, P \wedge Q < P \wedge Q$. Then by the Elimination Lemma, $\Gamma, P \wedge Q \leq \{P, Q\}$, and hence either (i) $P \wedge Q \preceq P$ or (ii) $P \wedge Q \preceq Q$. Suppose (i). Then since $P \prec P \wedge Q$, it follows by transitivity of P that $P \prec P$, in contradiction to P's irreflexivity.

For amalgamation, suppose that $\Gamma_1 < P \wedge Q, \Gamma_2 < P \wedge Q, \ldots$. For each Γ_i: $\Gamma_i \leq \{P, Q\}$. So let $\Gamma_{i_P} \leq P$ and $\Gamma_{i_Q} \leq Q$ with $\Gamma_i = \Gamma_{i_P} \cup \Gamma_{i_Q}$ for all i. Then by weak ground amalgamation for P and Q, $\Gamma_{1_P}, \Gamma_{2_P}, \ldots \leq P$, and $\Gamma_{1_Q}, \Gamma_{2_Q}, \ldots \leq Q$, so $\Gamma_1, \Gamma_2, \ldots \leq \{P, Q\}$, and hence $\Gamma_1, \Gamma_2, \ldots < P \wedge Q$.

For transitivity, suppose that $\Delta, R < P \wedge Q$ and $\Gamma < R$. Then $\Delta, R \leq \{P, Q\}$. So we may write Δ, R as the union of weak full grounds of respectively P and Q. R

will be a member of at least one of these. By replacing R with Γ, using the transitivity of P and Q, we may infer $\Gamma, \Delta < \{P, Q\}$, and hence $\Gamma, \Delta < P \wedge Q$. □

References

Batchelor, R. (2010). Grounds and consequences. *Grazer Philosophische Studien*, *80*, 65–77.
Correia, F. (2010). Grounding and truth-functions. *Logique et Analyse*, *53*(211), 251–279.
Correia, F. (2014b). Logical grounds. *The Review of Symbolic Logic*, *7*(1), 31–59.
Correia, F., & Schnieder, B. (2012). Grounding: An opinionated introduction. In F. Correia & B. Schnieder (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (pp. 1–36). Cambridge: Cambridge University Press.
deRosset, L. (2013). What is weak ground? *Essays in Philosophy*, *14*(1), 7–18.
deRosset, L. (2014). On weak ground. *The Review of Symbolic Logic*, *7*(4), 713–744.
Dixon, T. S. (2016). What is the well-foundedness of grounding? *Mind*, *125*(498), 439–468.
Fine, K. (2010). Some puzzles of ground. *Notre Dame Journal of Formal Logic*, *51*(1), 97–118.
Fine, K. (2012a). Counterfactuals without possible worlds. *The Journal of Philosophy*, *109*(3), 221–246.
Fine, K. (2012b). Guide to ground. In B. Schnieder & F. Correia (Eds.), *Metaphysical Grounding* (pp. 37–80). Cambridge: Cambridge University Press.
Fine, K. (2012c). The pure logic of ground. *The Review of Symbolic Logic*, *5*(1), 1–25.
Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, *43*, 549–577.
Fine, K. (2016). Angellic content. *Journal of Philosophical Logic*, *45*(2), 199–226.
Jenkins, C. (2011). Is metaphysical dependence irreflexive? *The Monist*, *94*(2), 267–276.
Krämer, S. (2013). A simpler puzzle of ground. Thought: A. *Journal of Philosophy*, *2*(2), 85–89.
Krämer, S., & Roski, S. (2015). A note on the logic of worldly ground. Thought: A. *Journal of Philosophy*, *4*(1), 59–68.
Krämer, S., & Roski, S. (2016). Difference-making grounds. *Philosophical Studies*, pp. 1–25.
Litland, J. E. (2013). On some counterexamples to the transitivity of grounding. *Essays in Philosophy*, *14*(1), 19–32.
Litland, J. E. (2016a). An infinitely descending chain of ground without a lower bound. *Philosophical Studies*, *173*(5), 1361–1369.
Litland, J. E. (2016). Pure logic of many-many ground. *Journal of Philosophical Logic*, pp 1–47. doi:10.1007/s11229-015-0923-x.
Poggiolesi, F. (2015). On defining the notion of complete and immediate formal grounding. *Synthese*, pp. 1–21.
Raven, M. (2013). Is ground a strict partial order? *American Philosophical Quarterly*, *50*(2), 193–201.
Reboul, A. (Ed.). (2014a). From grounding to truth-making: Some thoughts. *Mind, Values, and Metaphysics: Philosophical Essays in Honour of Kevin Mulligan* (Vol. 1, pp. 85–98). New York: Springer. Originally published online in 2011.
Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffmann (Eds.), *Modality: Metaphysics, Logic, and Epistemology* (pp. 109–136). Oxford: Oxford University Press.
Schaffer, J. (2009). On what grounds what. In D. J. Chalmers, D. Manley, & R. Wassermann (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology* (pp. 347–383). Oxford: Oxford University Press.
Schaffer, J. (2012). Grounding, transitivity, and contrastivity. In B. Schnieder & F. Correia (Eds.), *Metaphysical Grounding* (pp. 122–138). Cambridge: Cambridge University Press.
Schnieder, B. (2008). 'By': A refutation of the Anscombe Thesis. *Linguistics and Philosophy*, *31*(6), 649–669.
Schnieder, B. (2010). A Puzzle about 'Because'. *Logique et Analyse*, *53*(211), 317–343.
Schnieder, B. (2011). A logic for because. *The Review of Symbolic Logic*, *4*(3), 445–465.
Williamson, T. (1994). *Vagueness*. London: Routledge.

Synthese (2021) 198:1643–1683
https://doi.org/10.1007/s11229-019-02154-4

Ground-theoretic equivalence

Stephan Krämer[1]

Received: 13 July 2018 / Accepted: 20 February 2019 / Published online: 17 June 2019

Abstract
Say that two sentences are ground-theoretically equivalent iff they are interchangeable salva veritate in grounding contexts. Notoriously, ground-theoretic equivalence is a hyperintensional matter: even logically equivalent sentences may fail to be interchangeable in grounding contexts. Still, there seem to be some substantive, general principles of ground-theoretic equivalence. For example, it seems plausible that any sentences of the form $A \wedge B$ and $B \wedge A$ are ground-theoretically equivalent. What, then, are in general the conditions for two sentences to stand in the relation of ground-theoretic equivalence, and what are the logical features of that relation? This paper develops and defends an answer to these questions based on the mode-ified truthmaker theory of content presented in my recent paper 'Towards a theory of ground-theoretic content' (Krämer in Synthese 195(2):785–814, 2018).

Keywords Ground · Content · Equivalence · Truthmaking · Logic

1 Introduction

According to a widely held metaphysical picture, the world is not a mere aggregate of facts, but rather a *structured* whole, with some of its members obtaining *in virtue of* some other members. In such a case, the facts in virtue of which some other fact obtains are then said to collectively *(metaphysically) ground* the latter fact. Recently, considerable efforts have been made to try and clarify this picture by working out the

Part of the research for this paper was funded by the Deutsche Forschungsgemeinschaft (Grants KR 4516/1-1 and KR 4516/2-1). I gratefully acknowledge the support.

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

1 Department of Philosophy, University of Hamburg, Überseering 35, Postfach #4, 22297 Hamburg, Germany

general theory and logic of the grounding relations.[1] The present paper contributes to this project by addressing the question of *ground-theoretic equivalence*—roughly, the question under which conditions two sentences are interchangeable salva veritate in a grounding statement.

The question is important. First of all, it is of great significance for the general theory of ground, since one's options in developing the latter are often constrained by how one answers the former. For a particularly simple instance of this, note that given the widely accepted assumption that nothing helps grounds itself, maintaining that A and B are ground-theoretically equivalent bars one from holding that either helps ground the other. The question of ground-theoretic equivalence may also have an impact on more general issues about the nature of the grounding relation and its relata. For instance, it has been suggested that a very fine-grained view of ground-theoretic equivalence is incompatible with a conception of ground as a *worldly* phenomenon, but requires us to see ground as relating *representations* of the world rather than the world itself.[2] Beyond those connections within the theory of ground, the question of ground-theoretic equivalence may also have wider implications. For example, McDaniel (2015) has argued that we should individuate propositions in general in terms of their place in the network of grounding relations. On this view, the question of ground-theoretic equivalence turns into the more general question of propositional identity.

There is a broad consensus in the current debate on certain *partial* answers to the question of ground-theoretic equivalence. In particular, most participants in that debate agree that the logical equivalence of two sentences is *not* sufficient to render them ground-theoretically equivalent. For example, if A is true, then this is taken to entail that A grounds $A \vee \neg A$, but not that A grounds $B \vee \neg B$—even though $A \vee \neg A$ and $B \vee \neg B$ are of course logically equivalent. On the other hand, most parties will also agree that there are *some* non-trivial cases of ground-theoretic equivalence. Perhaps the least contentious examples arise from the commutativity of conjunction and disjunction: pairs of the form $A \wedge B$ and $B \wedge A$, or $A \vee B$ and $B \vee A$, are presumably always ground-theoretically equivalent.

But beyond these points, there is no consensus as regards the conditions that are necessary and sufficient for ground-theoretic equivalence. Indeed, it is not so much that people disagree about what conditions are necessary and what conditions are sufficient. The problem is rather that so far, the published literature contains only very few general and appropriately developed answers to the question. As a result, there is at present no clear understanding of the range of potentially viable answers, let alone their respective merits or demerits.

My aim in this paper is to improve on this situation in two ways. First, I shall develop and defend a novel account of ground-theoretic equivalence, based on the theory of ground-theoretic content that I recently put forward in Krämer (2018). To that end, I first clarify the question of ground-theoretic equivalence and lay down some desiderata on a satisfactory answer (Sect. 2). I then briefly introduce my theory of

[1] The central pioneering contributions to the study of the logic of ground are Batchelor (2010), Correia (2010), Fine (2010, 2012a, b), Rosen (2010) and Schnieder (2011). More recent work includes Correia (2014, 2016, 2017), deRosset (2013, 2014), Krämer (2013, 2018), Krämer and Roski (2015, 2017), Litland (2016), Poggiolesi (2016b, 2018).

[2] This has been urged in particular by Fabrice Correia, cf. his (2010, pp. 256f, 264ff, 2017, p. 508).

ground-theoretic content, describe the range of possible accounts of ground-theoretic equivalence that can be formulated within that theory, and provide some reasons for favouring a particular one of them (Sect. 3). Second, I shall compare the account to rival approaches and try to clarify their interrelations, as well as highlight what I take to be distinctive strengths of my own view (Sect. 4). I have aimed to keep the discussion in the main text as informal and accessible as possible without significantly compromising on accuracy. All technical details, including soundness and completeness proofs, can be found in a formal appendix.

2 Preliminaries

To a first approximation, and in keeping with common practice in the debate, ground will here be understood as the relation of one fact obtaining in virtue of others. Paradigmatic examples include: the fact that the ball is scarlet grounds the fact that the ball is red, the fact that the ball is red and the fact that the ball is round jointly ground the fact that the ball is red and round, and the fact that the ball is round grounds the fact that the ball is round or square. When some facts jointly ground another, each fact among the former may be said to *partially* ground the latter, and we may speak of a fact's or some facts' *fully* grounding another to distinguish the primary, full sense of ground from the partial one.

Two qualifications are in order. First, there is disagreement in the current debate concerning whether ground should be seen as relating facts or rather truths. Second, there is disagreement as to whether ground should, strictly speaking, be conceived as a relation at all. Some authors are of the view that ground is best expressed by means of a sentential connective akin to 'because' rather than by a relational predicate like 'ground(s)' and thus see no need to countenance a genuine relation of grounding.[3] For the purposes of this paper, we need not decide either issue. Informally and to facilitate presentation, I shall continue to speak in the relational mode, and describe ground as relating either facts or truths. Formally, I shall use sentential connectives symbolizing ground, for example writing $P_1, P_2, \ldots < Q$ to say that the fact that P_1, the fact that P_2, and … fully ground the fact that Q.

Above, I introduced the question of ground-theoretic equivalence as, roughly, the question under what conditions two sentences are interchangeable salva veritate in a grounding statement. I now wish to refine this question somewhat. First of all, the question I shall focus on is when sentences are interchangeable in an *argument place* of an expression of ground. The conditions under which this holds may differ from the conditions under which sentences are exchangeable when embedded within a larger sentence that occupies such an argument place. Second, the expressions of ground adverted to should be thought of as connectives like $<$, whose argument places may be filled by one or more sentences.[4] Third, I shall only be interested in cases when

[3] This view is held, for example, by Fine (2012a) and by Correia (2010). Rosen (2010) and Audi (2012) are among those defending the opposing view.

[4] If the predicational mode of formulating statements of ground is preferred, the argument places will be filled with one or more singular terms for facts, and the canonical form of the latter would be 'the fact that

sentences are so interchangeable for a particular kind of reason. For suppose that two sentences A and B are interchangeable salva veritate in the argument places of $<$. Then this may be so for at least three quite different kinds of reasons, and to obtain a natural and fruitful understanding of ground-theoretic equivalence, we should abstract from two of them. First, A and B may be interchangeable simply because they are both *false*. For in that case, due to the factivity of ground, any statements of the form $\Gamma, A < C$ and $\Gamma < A$ will be false, and so will the results $\Gamma, B < C$ and $\Gamma < B$ of replacing A by B. Second, A and B may be interchangeable even though the facts they state differ with respect to their grounds, or with respect to what they ground, simply because these differences cannot be expressed in the language under consideration. Thus, it may be that the fact stated by A is fully grounded by the fact that P, and that this fact is not a full ground of the fact stated by B, but since the background language lacks a sentence stating the fact that P, A and B are nevertheless freely interchangeable in the argument places of $<$ without change of truth-value. Third, it may be that A and B are alike in every respect of meaning or content to which ground is sensitive. This seems to hold, for example, of any pair of the form $A \wedge B$ and $B \wedge A$. To the extent that these differ semantically at all, the differences they exhibit seem to be of a kind to which ground is indifferent. I shall call sentences ground-theoretically equivalent iff they are interchangeable for this third kind of reason.

Having clarified the content of the question of ground-theoretic equivalence, let us consider what might be reasonable desiderata for an adequate answer to the question. Ideally, it seems, an answer would be both fully general and maximally instructive. That is, it would specify necessary and sufficient conditions for *any* two sentences to be ground-theoretically equivalent (full generality), and the conditions it specifies would be such that we can easily ascertain with respect to any particular two sentences whether they are satisfied (maximal instructiveness). However, a single answer with both features is too much to hope for. A fully general answer would need to apply to any pair of sentences from any language, but a maximally instructive answer can in general be expected at best for sentence-pairs from some well-defined formal language.

The next best thing would consist in two complementary answers, one achieving full generality at the cost of instructiveness, the other achieving maximal instructiveness at the cost of generality. This is analogous to what we have in the case of intensional equivalence. On the one hand, we can say with full generality that any given two sentences are intensionally equivalent just in case they are true in the same possible worlds. The notion of truth in a possible world has some intuitive content, so this is at least somewhat helpful. At the same time, it is not in general straightforward to ascertain, with respect to an arbitrary pair of sentences, whether the condition of truth in the same possible worlds is indeed satisfied.[5] The answer is therefore not maximally instructive. On the other hand, we can give a maximally instructive answer for certain special cases of the question. For instance, any standard propositional

Footnote 4 continued
P'. Our question would then have to be aimed at the interchangeability of two sentences in the position of 'P' in the canonical fact-designators.

[5] For example, it is extremely difficult to ascertain whether '$2+2=4$' is true in the same worlds as 'every even integer greater than 2 can be written as the sum of two primes', a sentence expressing the famously unsolved Goldbach conjecture.

modal logic provides us with a mechanical procedure for ascertaining whether two sentences of the relevant propositional language are intensionally equivalent in virtue of their logical form (under the account of intensional equivalence for which the logic is adequate).[6]

I shall aim for an account that is in this way analogous to the account of intensional equivalence. That is, I shall demand of an adequate answer to the question of ground-theoretic equivalence that it comprise two components: one general answer, analogous to the claim that two sentences are intensionally equivalent iff true in exactly the same possible worlds, and one restricted but formally precise answer, describing a deductive system allowing us to derive exactly those statements of ground-theoretic equivalence between pairs of a formal propositional language that obtain in virtue of the logical form of the sentences involved.[7] The next section develops an answer that meets these conditions based on the mode-ified truthmaker theory of content (Krämer 2018).[8]

3 The Mode-ified truthmaker account

For a sentence to be ground-theoretically equivalent to another is for the two to be semantically alike in every respect to which ground is sensitive. But to what semantic features of a sentence is ground sensitive? On the view developed in Krämer (2018), the answer is that ground is sensitive to (a) what states verify a sentence and (b) *how*, or in what *modes*, they do so. Sect. 3.1 clarifies how the key notions of verification, states, and modes are understood within modified truthmaker theory. On that basis, Sect. 3.2 then develops an answer to the general question of ground-theoretic equivalence, before Sect. 3.3 answers the formal question.

3.1 Mode-ified truthmaking

As its name suggests, the mode-ified truthmaker theory of content is a modification of the truthmaker theory of content developed by Kit Fine (see e.g. his 2016, 2017a, b). Most of the details of the truthmaker account are not essential for understanding the mode-fied version. It suffices to note the following basic points. The key concept of the approach is that of a proposition being (exactly[9]) verified or falsified by a state. A state is conceived as a part, or fragment, of a world, distinguished from the latter by the fact that it need not be complete, but may leave open the truth-value of many

[6] Thanks to an anonymous referee for pressing me to elaborate on the claims made in this and the previous paragraph.

[7] One might, of course, reasonably hope for a somewhat more general formal answer, perhaps one that is adequate for a richer language or which allows us to derive all the statements of ground-theoretic equivalence entailed by an arbitrary set of premises. For the purposes of this paper, however, we do better to focus on the simpler class of cases. This will simplify the formal side of things and it will facilitate comparison with previous relevant work such as Correia (2016) and Fine (2016), which also focuses on this class.

[8] The label ‘mode-ified’ is motivated by the fact that the theory crucially distinguishes different *modes* of truthmaking.

[9] Fine distinguishes several notions of truthmaking of which the exact notion is the narrowest or most demanding one. The other notions are of no special importance for our purposes. When speaking of truthmakers or verification, I shall henceforth always mean the exact variety.

propositions. The space of states comes equipped with a relation of part-whole, and given any set of states $T = \{s_1, s_2, \ldots\}$ we may form the fusion $\bigsqcup T = s_1 \sqcup s_2 \sqcup \ldots$ of its members, which is the smallest state of which they are all part. The operation of fusion is crucial to the treatment of conjunction: the states verifying a conjunction are exactly those states that may be obtained by fusing a verifier of one conjunct with a verifier of the other conjunct. A disjunction, by contrast, is verified by the states verifying either disjunct. Dually, a disjunction is falsified by any fusions of falsifiers of the disjuncts, and a conjunction is falsified by the falsifiers of the conjuncts.

The notion of verification (and analogously falsification) incorporates a strong requirement of *relevance*. In order for a state to verify a proposition, it must be wholly relevant to the truth of the proposition. Thus, the state of two plus two being four does not verify the proposition that grass is green or not: even though the state's obtaining necessitates the truth of the proposition, it is irrelevant to it. Moreover, the state of it being sunny and warm does not verify the proposition that it is sunny, since it is not *wholly* relevant to the latter: it contains as an irrelevant part the state of it being warm.

We turn finally to *modes* of verification. The basic observation motivating the modeified truthmaker approach is this. Intuitively, for some propositions, there are several ways, or *modes*, in which a state might verify the proposition (analogous remarks apply with respect to falsification). Disjunctive propositions are perhaps the most compelling example. Thus, consider the proposition that this ball is red or blue. This proposition might be made true, for example, by the state of the ball being scarlet—call this state s—, or by the state of the ball being navy blue—call this state n. Now note that the state s also verifies the proposition that this ball is red, and it seems plausible to say that s verifies the proposition that this ball is red or blue *by* verifying the proposition that this ball is red. Similarly, the state n also verifies the proposition that this ball is blue, and it seems plausible to say that n verifies the proposition that this ball is red or blue *by* verifying the proposition that this ball is blue. In this way, we may distinguish between at least two different modes of verifying the proposition that this ball is red or blue: the mode of verifying it by verifying the one disjunct, that this ball is red, and the mode of verifying it by verifying the other disjunct, that this ball is blue.[10]

The key features of modes of verification, in bullet point fashion, are these:

(1) Modes are subject to a fullness or sufficiency condition. So verifying a proposition P is a mode of verifying a proposition Q only if verifying P is *sufficient* for verifying Q.[11]
(2) Some propositions may be verified by verifying not a single proposition P, but a *multitude* of propositions $P_1, P_2, \ldots$ The most compelling example is that of a conjunction $P \wedge Q$ which may be verified by verifying both P and Q.
(3) Modes can be combined. For instance, if P may be verified by verifying $P_1, P_2, \ldots$ and Q may be verified by verifying $Q_1, Q_2, \ldots$ then plausibly,

[10] In order to further clarify the notion of a mode of verification, one might consider analysing it in terms of ground: prima facie, one might think that for a state s to verify P by verifying $Q, R, \ldots$ is for s's verifying $Q, R, \ldots$ to ground s's verifying P. But this is not the place to further pursue this idea. (Thanks here to Fabrice Correia.)

[11] This bears highlighting since in ordinary language, 'by' tolerates partial ways of doing something: by saying that someone got into the house by breaking a window we do not imply that breaking the window was on its own sufficient for the person to get into the house.

$P \wedge Q$ may be verified in a combination of these two modes, by verifying $P_1, P_2, \ldots, Q_1, Q_2, \ldots$

(4) There may be propositions P for which there are no propositions $P_1, P_2, \ldots$ such that P may be verified by verifying $P_1, P_2, \ldots$ P is then said to be verified in a *direct*, unmediated way by any verifying states.

(5) Talk of verification, and modes of verification, is to be understood *non-factively*: that verifying P is a mode of verifying Q does not imply that P or Q are true. Indeed, in the intended sense, it does not even imply that P or Q are *possibly* true. For example, verifying $P, \neg P$ is a mode of verifying $P \wedge \neg P$, even though it is logically impossible for $P \wedge \neg P$ to be true.[12]

(6) The relation of verification in a mode has certain important structural properties similar to those standardly associated with grounding. For example, like grounding, it is non-monotonic. Thus, it may happen that a proposition can be verified by verifying P, but not by verifying P, Q. Further structural properties that modified verification plausibly exhibits include irreflexivity—no proposition is verified by verifying itself—and transitivity—if P is verified by verifying Q, and Q by verifying R, then P is verified by verifying R.[13]

For ease of expression, we informally refer to the mode of verifying a proposition by verifying $P_1, P_2, \ldots$ as the mode via-$P_1, P_2, \ldots$ Note that to each mode $m =$ via-$P_1, P_2, \ldots$ there corresponds the set $|m|$ of propositions $\{P_1, P_2, \ldots\}$. It is by reference to this set of propositions that relationships of *ground* may ultimately be defined.

First, though, we need to say how the notion of a proposition is understood within the mode-ified truthmaker account. Any non-empty set of modes is a *unilateral* proposition. The presence of a mode m in a unilateral proposition P represents that m is a mode of verifying P. More precisely, if $m =$ via-$P_1, P_2, \ldots$, then whenever $s = s_1 \sqcup s_2 \sqcup \ldots$ with each s_i verifying the corresponding P_i, s verifies P by verifying $P_1, P_2, \ldots$ If m is the mode of being directly verified by state s, then the presence of m in a proposition P represents that P is directly verified by s.

Any pair of unilateral propositions is a *bilateral* proposition. While unilateral propositions are referred to by uppercase letters $P, Q, \ldots$, bilateral propositions are referred to by bold-face versions of these. The first (second) component of a bilateral proposition $\mathbf{P}$ is called its positive (negative) content and denoted $\mathbf{P}^+$ ($\mathbf{P}^-$). That a mode m is a member of $\mathbf{P}^+$ ($\mathbf{P}^-$) represents that m is a mode of verifying (falsifying) $\mathbf{P}$.

Finally, relations of full, partial, weak, and strict ground for unilateral contents are defined as follows:

Γ strictly fully grounds P ($\Gamma < P$)	$:\leftrightarrow \Gamma = \lvert m\rvert$ for some mode $m \in P$
Q strictly partially grounds P ($Q \prec P$)	$:\leftrightarrow Q \in \Gamma$ for some $\Gamma < P$
Γ weakly full grounds P ($\Gamma \leq P$)	$:\leftrightarrow \Gamma = \{P\}$ or $\Gamma < P$ or $\Gamma \setminus \{P\} < P$
Q weakly partially grounds P ($Q \preceq P$)	$:\leftrightarrow Q = P$ or $Q \prec P$

[12] Similarly, when mention is made of states (of affairs) verifying certain propositions, the pertinent notion of a state should be taken to encompass non-actual and even impossible states as well as actual ones.

[13] Under the ground-theoretic analysis of modes of verification suggested in footnote 10, we might be able to derive that modes of verification have these properties from the corresponding ground-theoretic assumptions.

Grounding between bilateral propositions holds just in case it obtains between the positive components. Note that these definitions target a *non-factive* understanding of ground on which falsities may ground each other just as much as truths may. The more familiar factive notion can easily be defined in terms of the non-factive one by adding the condition that grounds and groundee be true.[14] For present purposes, however, there is no need to do so.

Krämer (2018) then shows that under a natural account of conjunction, disjunction, and negation, grounding interacts with these truth-functions in accordance with some very attractive and intuitively compelling principles.[15] To state them succinctly, let $\boldsymbol{\Gamma} \leq \{\mathbf{P}_1, \mathbf{P}_2, \ldots\}$ abbreviate that for some sets $\boldsymbol{\Gamma}_1, \boldsymbol{\Gamma}_2, \ldots$ with $\boldsymbol{\Gamma} = \boldsymbol{\Gamma}_1 \cup \boldsymbol{\Gamma}_2 \cup \ldots$, $\boldsymbol{\Gamma}_1 \leq \mathbf{P}_1$ and $\boldsymbol{\Gamma}_2 \leq \mathbf{P}_2$ and ... Then the central result is the following set of equivalences (theorems 5 and 6 of the appendix in Krämer 2018):

($<\wedge$) $\boldsymbol{\Gamma} < \mathbf{P} \wedge \mathbf{Q}$ iff $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
($<\vee$) $\boldsymbol{\Gamma} < \mathbf{P} \vee \mathbf{Q}$ iff $\boldsymbol{\Gamma} \leq \mathbf{P}$ or $\boldsymbol{\Gamma} \leq \mathbf{Q}$ or $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
($<\neg\neg$) $\boldsymbol{\Gamma} < \neg\neg\mathbf{P}$ iff $\boldsymbol{\Gamma} \leq \mathbf{P}$
($<\neg\wedge$) $\boldsymbol{\Gamma} < \neg(\mathbf{P} \wedge \mathbf{Q})$ iff $\boldsymbol{\Gamma} \leq \neg\mathbf{P}$ or $\boldsymbol{\Gamma} \leq \neg\mathbf{Q}$ or $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$
($<\neg\vee$) $\boldsymbol{\Gamma} < \neg(\mathbf{P} \vee \mathbf{Q})$ iff $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$

I shall sometimes refer to these principles as introduction (elimination) principles for the relevant connectives, when only the right-to-left (left-to-right) direction of the bi-conditional is at issue.

3.2 The general account of ground-theoretic equivalence

The general Mode-ified Truthmaker Account of ground-theoretic equivalence may be stated as follows:[16]

MTA Sentences S and T are ground-theoretically equivalent if and only if S and T are verified by the same states in the same modes.

Given the account of the previous section, this simplifies to the claim that sentences are ground-theoretically equivalent iff the positive components of the bilateral propositions expressed by the sentences are identical. Writing $[S]$ for the proposition expressed by S:

(A) Sentences S and T are ground-theoretically equivalent iff $[S]^+ = [T]^+$

I shall make the widely accepted assumption that no proposition is a strict full ground of itself. Given that assumption, the condition of identity of positive content is equivalent

14 Cf. Krämer (2018, pp. 795f, 798). On the distinction between factive and non-factive grounding, see also Fine (2012a, p. 48ff) and Correia (2014, p. 36).—To define a notion of truth in the present framework, some modes may be designated as *actual*, i.e. such that some state verifies in the relevant mode, and a proposition may then be counted as true just in case it is verified in some actual mode.

15 These principles, in a slightly different form, were first proposed by Fine (2012a, sections 7 and 8). Versions of in particular the left-to-right directions of the bi-conditionals are endorsed in numerous works on the logic of ground, including Batchelor (2010), Rosen (2010), Schnieder (2011), Correia (2014, 2017).

16 Here and in what follows, talk of sentences being verified (in a certain mode) is to be read as shorthand for talk of the sentence's positive content being verified (in the relevant mode).

to what we may call sameness of overall ground-theoretic profile.[17] More precisely, let us write $\approx$ for the relation defined as follows:

$\mathbf{P} \approx \mathbf{Q}$ iff

(i) for all $\boldsymbol{\Gamma}$: $\boldsymbol{\Gamma} < \mathbf{P}$ iff $\boldsymbol{\Gamma} < \mathbf{Q}$, and
(ii) for all $\boldsymbol{\Delta}$ and $\mathbf{R}$: $\boldsymbol{\Delta}, \mathbf{P} < \mathbf{R}$ iff $\boldsymbol{\Delta}, \mathbf{Q} < \mathbf{R}$.

Thus, $\mathbf{P} \approx \mathbf{Q}$ iff $\mathbf{P}$ and $\mathbf{Q}$ participate in exactly the same relationships of non-factive strict full ground.[18] Since all the other relationships of ground are defined in terms of non-factive strict full ground, $\mathbf{P} \approx \mathbf{Q}$ implies overall sameness of ground-theoretic profile. Then assuming the irreflexivity of ground,

(B) $\mathbf{P}^+ = \mathbf{Q}^+$ iff $\mathbf{P} \approx \mathbf{Q}$

I shall henceforth use 'ground-theoretic equivalence' both for the relationship between sentences characterized in the beginning of the paper and the relation $\approx$ between propositions, and I shall often tacitly rely on the equivalences (A) and (B). For convenience, I mostly carry out the subsequent investigation purely on the level of content, avoiding the distraction of a detour through language whenever possible.

In order to determine whether two given propositions are ground-theoretically equivalent, we may need to decide whether a certain mode m = via-$P_1, P_2, \ldots$ is *the same mode* as a mode n = via-$Q_1, Q_2, \ldots$ For instance, suppose we wish to know whether $\mathbf{P} \wedge \mathbf{Q} \approx \mathbf{Q} \wedge \mathbf{P}$. Under a natural account of conjunction, this will be so just in case the mode via-$\mathbf{P}^+, \mathbf{Q}^+$ is identical to the mode via-$\mathbf{Q}^+, \mathbf{P}^+$.

It is clear, first of all, that the modes via-$P_1, P_2, \ldots$ and via-$Q_1, Q_2, \ldots$ are identical if the *sequences* $\langle P_1, P_2, \ldots\rangle$ and $\langle Q_1, Q_2, \ldots\rangle$ are identical. We may assume, furthermore, that the modes via-$P_1, P_2, \ldots$ and via-$Q_1, Q_2, \ldots$ are identical *only if* the *sets* $\{P_1, P_2, \ldots\}$ and $\{Q_1, Q_2, \ldots\}$ are identical.[19] In light of this, there are two natural further principles concerning the individuation of modes to consider. The first is that modes are *order-insensitive*: the modes via-$P_1, P_2, \ldots$ and via-$Q_1, Q_2, \ldots$ are identical whenever the sequences $\langle P_1, P_2, \ldots\rangle$ and $\langle Q_1, Q_2, \ldots\rangle$ differ at most by the order of the propositions they contain, i.e. when the *multi-sets* $\lfloor P_1, P_2, \ldots\rfloor$ and $\lfloor Q_1, Q_2, \ldots\rfloor$ are the same.[20] The second is that modes are order- and *repetition-*

[17] This is shown in Krämer (2018, p. 805f).

[18] An anonymous referee has raised the question whether condition (ii) is needed in the definition of $\approx$. Condition (ii) is redundant just in case whenever two propositions have the same grounds, they ground the same things. This is not in general so under the mode-theoretic approach, it depends on just how modes are individuated—a matter we shall take up momentarily. We shall see then that under the most coarse-grained of the available views, condition (ii) is indeed redundant, but not on the others. As it happens, however, condition (i) *is* in general redundant within the mode-theoretic framework when irreflexivity of $<$ is assumed. For suppose $\mathbf{P}$ and $\mathbf{Q}$ ground the same things. By ($< \neg\neg$), $\mathbf{Q} < \neg\neg\mathbf{Q}$, and so $\mathbf{P} < \neg\neg\mathbf{Q}$, which entails again by ($< \neg\neg$) that $\mathbf{P} \leq \mathbf{Q}$. By parallel reasoning, $\mathbf{Q} \leq \mathbf{P}$. but given irreflexivity, it is easy to verify that mutual weak full ground implies sameness of positive content and thereby also sameness of grounds. I have nevertheless included condition (i) since this is how the relation is defined in my earlier paper, and since this makes the conceptual connection to ground-theoretic equivalence as introduced in the beginning of the paper much more explicit.

[19] Indeed, we have assumed as much when we claimed that to each mode m there corresponds a unique set $|m|$ of propositions. The assumption is crucial for the derivation of the main results in the logic of ground of Krämer (2018); see esp. the appendix of that paper.

[20] A multi-set is like a set except in that it may contain the same item more than once. I write $\lfloor P, Q \ldots\rfloor$ for the multi-set including exactly $P, Q, \ldots$, each exactly as many times as it is listed.

insensitive: the modes via-$P_1, P_2, \ldots$ and via-$Q_1, Q_2, \ldots$ are identical whenever the sequences $\langle P_1, P_2, \ldots\rangle$ and $\langle Q_1, Q_2, \ldots\rangle$ differ at most by the order or the number of occurrences of the propositions they contain, i.e. when the *sets* $\{P_1, P_2, \ldots\}$ and $\{Q_1, Q_2, \ldots\}$ are the same.

I shall defend the view that our conception of modes should be order-insensitive but repetition-sensitive. Section 3.2.1 makes the case for order-insensitivity, and Sect. 3.2.2 argues for repetition-sensitivity. Section 3.2.3 develops and responds to a possible objection.

3.2.1 Against order-sensitivity

For the purposes of this discussion, I shall assume that if an order-sensitive conception of modes is assumed, then $\mathbf{P} \wedge \mathbf{Q}$ will not in general be verified via $\mathbf{Q}^+, \mathbf{P}^+$, although it will be verified via $\mathbf{P}^+, \mathbf{Q}^+$. This seems reasonable. First, the assumption holds under the account of conjunction in Krämer (2018). Second, although it would be possible to adjust that account of conjunction so as to avoid this result, it is hard to see what point the distinction between the modes via-$\mathbf{P}^+, \mathbf{Q}^+$ and via-$\mathbf{Q}^+, \mathbf{P}^+$ could have if not to distinguish between the conjunctions $\mathbf{P} \wedge \mathbf{Q}$ and $\mathbf{Q} \wedge \mathbf{P}$.

The first point in favour of order-insensitivity is that an order-sensitive conception of modes has counter-intuitive consequences, such as that $\mathbf{P} \wedge \mathbf{Q}$ grounds $\neg\neg(\mathbf{P} \wedge \mathbf{Q})$ whereas $\mathbf{Q} \wedge \mathbf{P}$ does not.[21] In addition, we can defend order-insensitivity by appeal to the way that modes of verification have informally been introduced. Suppose that Q and R are distinct unilateral propositions, and that via-Q, R and via-R, Q are different modes. We are then making a distinction between the mode of verifying a proposition by verifying Q, R and the mode of verifying the proposition by verifying R, Q. This is plausible only to the extent that some difference can be made out between a state's verifying Q, R and the same state's verifying R, Q.

The relevant notion of verifying a sequence of propositions is explicitly defined in terms of the notion of verifying a single proposition in the following way: 'the propositions $P_1, P_2, \ldots$ are verified by a [state] iff the [state] is the fusion of some [states] $s_1, s_2, \ldots$ verifying $P_1, P_2, \ldots$ respectively' (Krämer 2018, p. 792).[22] Under this definition, then, to say that a state s verifies a given proposition P via-Q, R is in effect to say that s verifies P by being the fusion of some states t, u verifying Q, R, respectively. And to say, in constrast, that state s verifies P via R, Q is to say that s verifies P by being the fusion of some states t, u verifying R, Q, respectively. Given that the operation of fusion is an operation on sets of states and thus insensitive to order, it would seem wholly implausible to take these two statements to describe two distinct ways in which s verifies P. Moreover, it is hard to see how any plausible alternative definition of the verification of a sequence of propositions could lead to a different assessment.

[21] Suppose that via-$\mathbf{P}^+, \mathbf{Q}^+$ and via-$\mathbf{Q}^+, \mathbf{P}^+$ are distinct modes, so $\mathbf{P} \wedge \mathbf{Q}$ and $\mathbf{Q} \wedge \mathbf{P}$ are ground-theoretically inequivalent. Then $\mathbf{P}\wedge\mathbf{Q} < \neg\neg(\mathbf{P}\wedge\mathbf{Q})$. But suppose for reductio that also $\mathbf{Q}\wedge\mathbf{P} < \neg\neg(\mathbf{P}\wedge\mathbf{Q})$. By the elimination principle for double negation, it follows that $\mathbf{Q} \wedge \mathbf{P} \leq \mathbf{P} \wedge \mathbf{Q}$. Since $\mathbf{P} \wedge \mathbf{Q} \not\approx \mathbf{Q} \wedge \mathbf{P}$, it follows that $\mathbf{Q} \wedge \mathbf{P}$ is *strict* full ground of $\mathbf{P} \wedge \mathbf{Q}$, which is absurd. So $\mathbf{Q} \wedge \mathbf{P} \not< \neg\neg(\mathbf{P} \wedge \mathbf{Q})$.

[22] In the quoted passage, I spoke of facts—i.e. states that actually obtain—rather than states, but as the view is later developed, the general definition must appeal to states (cf. Krämer 2018, p. 794f).

3.2.2 For repetition-sensitivity

Similarly as in the previous case of order-sensitivity, for the purposes of our discussion of repetition-sensitivity, we shall make an assumption about how the matter manifests itself in the theory of ground-theoretic equivalence. Specifically, we shall assume that if a repetition-sensitive conception of modes is adopted, then the propositions $\mathbf{P} \wedge \mathbf{P}$, $\mathbf{P} \vee \mathbf{P}$, and $\neg\neg\mathbf{P}$ will be pairwise inequivalent, whereas they are pairwise equivalent under a repetition-insensitive conception. Again, this seems reasonable. First, the assumption holds under the account of the truth-functional operations in Krämer (2018): $\mathbf{P} \wedge \mathbf{P}$ is then verified via $\mathbf{P}^+, \mathbf{P}^+$ but not via $\mathbf{P}^+$, the opposite is true of $\neg\neg\mathbf{P}$, while $\mathbf{P} \vee \mathbf{P}$ is verified both via $\mathbf{P}^+$ and via $\mathbf{P}^+, \mathbf{P}^+$. Second, although it might be possible to adjust the definitions so as to avoid this result, it is hard to see what point the distinction between modes differing only with respect to repetition could retain under such an alternative set of definitions.[23] Third, as we shall see, the ability to render the above pairs inequivalent is precisely what makes the repetition-sensitive conception seem preferable.

We note first that a repetition-sensitive conception of modes seems to yield a more intuitive account than a repetition-insensitive one. For consider the truth that snow is black or (snow is white and snow is white). It obtains because (snow is white and snow is white). From the ground-theoretic equivalence between the proposition that (snow is white and snow is white) and the proposition that (it is not the case that snow is not white), we could infer that the truth that snow is black or (snow is white and snow is white) holds because it is not the case that snow is not white. This seems a rather odd thing to say.[24] So the relevant intuitions about the plausibility of 'because'-statements seem to speak against imposing the requirement of repetition-insensitivity.

In addition, we can defend repetition-sensitivity based on the informal conception of a mode of verification. A conception of modes that is repetition-sensitive embodies a distinction between pairs of modes like via-P and via-P, P. This distinction is justified just in case there is a relevant *difference* between a state's verifying P and a state's verifying P, P. I shall now argue that there is such a difference.

Similarly as before, we proceed by applying the official definition of the notion of a state's verifying a list of proposition. We then obtain that for a state s to verify P, P is for s to be the fusion of some states t, u such that t verifies P and u verifies P. At first glance, this condition seems importantly different from the simple condition of s being a verifier of P: being the fusion of items with a certain property is not the same as being an item with that property.

---

23 It might perhaps be argued that $\mathbf{P} \vee \mathbf{P}$ should not be taken to be verified via $\mathbf{P}^+, \mathbf{P}^+$, but only via $\mathbf{P}^+$. It would then turn out ground-theoretically equivalent to $\neg\neg\mathbf{P}$ even under a repetition-sensitive conception of modes. A version of the central point above would then still remain, however. For the purposes of our arguments below, we only need the assumption that $\mathbf{P} \wedge \mathbf{P}$ is ground-theoretically distinguished from both $\mathbf{P} \vee \mathbf{P}$ and $\neg\neg\mathbf{P}$ just in case a repetition-sensitive conception is adopted. And it would be hard to see what point the distinction between modes via-P, P and via-P could retain if not to allow for this distinction.

24 The intuition may be even more forceful when one discusses the case schematically—even if intuitions about these sorts of statements are perhaps not in any strong sense *pre-theoretical* intuitions. If it is true that A, then it seems plausible that B or (A and A) because A and A, but it seems considerably stranger to suggest that B or (A and A) because it is not the case that not-A.

One might try to object that the case of the property of verifying a proposition P is special in this regard. In particular, one can argue that the conditions in question are equivalent. Clearly, if s is a verifier of P, then since $s = s \sqcup s$, s is also the fusion of a verifier of P with a verifier of P. What about the converse? In Fine's truthmaker semantics for ground, it is assumed that whenever a state s is a fusion of states t, u both verifying P, s itself also verifies P. In that context, this assumption corresponds to the ground-theoretic principle of *amalgamation*, which says that if $\Gamma < P$ and $\Delta < P$ then $\Gamma \cup \Delta < P$. An analogous condition on modes, playing the same role with respect to the logic of ground, may also be imposed on the mode-ified truthmaker account (cf. Krämer 2018, p. 804f). So by assuming the principle of amalgamation, one might indeed argue that any fusion of verifiers of P must itself be a verifier of P.[25]

But even granting that the conditions are satisfied by the same states, this is consistent with the claim that there are important differences between them. In particular, as I shall now argue, the conditions differ with respect to what propositions a state may verify *by* satisfying them. For let s be a state verifying a proposition P, and consider the conjunction $P \wedge P$. Under the truthmaker approach, mode-ified or otherwise, the set of states verifying a conjunction $P \wedge Q$ is defined as the set of states that can be obtained by taking the fusion of a verifier of P and a verifier of Q. By that definition, what it takes for a state to verify $P \wedge P$ is for it to be the fusion of a verifier of P and a verifier of P. The state s satisfies this condition, and therefore verifies $P \wedge P$. Moreover, so the basic insight underlying the mode-ified approach, s verifies $P \wedge P$ *by* being the fusion of a verifier of P and a verifier of P. By the definition of verification of a list of propositions, we may infer that s verifies $P \wedge P$ *by* verifying P, P.

Consider now the double negation $\neg\neg P$.[26] The set of states verifying $\neg\neg P$ is defined as the set of states verifying P. Since s verifies P, it also verifies $\neg\neg P$, and under the account defended in Krämer (2018), it does so *by* verifying P. However, if no distinction is made between verifying P and verifying P, P, we may infer from this that s verifies $\neg\neg P$ by verifying P, P, i.e. by being the fusion of a verifier of P and a verifier of P—just like in the case of $P \wedge P$. But this seems implausible. Although s does satisfy this condition, it is not *in virtue of* satisfying *this* condition that s qualifies as a verifier of $\neg\neg P$. If we look at what makes it the case that a state verifies $\neg\neg P$, fusions simply don't come into it. So we have a good reason to distinguish between the modes via-P and via-P, P: $\neg\neg P$ should include the former but not the latter.

3.2.3 Objection: individuation from below

Given the usual conception of grounding as relating a plurality or set of propositions to a single proposition, adopting a repetition-sensitive conception of modes creates a

---

[25] It should be mentioned, though, that the principle of amalgamation is among the more contentious principles in the logic of ground advocated in Fine (2012b), and that Fine himself elsewhere professes some uneasiness about it; cf. his (2012a, 59n16).

[26] Strictly speaking, under the present approach, negation is only defined on bilateral contents. However, if P is the positive content of $\mathbf{P}$, the positive content of $\neg\neg\mathbf{P}$ depends only on P—it is the result of applying the so-called *raising* operation to P. As a harmless simplification, I here speak of that unilateral content as the double negation of P.

kind of mismatch in the individuation of modes on the one hand and grounds on the other. The strongest argument against the repetition-sensitive conception of modes, it seems to me, holds that this mismatch is undesirable. The aim of this section is to sketch such an argument, and to say how repetition-sensitivity may be defended against the objection.

Under a repetition-insensitive view of modes, the indirect modes in which a given proposition may be verified correspond one-to-one to non-factive strict full grounds of the proposition. As a result, given any propositions **P** and **Q** that are not fundamental—i.e. they have at least one non-factive strict full ground—**P** is ground-theoretically equivalent to **Q** just in case **P** and **Q** have exactly the same grounds. So for non-fundamental propositions, sameness of grounds implies sameness of groundees. In this sense, the items in the ground-theoretic hierarchy then obey a principle of *Individuation from Below*: the identity of any given item in the hierarchy is fully determined by the part of the hierarchy below it.

I can think of two reasons to be attracted to the principle of Individuation from Below, one logical and one metaphysical. The logical reason is that it opens up the possibility of an alternative, and potentially particularly neat and elegant axiomatization of the logic of ground-theoretic equivalence, given a suitable background logic of ground. Essentially, we could utilize familiar techniques from standard quantificational logic to allow the inference to $A \leq B$ whenever $\Gamma < B$ has been derived from $\Gamma < A$ for *arbitrary* Γ. We would then only need to add a rule allowing us to infer $A \approx B$ given both $A \leq B$ and $B \leq A$.[27]

Although this is an intriguing idea, more work would need to be done before it could justify an assessment of repetition-insensitivity as generating the overall more attractive logic of ground and ground-theoretic equivalence. And even if such an assessment could be justified, there remains the question of how this fact should be weighed up against the intuitive benefits offered by the repetition-sensitive view. So far, then, the logical considerations in favour of Individuation from Below do not constitute a strong case against the repetition-sensitive conception.

The metaphysical reason to find the principle of Individuation from Below attractive is that it yields a purer, more self-contained picture of the ground-theoretic hierarchy, similar in respect of individuation to the set-theoretic hierarchy. In the latter case, the identity of each item in the hierarchy is fully accounted for in terms of its set-theoretic relationship to items that are below it. In the same way, on the repetition-insensitive view, the identity of each item in the ground-theoretic hierarchy is fully accounted for in terms of its ground-theoretic relationship to items below it. Under the repetition-sensitive conception, in contrast, this is not so. Although the identity of an element in the grounding hierarchy may still be determined without reference to items *above* it, we now have to appeal to something *extraneous* to the ground-theoretic hierarchy, namely *multi-sets* of elements below the given element. In that sense, we are introducing

[27] For a simple illustration of the idea, assume for arbitrary Γ that $\Gamma < A \wedge B$. Using the elimination rule for conjunction, we infer $\Gamma \leq \{A, B\}$. But since $\{A, B\} = \{B, A\}$, this is just $\Gamma \leq \{B, A\}$. Using an obvious generalization of the usual introduction rule for conjunction, we infer $\Gamma < B \wedge A$. Since Γ was arbitrary, by the rule sketched above, we may then infer $A \wedge B \leq B \wedge A$. Parallel reasoning establishes $B \wedge A \leq A \wedge B$, whence we may infer $A \wedge B \approx B \wedge A$.

distinctions into the ground-theoretic hierarchy with no purely ground-theoretic basis, and it might be thought that this should be avoided.

Three things may be said in response. The first is that the argument relies on a contentious 'purity' assumption concerning legitimate ways to individuate ground-theoretic content, for which a sustained defence has yet to be given. The second is to challenge the claim that the relevant distinctions have no purely ground-theoretic basis. For it seems that one natural way to understand the argument of the previous section is precisely as identifying such a basis. We argued there, recall, that $P \wedge P$ may be verified by verifying $P,\ P$, whereas $\neg\neg P$ may not be so verified. It is not implausible to regard these claims as equivalent to corresponding grounding claims: that s verifies $P,\ P$ grounds that s verifies $P \wedge P$, but does not ground that s verifies $\neg\neg P$. A third response might be to accept Individuation from Below, but to insist that rather than adopt a repetition-insensitive conception of modes, we should adopt a repetition-sensitive conception of grounds. I myself have some sympathy for this suggestion, but I lack the space to further pursue it here.[28]

3.3 The formal account of ground-theoretic equivalence

This section answers the formal question of ground-theoretic equivalence. Since the case made in the previous section for repetition-sensitivity is not fully decisive, and since we found the repetition-insensitive conception to carry some theoretical interest, I shall do this for both the repetition-sensitive and the repetition-insensitive conception of modes. The formal question I posed, recall, is under what conditions two sentences are ground-theoretically equivalent in virtue of their propositional logical form. To answer this question, we shall make use of a formal language $\mathscr{L}_{\approx}$ in which statements of ground-theoretic equivalence can be formulated. In our choice of a language, we shall follow the example of Correia (2016) and use a standard propositional language with connectives $\wedge$, $\vee$, and $\neg$, augmented by all expressions of the form $A \approx B$ where $A,\ B$ are sentences of the propositional language (and thus do not already include occurrences of $\approx$).

The following system of axioms and rules for $\mathscr{L}_{\approx}$—I shall call it $\mathfrak{I}$ or *the intermediate system*—is sound and complete relative to the order-insensitive but repetition-sensitive conception of modes.[29, 30]

(Commutativity $\vee$)	$A \vee B \approx B \vee A$
(Commutativity $\wedge$)	$A \wedge B \approx B \wedge A$
(De Morgan 1)	$\neg(A \vee B) \approx \neg A \wedge \neg B$
(De Morgan 2)	$\neg(A \wedge B) \approx \neg A \vee \neg B$
(Reflexivity)	$A \approx A$

[28] Note that Poggiolesi (2016b, 2018) also works with a conception of grounds as multi-sets.

[29] The principles are to be read as follows. The first five principles are axiom schemata, so any instance obtained by systematically replacing the sentence letters $A,\ B,\ C$ by propositional formulae is an axiom. The last five principles are (schematic) rules, starting with a comma-separated list of the premises, separated by a forward slash from the conclusion.

[30] The proof of this result and the next is given in appendix A. The motivation for the label 'intermediate' is that relative to the range of possible standards of individuation in the mode-theoretic framework, $\mathfrak{I}$ occupies an intermediate position.

(Symmetry) $A \approx B \ / \ B \approx A$
(Transitivity) $A \approx B, \ B \approx C \ / \ A \approx C$
(Preservation $\vee$) $A \approx B \ / \ A \vee C \approx B \vee C$
(Preservation $\wedge$) $A \approx B \ / \ A \wedge C \approx B \wedge C$
(Preservation $\neg\neg$) $A \approx B \ / \ \neg\neg A \approx \neg\neg B$

Under a repetition-insensitive conception of modes, corresponding self-disjunctions, self-conjunctions, and double negations are ground-theoretically equivalent. As a result, this conception leads to the validation of the following additional axioms:

(Collapse $\wedge/\vee$) $A \wedge A \approx A \vee A$
(Collapse $\vee/\neg\neg$) $A \vee A \approx \neg\neg A$

More strikingly, under a repetition-insensitive conception, the condition of a proposition **P** being a weak full ground of a proposition **Q** is equivalent to the condition that $\mathbf{P} \vee \mathbf{Q} \approx \mathbf{Q} \vee \mathbf{Q}$ (cf. Krämer 2018, p. 807). Therefore, letting $A \leq B$ abbreviate $A \vee B \approx B \vee B$, we can then state within $\mathscr{L}_{\approx}$ a form of introduction rules for $\leq$:

(Introduction $\leq\wedge$) $A \leq B, \ A \leq C \ / \ A \leq B \wedge C$
(Introduction $\leq\vee$) $A \leq B \ / \ A \leq B \vee C$

Adding these four principles to $\mathfrak{I}$ results in a system I shall call $\mathfrak{E}$, or *the extensional system*.[31] It is sound and complete relative to the order- and repetition-insensitive conception of modes.

4 Comparison to other approaches

So far, there is effectively only one other approach that offers both a general account of ground-theoretic equivalence and an answer to the formal question: the (unmodified) truthmaker account, which is developed in slightly different versions in Fine (2016) and Correia (2016). Beyond that, three partial accounts have been proposed. Correia (2017) develops, among other things, a formal account of ground-theoretic equivalence, but no informative general account that could serve to motivate the formal account. Correia (2018), in contrast, formulates an alternative account of ground, and ground-theoretic equivalence, but only partially specifies the resulting propositional logic of ground-theoretic equivalence. Poggiolesi (2016b, 2018) develops a further and interestingly different account of logical grounding and ground-theoretic equivalence. The purpose of this section is to identify both similarities and dissimilarities between these other approaches and the one defended in this paper, and to point out where I think the latter holds an advantage.

4.1 Fine's and Correia's truthmaker accounts

The *truthmaker* account may be stated as follows:

[31] The motivation for the label is that given the principle of Individuation from Below, non-fundamental ground-theoretic contents are extensional with respect to their grounds in the sense that they are ground-theoretically equivalent if they have the same grounds.

(TM) Sentence S is ground-theoretically equivalent to sentence T iff S and T are exactly *verified* by exactly the same states.

This account of ground-theoretic equivalence is implicit in the truthmaker account of ground presented in Fine (2017b, §6). Here ground is defined as relating contents identified with sets of exact truthmakers. Sentences accordingly will turn out ground-theoretically equivalent just in case they have the same exact truthmakers.[32]

A version of this view is given a more explicit and detailed exposition in Correia (2016). Here, Correia discusses a relation he calls *factual equivalence*, which he defines as obtaining between two sentences just in case they describe the same situations. The notion of a sentence's describing a situation is proposed by Correia as a more concrete specification of Fine's notion of a sentence being verified by a state.[33] In the final section of the paper, he then proposes a substitutivity principle for the logic of ground that permits the substitution of sentences in an argument place of the grounding operator in case the sentences stand in the relation of factual equivalence, thereby assigning to factual equivalence the role of ground-theoretic equivalence as I have used the term.[34]

Fine (2016) and Correia (2016) determine the logics of ground-theoretic equivalence under a couple of closely related semantic implementations of (TM). Fine obtains a logic that can very naturally be axiomatized as follows:

(Collapse $\neg\neg$)	$A \approx \neg\neg A$
(Collapse $\wedge$)	$A \approx A \wedge A$
(Commutativity $\wedge$)	$A \wedge B \approx B \wedge A$
(Associativity $\wedge$)	$A \wedge (B \wedge C) \approx (A \wedge B) \wedge C$
(Collapse $\vee$)	$A \approx A \vee A$
(Commutativity $\vee$)	$A \vee B \approx B \vee A$
(Associativity $\vee$)	$A \vee (B \vee C) \approx (A \vee B) \vee C$
(De Morgan 1)	$\neg(A \vee B) \approx \neg A \wedge \neg B$
(De Morgan 2)	$\neg(A \wedge B) \approx \neg A \vee \neg B$
(Distributivity 1)	$A \wedge (B \vee C) \approx (A \wedge B) \vee (A \wedge C)$
(Distributivity 2)	$A \vee (B \wedge C) \approx (A \vee B) \wedge (A \vee C)$
(Symmetry)	$A \approx B \ / \ B \approx A$
(Transitivity)	$A \approx B,\ B \approx C \ / \ A \approx C$

[32] Versions of the same view are also implicit in the other contributions of Fine's in which he formulates a truthmaker semantics for ground, specifically (Fine 2012a, b). The discussion in Fine (2017b) is more pertinent however. In contrast to the Fine (2012b), it deals not just with the purely structural features of ground but also with its interaction with the truth-functional connectives. In contrast to the discussion in Fine (2012a), it explicitly works with a conception of states encompassing merely possible and indeed impossible states rather than just actually obtaining facts, resulting in a much more plausible view.

[33] Cf. Correia (2016, p. 107). More accurately, Correia offers the notion of a sentence's *fittingly* describing a situation as a specification of Fine's notion of a sentence being *exactly* verified by a state. Similar to Fine's looser notions of verification, Correia also discusses looser versions of describing. Again, these are of no importance for our present purposes.

[34] Strictly speaking, this is not quite right, since Correia is only committed to the claim that factual equivalence implies ground-theoretic equivalence, not to the converse. To simplify presentation, I shall conduct my discussion under the assumption that the converse claim is accepted, too. However, nothing substantive hangs on this, as my critical comments target exclusively the claim that factual equivalence implies ground-theoretic equivalence.

(Preservation $\vee$) $A \approx B \;/\; A \vee C \approx B \vee C$
(Preservation $\wedge$) $A \approx B \;/\; A \wedge C \approx B \wedge C$

This logic—let us call it $\mathfrak{W}_1$—coincides with R. B. Angell's logic of analytic equivalence (Angell 1977), which had also been proposed as the logic of ground-theoretic equivalence in Correia (2010). The axiom (Distributivity 2) has since been independently criticized on similar grounds by Krämer and Roski (2015) and (Correia 2016, p. 119). In the latter paper, Correia therefore drops (Distributivity 2), and proves the resulting weaker logic—call it $\mathfrak{W}_2$—sound and complete with respect to a slightly modified form of the truthmaker semantics used in Fine (2016).

Although in many ways attractive, the truthmaker account of ground-theoretic equivalence suffers from an important limitation: it is unable to validate the above described introduction and elimination principles for ground—indeed, it is precisely this limitation that provides the motivation for the development in Krämer (2018) of the mode-ified truthmaker account. The most obvious instance of the problem concerns the principle that P grounds $\neg\neg P$: in conjunction with the axiom (Collapse $\neg\neg$), it would enable us to derive $A < A$ for any A, rendering ground not only not irreflexive, but reflexive, which is an absurd result. Similar comments apply to the other two collapse rules (Collapse $\wedge$) and (Collapse $\vee$).

Beyond the three collapse principles, there are four more axioms of $\mathfrak{W}_1$ that are invalid even under the most coarse-grained mode-theoretic account $\mathfrak{C}$, namely the two associativity principles and the two distributivity principles. It is worth considering briefly what may be said in defence of their rejection, since especially the associativity principles may appear quite plausible at first glance. The first and most important point is that we may reject these principles on essentially the same grounds that motivate the rejection of the collapse principles: they are incompatible with the introduction and elimination principles.[35] I shall illustrate the problem for (Associativity $\vee$), the other cases are similar. By the introduction principle for $\vee$, we have $B \vee C < A \vee (B \vee C)$. By (Associativity $\vee$), we could then infer that $B \vee C < (A \vee B) \vee C$. By the elimination principle for $\vee$, it follows that either (a) $B \vee C \leq (A \vee B)$ or (b) $B \vee C \leq C$. But from (b) we could infer, using the introduction principle for $\vee$ and the suitable transitivity principle, that $C < C$, again rendering ground reflexive. From (a), however, we could infer by the same principles that $C < (A \vee B)$, and thus that any disjunction is grounded by everything, which is absurd.

Fine and Correia are of course aware that the proposed account of ground-theoretic equivalence is not compatible with these introduction and elimination principles. Their response is not to simply reject the rules and to insist on the correctness of the truthmaker account, but instead to invoke a distinction between two alternative and equally legitimate conceptions of ground. On a conception of ground as *worldly*, they suggest, the truthmaker account of ground and equivalence is adequate, and the introduction and elimination rules invalid. On a conception of ground as *representational*, in con-

[35] Of course, these principles are not beyond reasonable doubt themselves, and it would be of interest to explore both how they might be modified to allow for Associativity to hold, and whether the mode-theoretic account might be modified accordingly. However, since the introduction and elimination principles are fairly widely accepted and have considerable initial appeal, their incompatibility with Associativity gives us at least some reason to reject Associativity.

trast, the opposite is the case: the introduction and elimination rules hold, and the truthmaker account constitutes an insufficiently fine-grained view of ground-theoretic equivalence.[36]

I am happy to grant that there is a legitimate conception of ground, and ground-theoretic equivalence, for which the truthmaker account is adequate. I also maintain, naturally, that there is another legitimate conception of the notions for which the mode-theoretic account is adequate. I shall return to the question of its classification with respect to worldliness or representationality in the final subsection of this section.

4.2 Correia's representational account

Correia (2017) presents the so far only logic of ground that validates the introduction and elimination principles we stated above, and for which a soundness and completeness results exists. The logic of ground incorporates a logic of what Correia calls *propositional equivalence*, which is the relation that, within his system, plays the role of ground-theoretic equivalence: it obtains just in case the relevant propositions are exactly alike in all respects to which the notion of ground is sensitive.

Transposed to our setting, Correia's account may be described as comprising the following rules, constituting the *representational*[37] system $\mathfrak{R}$:[38]

(Commutativity $\vee$)	$A \vee B \approx B \vee A$
(Commutativity $\wedge$)	$A \wedge B \approx B \wedge A$
(Reflexivity)	$A \approx A$
(Symmetry)	$A \approx B \ / \ B \approx A$
(Transitivity)	$A \approx B,\ B \approx C \ / \ A \approx C$
(Preservation $\vee$)	$A \approx B \ / \ A \vee C \approx B \vee C$
(Preservation $\wedge$)	$A \approx B \ / \ A \wedge C \approx B \wedge C$
(Preservation $\neg$)	$A \approx B \ / \ \neg A \approx \neg B$

This system differs from $\mathfrak{I}$ by dropping the DeMorgan rules and replacing the rule of preservation under double negation by the stronger rule of preservation under single negation. Since $\mathfrak{R}$ lacks some rules of $\mathfrak{I}$ but also contains a rule that is not valid in $\mathfrak{I}$, it is not as straightforward to compare the systems as in the previous cases. Still, it can be shown that $\mathfrak{R}$ yields a strictly more fine-grained conception of ground-theoretic equivalence in the sense that the set of equivalences it proves is a proper subset of the equivalences proved by $\mathfrak{I}$ (the proof is in appendix B).

[36] The distinction between a worldly and a conceptual or representational conception of ground was first introduced by Correia (2010, p. 256f). In that paper, Correia argued against the representational view, but he has since come to view it as a legitimate alternative conception (cf. e.g. Correia 2018, p. 18n16). For Fine's view of the matter, see esp. (Fine 2017b, p. 685f).

[37] The label is motivated by the fact that Correia explicitly advocates the logic as adequate relative to a conception of ground as a representational rather than worldly relation.

[38] The background language in Correia (2017) is quite different from $\mathscr{L}_{\approx}$, so the rules he writes down look a bit different. (For instance, his basic expression for ground-theoretic equivalence connects a set of sentences Δ with a single sentence φ, and is supposed to say that Δ is non-empty, and every sentence in Δ is equivalent to φ.) He also adds various elimination rules for ground-theoretic equivalence which are redundant for the purpose of deriving equivalences true in virtue of propositional logical form, though not for the broader purposes pursued in Correia (2017).

As indicated above, Correia does not offer an informative general account of ground-theoretic equivalence. It is of course correct to say, on his approach, that sentences S and T are ground-theoretically equivalent iff they express the same proposition, so in this way we can give general necessary and sufficient conditions for ground-theoretic equivalence. But Correia is explicit that the relevant notion of a proposition is to be understood as introduced in a purely functional way: ‘they are the items that play the role of the relata of the relation of strict grounding’ (Correia 2017, p. 515). This functional characterization cannot serve to adequately motivate the specific principles Correia lays down with respect to ground-theoretic equivalence, since it is consistent with propositions being mode-theoretic contents, and thus individuated in a more coarse-grained way.[39]

Can Correia’s formal account be supplemented by a general account appropriately motivating the former? A natural idea is to offer an account along the following lines. Two sentences are ground-theoretically equivalent iff they express the same proposition. Propositions have internal, quasi-syntactic structure. For example, any conjunction $P \wedge Q$ has a unique decomposition into two immediate constituents: the concept of conjunction and the doubleton set of P and Q. Similarly for disjunction and negation, and perhaps other proposition-forming concepts.

Although such an account could certainly be developed, and appears by no means absurd, I think that there is at least one important way in which it would inevitably be less satisfying than either the mode-theoretic or indeed the truthmaker-theoretic account. On both these approaches, we can see *why* ground should be sensitive to the features that we are taking it to be sensitive to. For the same kinds of semantic features of a sentence are used to say *both* when a statement of ground is true and when sentences are ground-theoretically equivalent. Thus, on the truthmaker-theoretic account, ground is sensitive to what states a given proposition is verified by, because to be (weakly) grounded by another proposition is just to be verified by all the states verifying that proposition. And on the mode-theoretic account, ground is sensitive to the modes in which a proposition may be verified because for Γ to ground P just is for it to be the case that verifying Γ is a mode of verifying P. Since one cannot in this way give an attractive general account of ground in terms of the quasi-syntactic conceptual structure of propositions, one cannot in this way explain why ground should be sensitive to differences pertaining to that sort of structure.

4.3 Correia’s relative-fundamentality-based account

In his recent paper on the logic of relative fundamentality (2018), Correia discusses the possibility of giving a reductive account of ground in terms of an independent notion of relative fundamentality and the notion of (necessary) entailment.[40] The basic idea is that some facts ground another just in case the former necessitate both that the latter obtains and that it is less fundamental than they are. Relative fundamentality is represented via a totally ordered system of *levels*. Each fact is assigned a unique

[39] Not that Correia claims otherwise, I hasten to add.

[40] I should emphasize that Correia does not endorse this account, but only makes the far weaker claim that it merits discussion.

level, and a fact is more fundamental than another just in case it belongs to a lower level. Correia shows that the logic of ground he obtains validates versions of our above introduction rules.[41]

He then points out that within the resulting system, there is a natural way of defining a relation of propositional equivalence, which has the property that whenever two formulas are propositionally equivalent, they are intersubstitutable in the argument places of the grounding operator. In this way, propositional equivalence plays roughly the role of ground-theoretic equivalence.[42] That relation is defined as obtaining between two propositions just in case they are necessarily equivalent and necessarily equally fundamental if true. We may thus consider the following proposal for a general Relative Fundamentality-Based Account of ground-theoretic equivalence (cf. Correia 2018, p. 19f):

RFBA Sentences S and T are ground-theoretically equivalent iff the propositions expressed by S and T are (i) necessarily equivalent and (ii) necessarily equally fundamental if true.

Although Correia does not determine the exact resulting logic of ground-theoretic equivalence, he does specify for a number of important principles whether they hold or not. Indeed, it is not hard to verify that all of the rules of our extensional system turn out to be valid under the interpretation of $\approx$ as expressing the relative-fundamentality-based notion of equivalence.[43]

This equivalence relation is not co-extensional with the mode-theoretically defined one, though: it is strictly wider. Consider, for example, the pair of $\bot \vee (\top \vee \bot')$ and $(\bot \vee \top) \vee \bot'$, where $\top$ is some necessary truth, and $\bot$ and $\bot'$ are distinct necessary falsehoods. Evidently, the two propositions are necessarily true. Moreover, the level of both propositions is necessarily that of $\top$ plus 2. This is because the level of a disjunction with just one true disjunct is defined by Correia to be that of the true disjunct plus 1 (cf. Correia 2018, p. 6). Hence, $\bot \vee (\top \vee \bot')$ and $(\bot \vee \top) \vee \bot'$ are equivalent under Correia's account. But they are inequivalent under the mode-theoretic account, since only the former may be verified via $\top \vee \bot'$, and only the latter via $\bot \vee \top$.

It seems to me that the relative-fundamentality-based account faces serious objections that concern necessary falsehoods and necessary equivalence.[44] Firstly, the account permits no distinction between necessary falsehoods: they are all necessarily equivalent and necessarily equally fundamental if true. As a result, the account also

[41] Cf. Correia (2018, p. 17f). I say versions of, because they target a factive understanding of ground, and therefore require the truth of the putative grounds as premises.

[42] As in the case of factual equivalence discussed in 4.1, Correia only claims, in effect, that propositional equivalence implies ground-theoretic equivalence, and does not commit to the converse. And as before, for simplicity I focus on the view that results from also endorsing the converse, although again all my critical comments pertain purely to the direction of implication that holds on the account Correia describes.

[43] Correia himself mentions the DeMorgan principles and the commutativity principles. The structural rules and preservation rules are unproblematic. It is moreover clear that the collapse principles $A \vee A \approx A \wedge A \approx \neg\neg A$ turn out valid since the relevant propositions are clearly necessarily equivalent and necessarily equally fundamental. It is only slightly harder to verify that the principles I have described as introduction rules for $\leq$, with the abbreviations suitably unpacked, also turn out valid in Correia's system.

[44] Of course, this is not to say that the account could not be modified in such a way as to avoid these problems. At the very end of his paper, Correia himself mentions potential problems similar to those discussed here, and hints at a possible refinement of the view that avoids them, and potentially also the

cannot respect some intuitive ground-theoretic distinctions in the realm of truths. For suppose P is true. Then (P or $2 + 2 = 5$) grounds that (P or $2 + 2 = 5$) or Q, for arbitrary Q. But it then follows that likewise (P or snow is white and not white) grounds that (P or $2 + 2 = 5$) or Q, which seems counter-intuitive.

The second problem concerns necessary equivalence, and seems even more serious to me. It is based on the widely held assumptions that the existence of a set is necessary given the existence of its members, and indeed strictly and fully grounded by the existence of its members. It is then also very plausible to take the fact that a certain set exists to belong to level n+1, where n is the highest level occupied by any fact to the effect that a certain member of the set exists. But then consider the following truths:

1. {{Socrates}} exists
2. {{Socrates}, Socrates} exists

By our assumptions, they are necessarily equivalent and they are necessarily equally fundamental—hence they turn out to be ground-theoretically equivalent. But this is a very implausible result. For instance, only the former truth seems to ground that {{{Socrates}}} exists.[45]

4.4 Poggiolesi's account

The final account I want to consider and compare to my own is that recently developed by Francesca Poggiolesi. In her (2016b), Poggiolesi proposes a definition, in syntactic form and for a standard propositional language, of a relationship of grounding. In her (2018), she then specifies a suitable deductive system that she proves to be sound and complete relative to the definition of grounding. As she makes clear, implicit in her account of grounding is a distinctive view of ground-theoretic equivalence. Poggiolesi's account is in some ways very close to my own. In particular, it is characterized by repetition-sensitivity and order-insensitivity, and it still counts as ground-theoretically equivalent certain sentences with different connectives and overall different representational structures. Still, as we shall see, there are also significant differences between our views.

Before addressing specifically the matter of ground-theoretic equivalence, we should note that the notion of *grounding* that Poggiolesi investigates in these papers is somewhat different from the notion I have worked with, and that is in focus in most other recent contributions to the grounding debate. For the notion Poggiolesi discusses is explicitly a notion of *complete* and *immediate formal* grounding (2016b, p. 3150f). Each of the highlighted adjectives requires comment. First, by *formal* grounding, Poggiolesi means what others have called *logical* grounding, i.e. grounding that holds in virtue of the logical forms of the relevant sentences alone. Second, whereas most

Footnote 44 continued
ones I describe. The crucial move would be to appeal to a richer conception of worlds that allows both for incomplete and for inconsistent worlds. Still, as far as the version that is given a detailed development in the paper is concerned, my objections apply.

[45] The same example, given my assumptions, also constitutes a direct objection to the account of ground that Correia describes, for the account then immediately implies that the truth that {{Socrates}, Socrates} exists grounds that {{{Socrates}}} exists. Thanks here to Fabrice Correia for helpful discussion.

authors think of grounding as a transitive (or normally transitive) relation, Poggiolesi's focus is on the special case of *immediate* grounding.[46] Third, by a *complete* immediate ground of a truth, Poggiolesi understands the collection of *all* partial immediate grounds of the truth. Thus a complete immediate ground is not the same as a *full* immediate ground in our sense. For instance, if P and Q are both true, then each of them would normally be seen as a full immediate ground of $P \vee Q$. But only their collection $\{P, Q\}$ is a complete immediate ground of $P \vee Q$.[47] The case of disjunction also illustrates the motivation for another deviation in Poggiolesi's account from previous views of ground, namely that ground is conceived as relative to a background 'robust condition' (2016b, p. 3159ff): P is a complete and immediate formal ground of $P \vee Q$ under the condition that $\neg Q$ holds, but not under the condition that Q holds.[48]

It would take us too far here to explain the details of Poggiolesi's elaborate definition of ground, but we can perhaps convey some of the spirit behind it. The fundamental idea is that logical grounding is a matter of (a) the groundee being derivable from the grounds, (b) the grounds being less complex, in a particular sense, than the groundee, and (c) the negation of the groundee being derivable from the collection of the negation of the grounds, together with the (possibly empty) robust condition. For instance, P grounds $P \vee Q$ under robust condition $\neg Q$ since $\{P, \neg Q\}$ is less complex, in the relevant sense, than $P \vee Q$, $P \vee Q$ is derivable from P, and $\neg(P \vee Q)$ is derivable from $\neg P$ together with the robust condition $\neg Q$.[49]

Despite the differences in the targeted concept of ground, some of the principles Poggiolesi accepts or rejects correspond in a clear way to principles in our own framework, so that a meaningful comparison is possible. In fact, Poggiolesi herself helpfully highlights two significant differences between her view and the views considered in Correia (2014), Fine (2012a) and Schnieder (2011), which in the relevant respects agree with the account defended here (cf. her 2016b, sec. 7, 2018, sections 2, 6). The *first* difference concerns the question of the immediate grounds of a proposition of the form $\neg(\neg P \vee \neg Q)$. On my view, such a proposition is immediately grounded by $\neg\neg P, \neg\neg Q$. On Poggiolesi's view, in contrast, it is not grounded by that collection at all. Instead, its immediate ground is P, Q. This is connected to a difference explicitly concerning ground-theoretic equivalence. On my view, the De Morgan equivalents $\neg(P \vee Q)$ and $\neg P \wedge \neg Q$ are ground-theoretically equivalent. On Poggiolesi's view, they are not. The *second* difference concerns the principles of associativity for con-

[46] A more standard notion might then of course be obtained from the immediate one by closing it under a suitable transitivity principle.

[47] As Poggiolesi highlights, both in imposing an immediacy and a completeness requirement, she is following in the footsteps of Bernard Bolzano, who was the first to develop a general and systematic theory of a notion of grounding that is very close to the contemporary notion(s); see esp. his Bolzano (1837), translated as Bolzano (2014). For a recent book-length study of Bolzano on grounding, see Roski (2017).

[48] Strictly speaking, the condition is not that $\neg Q$ is true, but that what Poggiolesi calls the *converse* of Q is true, which is sometimes but not always identical to $\neg Q$. The difference does not matter for our purposes. See Poggiolesi (2016b, p. 3155) for the precise definition of the converse of a formula. (Thanks here to an anonymous referee.)

[49] The caveat of the previous footnote with respect to $\neg Q$ applies here as well. Derivability here is simply classical derivability. The relevant sense of comparative complexity—what Poggiolesi calls 'completely and immediately less g-complex than'—has a somewhat complicated syntactic definition, for the details of which I have to refer the reader to Poggiolesi's article, esp. definition 4.8 in (2016b, p. 3158).

junction and disjunction. On Poggiolesi's view, the propositions $P \wedge (Q \wedge R)$ and $(P \wedge Q) \wedge R$ are ground-theoretically equivalent, and likewise for disjunction. On my view, as we have discussed, they are not. Correspondingly, Poggiolesi takes $P,\ Q \wedge R$ to be an immediate ground of $(P \wedge Q) \wedge R$, whereas I do not.

So like my proposed account, Poggiolesi endorses a more fine-grained picture of ground-theoretic equivalence than the worldly, truthmaker-based account, rejecting, for example, that $A \approx \neg\neg A$ and that $A \wedge (B \vee C) \approx (A \wedge B) \vee (A \wedge C)$. But whereas my account retains the DeMorgan equivalences and gives up Associativity, Poggiolesi's account gives up DeMorgan and retains Associativity. This lends to her view a very distinctive and novel character, and it would be interesting to see if something like the present mode-theoretic framework could also be used to provide a semantics for her view.

I know of no decisive reasons to prefer my account to Poggiolesi's, or hers over mine. The combination of principles accepted under my approach, in conjunction with the introduction and elimination principles for ground, strikes me as considerably more intuitive than Poggiolesi's. But Poggiolesi has different intuitions (cf. her 2016b, p. 3165f), and so appeals to intuitions do not seem to provide a way forward here. A justified choice between the accounts, I suspect, will have to be based on a comprehensive evaluation of their overall theoretical virtues, which it is beyond the scope of this paper to carry out.[50]

However, Poggiolesi's work also contains the material for a direct challenge to my account. For I have not only rejected the associativity principles, I have also endorsed the principles of *commutativity* as principles of ground-theoretic equivalence. And the discussion in Poggiolesi (2016b, p. 3156) strongly suggests that she takes these two kinds of principles to stand and fall together. It is therefore worth considering whether specifically the combination of the commutativity principles with the rejection of the associativity principles is problematic.

Why might one take commutativity and associativity to stand and fall together? One idea might be that both simply amount to an insensitivity to order in the relevant expressions.[51] As I have stressed, the sentences $A \wedge B$ and $B \wedge A$ differ only in the order of the arguments of $\wedge$. Then perhaps in the case of $A \wedge (B \wedge C)$ and $(A \wedge B) \wedge C$, one might say the difference is purely whether $\wedge$ is first applied to B and C, and then also to A, or first to A and B, and then also to C. But this is a misleading description of the case. For it is *not* the case that in both sentences, the same formulas show up as arguments of $\wedge$, only in a different order. In $A \wedge (B \wedge C)$, the complex formula $B \wedge C$ is an argument of $\wedge$, whereas in $(A \wedge B) \wedge C$, that formula does not occur at all. And it is just this difference that is reflected, on my view, in the grounds of the propositions expressed by the formulas: the proposition expressed by $B \wedge C$ is a partial ground

[50] Here are two considerations that would seem relevant to such an evaluation. Poggiolesi (2016a) points out that her account underwrites certain connections between ground and the normality of proofs. To the extent that it is indepedently plausible that there should be such connections, this would be a point in favour of her account. On the other hand, Poggiolesi's logic of ground appears not to be closed under substitution, which is also a property that is often seen as desirable feature of a logic. (Although $\neg(\neg P \vee \neg Q)$ is not grounded by $\neg\neg P, \neg\neg Q$, when P and Q are not themselves negations, then even on Poggiolesi's account, $\neg(P \vee Q)$ *is* grounded by $\neg P,\ \neg Q$.) So if a case can be made that the logic of ground should be closed under substitution, this would favour the present account.

[51] This was suggested to me by an anonymous referee.

of the proposition expressed by $A \wedge (B \wedge C)$, but not of the proposition expressed by $(A \wedge B) \wedge C$. Of course, one could try to argue that this difference should not be reflected ground-theoretically. But one cannot do this simply on the grounds that the differences between $A \wedge B$ and $B \wedge A$ should not be reflected ground-theoretically. These are a different *kind* of difference.

Another idea might be that with respect to both commutativity and associativity, the crucial point is that the differences between the relevant pairs of sentences are purely *notational*, and thus not the kind of difference to which ground can plausibly be considered sensitive.[52] Thus, in the case of $A \wedge B$ and $B \wedge A$, one might say that these are merely alternative notations for the application of conjunction to *one and the same unordered pair* of contents. In the case of $A \wedge (B \wedge C)$ and $(A \wedge B) \wedge C$, the claim might then be that these are merely alternative notations for the application of conjunction to *one and the same unordered triple* of contents. However, on closer inspection, the second claim is not plausible. For $A \wedge (B \wedge C)$ and $(A \wedge B) \wedge C$ both involve *two* applications of conjunction, and each of these applications is to a pair, not a triple. This is simply a consequence of the grammar of our standard propositional language, in which $\wedge$ is a two-place operator.

Still, one may feel that there is something to the thought that somehow $A \wedge (B \wedge C)$ and $(A \wedge B) \wedge C$ are just alternative notations for the same thing. Fortunately, this can be explained in a different way, compatible with the rejection of the associativity principles. For suppose that one thinks that each of A, B, and C are true, and simply wishes to say as much. It then feels artificial to have to choose between $(A \wedge B) \wedge C$ and $A \wedge (B \wedge C)$ as vehicles for saying what one wants to say. Relative to one's communicative intentions, the difference between the two is irrelevant. However, it doesn't follow that there are not ground-theoretic differences between the contents of $(A \wedge B) \wedge C$ and $A \wedge (B \wedge C)$. It may simply be that relative to one's communicative goals, these ground-theoretic differences, too, are irrelevant. Moreover, if it is seen as a bad thing that one has to choose between different, ground-theoretically inequivalent sentences to say what one wants to say, what this seems to suggest is that one should use a different *language*, namely one with a conjunction-operator $\bigwedge$ operating on a set or multi-set. Then one could simply utter $\bigwedge\{A, B, C\}$, without needing to make an arbitrary choice between alternative sentences. I think there are advantages to using such a language, but to facilitate comparison with previous approaches, I here focus on a more standard language with only binary conjunction and disjunction operators.

Finally, some may simply find it highly counter-intuitive that conjunction and disjunction should be commutative without being associative with respect to ground-theoretic equivalence.[53] So to close this discussion, let me try to offer an analogy that may help make the rejection of associativity more intelligible to such an opponent. I think there is a very natural rough picture of how truth-functional operations interact with the ground-theoretic hierarchy which should lead one to *expect* that associativity, but not commutativity, will fail.[54] Consider the set-theoretic operation of *pairing*,

[52] Thanks to two anonymous referees for pressing me on this point.

[53] An anonymous reviewer has expressed this sentiment.

[54] I do not mean to deny that there are also natural rough pictures which would not lead one to expect this. Indeed, I believe there are such alternative pictures. But that is okay, for my aim here is not to establish that

mapping arguments x and y to the set $\{x, y\}$. Note that it is commutative but not associative: if $x \neq z$, then $\{x, \{y, z\}\} \neq \{\{x, y\}, z\}$. Now one natural picture has it that the truth-functions behave with respect to the ground-theoretic hierarchy somewhat analogously to how pairing behaves with respect to the set-theoretic hierarchy. In particular, there is an analogous reason for the failure of associativity: it lies in the fact that, roughly speaking, each application of the relevant function involves "raising the level"—either set-theoretic or ground-theoretic—of the arguments of the function. If a function f has this level-raising character, it is easy to see how associativity may fail, for in $f(f(x, y), z)$ the arguments x and y are raised twice, as it were, and z is raised only once, whereas in $f(x, f(y, z))$, x is raised only once, and y and z twice.[55] Now if one thinks of ground and the truth-functions on something like this level-raising model, then it does not seem at all counter-intuitive to reject associativity and accept commutativity. For the level-raising nature of the truth-functions naturally leads to a violation of associativity, without presenting any threat whatsoever to commutativity.

4.5 Worldly versus representational ground

Finally, I want to briefly examine whether the mode-theoretic account justifies a view of the conception of ground it is adequate for as *representational*. The answer to this naturally depends on how the worldly/representational distinction is spelled out. I think there are at least three prima facie natural ways of doing this. On the first, ground is worldly (representational) just in case the *relata* of ground are worldly (representational) entities.[56] If we consider the mode-theoretic account as a guide to the nature of the relata of ground, then it would appear justified to say that it yields a representational conception of ground. For the relata of ground are then mode-theoretic propositions, which are clearly representational entities—roughly speaking, they represent that at least one of their modes is actual. But then by parallel reasoning it seems that if we consider the truthmaker-theoretic account as a guide, we would also appear to be justified in saying that it yields a representational conception of ground. For here, too, the relata are clearly representational entities, representing that at least one of their verifiers obtains.

The second explication tracks not the worldliness or otherwise of the relata of ground, but the worldliness or otherwise of the *differences between the relata* that ground is sensitive to.[57] Since a difference in verifying states is plausibly seen as a worldly difference between propositions—it concerns purely how the proposition relates to worldly entities—, the truthmaker account is then classified as worldly. A difference in modes of verification, in contrast, is not plausibly seen as worldly, since it typically concerns how the proposition relates to other propositions as well as the world

Footnote 54 continued
commutativity holds and associativity not. It is merely to show that one may *reasonably* take this to be so, and thus that a view of this sort should be considered a serious contender.

55 I am not claiming that it is *impossible* to define a level-raising associative function. But starting from the idea of a level-raising combination like pairing, one very *naturally* ends up with a non-associative function.

56 This explication fits nicely with a number of formulations found in Correia's papers, cf. e.g. (Correia 2010, p. 256f, 2017, p. 508).

57 This explication was suggested to me by Kit Fine.

(the exception being a difference in direct modes of verification). The mode-theoretic account would accordingly be counted as representational.

The third explication appeals to whether the relata of ground are individuated in terms of *representational structure* or not. Under that explication, even the mode-theoretic views plausibly qualify as worldly. For even under the repetition-sensitive conception of modes, since the DeMorgan rules hold for ground-theoretic equivalence, sentences with significantly different representational structures are still seen as ground-theoretically equivalent. In particular, as the case illustrates, the account does not permit an exclusive division between negative and positive propositions, as one would expect of an account that individuated propositions in terms of representational structure. Under the repetition-insensitive conception, not even an exclusive division between conjunctive and disjunctive propositions can be made, since no distinction is made between self-conjunctions and self-disjunctions.

How should we choose between these explications? Clearly, an explication will be unsatisfactory if it yields a distinction that carries no theoretical significance. If the worldly/representational distinction is to have any theoretical interest, it must presumably be correlated with a difference in theoretical roles that the conceptions of ground either side of the divide are suited to play. I have no firm view of the matter, and I lack the space to pursue it here. It may well be, however, that all of the dividing lines just sketched correspond to a difference in plausible theoretical roles. In that case, the most informative description of the status of the mode-theoretic accounts is that they occupy an intermediate position in regard to worldliness between the truthmaker account and the representational account of Correia's.

Acknowledgements I would like to thank Fabrice Correia, Kit Fine, Martin Glazier, and Stefan Roski for very helpful discussion and feedback. I am also grateful to two anonymous referees for this journal whose comments have led to several substantial improvements to the paper.

A Soundness and completeness

A.1 Definitions

Recall the definition of $\mathscr{L}_{\approx}$ as based on a propositional language with connectives $\wedge$, $\vee$, and $\neg$, augmented by all expressions of the form $A \approx B$ where $A,\ B$ are sentences of the propositional language (and thus do not already include occurrences of $\approx$). We call expressions of the latter form *equivalences*, and reserve the label 'sentence' for the purely truth-functional sentences. The language comprising only the sentences of $\mathscr{L}_{\approx}$ will be called $\mathscr{L}_B$.

For ease of reference, we repeat some relevant definitions from Krämer (2018):

Definition 1 *(Mode-Space)* A *mode-space* is a pair $\langle M, V\rangle$ such that

1. M is a non-empty set
2. V is a non-empty, partial function taking non-empty countable sequences of non-empty subsets of M into members of M
3. the domain of V is closed under non-empty subsequences and countable concatenation of sequences

4. $V(\gamma_1^\frown \gamma_2^\frown \ldots) = V(\delta_1^\frown \delta_2^\frown \ldots)$ whenever $V(\gamma_1) = V(\delta_1)$, $V(\gamma_2) = V(\delta_2), \ldots$

Informally, M is the set of modes. Any non-empty set of modes is a proposition. V is the via-function, mapping some sequences of propositions $P_1, P_2, \ldots$ to the mode of verifying via $P_1, P_2, \ldots$. Modes which are never the value of V are called *fundamental*, and their set is denoted M^F. All other modes are called derivative and their set is denoted M^D. Of any derivative modes m and n we can form the fusion $m \sqcup n$ which is $V\langle P_1, P_2, \ldots, Q_1, Q_2, \ldots\rangle$ when $V\langle P_1, P_2, \ldots\rangle = m$ and $V\langle Q_1, Q_2, \ldots\rangle = n$. If V is defined for $\langle P\rangle$, P is called *raisable*, and the set of raisable contents is denoted $\mathscr{R}$.

Conjunction, disjunction, and negation are now defined as follows. First, we define binary operations of fusion ($\sqcup$) and (disjunctive) addition ($+$) for propositions that are subsets of M^D:

$$\begin{aligned} P \sqcup Q &:= \{m \sqcup n : m \in P \text{ and } n \in Q\} \\ P + Q &:= (P \cup Q) \cup (P \sqcup Q) \end{aligned}$$

Next, for raisable contents P, we define an operation of *raising* on unilateral contents, which given an input P yields a content $\uparrow P$ just like P except in that it may be verified via P, and in terms of raising, fusion, and addition, we define conjunction and disjunction on raisable unilateral contents:

$$\begin{aligned} \uparrow P &:= \{V\langle P\rangle\} \text{ if } P \cap M^D \text{ is empty, } \{V\langle P\rangle\} + (P \cap M^D) \text{ otherwise} \\ P \wedge Q &:= \uparrow P \sqcup \uparrow Q \\ P \vee Q &:= \uparrow P + \uparrow Q \end{aligned}$$

Finally, we define conjunction, disjunction, and negation on bilateral contents in $\mathscr{R} \times \mathscr{R}$:

$$\begin{aligned} \neg\mathbf{P} &:= \langle \mathbf{P}^-, \uparrow\mathbf{P}^+\rangle \\ \mathbf{P} \wedge \mathbf{Q} &:= \langle \mathbf{P}^+ \wedge \mathbf{Q}^+, \mathbf{P}^- \vee \mathbf{Q}^-\rangle \\ \mathbf{P} \vee \mathbf{Q} &:= \langle \mathbf{P}^+ \vee \mathbf{Q}^+, \mathbf{P}^- \wedge \mathbf{Q}^-\rangle \end{aligned}$$

For the operations so defined to behave as desired, the background mode-space needs to satisfy two important conditions:

Definition 2 A mode-space $\langle M, V\rangle$ is called *complete* iff $P \sqcup Q \in \mathscr{R}$, $P + Q \in \mathscr{R}$, and $\uparrow P \in \mathscr{R}$ whenever $P, Q \in \mathscr{R}$.

Definition 3 A mode-space $\langle M, V\rangle$ is called *constrained* iff $V(\gamma) = V(\delta)$ only if the same ground-set corresponds to γ and δ.

Note that in a constrained mode-space, every derivative mode m corresponds to a unique ground-set, which is denoted by $|m|$. We shall henceforth deal only with complete and constrained mode-spaces.

Definition 4 A unilateral proposition P is

- *irreflexive* iff: P does not occur in γ whenever $V(\gamma) \in P$,
- *closed* iff: P contains a mode with ground-set $\Gamma_1 \cup \Gamma_2 \cup \ldots$ whenever P contains modes with ground-sets $\Gamma_1, \Gamma_2, \ldots$,

- *transitive* iff: P includes a mode with ground-set Γ, Δ whenever P includes a mode with ground-set Δ, Q and Q includes a mode with ground-set Γ.
- *normal* iff: irreflexive, closed, and transitive.

We define the two classes of mode-spaces with respect to which we shall establish soundness and completeness results.

Definition 5 A mode-space $\langle M, V \rangle$ is

- *intermediate* iff: $V(\gamma) = V(\delta)$ whenever the same multi-set underlies γ and δ
- *extensional* iff: $V(\gamma) = V(\delta)$ whenever the same set underlies γ and δ

Note that the class of intermediate mode-spaces is exactly the class of mode-spaces compatible with an order-insensitive but repetition-sensitive conception of modes, whereas the class of extensional mode-spaces correspondingly reflects the repetition-insensitive conception of modes.

We are now in a position to give a mode-space *semantics* for $\mathscr{L}_{\approx}$ by defining the notions of a model, truth in a model, and validity in a class of models.

Definition 6 If $\langle M, V \rangle$ is a mode-space, then $\mathscr{M} = \langle M, V, [\cdot] \rangle$ is a *model based on* $\langle M, V \rangle$ just in case $[\cdot]$ is a function mapping every sentence in $\mathscr{L}_{\approx}$ to a member of $\mathscr{R} \times \mathscr{R}$ so that for all sentences $A, B \in \mathscr{L}_{\approx}$:

- $[\neg A] = \neg[A]$
- $[A \wedge B] = [A] \wedge [B]$
- $[A \vee B] = [A] \vee [B]$

$\mathscr{M}$ is a model just in case $\mathscr{M}$ is a model based on some mode-space.

Definition 7 (Truth and Validity) For any equivalence $A \approx B \in \mathscr{L}_{\approx}$:

- $A \approx B$ is *true in a model* $\mathscr{M}$ ($\mathscr{M} \models A \approx B$) iff $[A] \approx [B]$
- $A \approx B$ is *valid in a class of models* $\mathscr{C}$ ($\models_{\mathscr{C}} A \approx B$) iff true in every model in $\mathscr{C}$.

A.2 Preparatory results

For ease of reference, we repeat the central theorems of Krämer (2018) that we shall need. Throughout, we tacitly restrict attention to complete and constrained mode-spaces.

Lemma 1 *The introduction and elimination principles for bilateral propositions hold (theorems 5 and 6):*

$(<\wedge)$ $\boldsymbol{\Gamma} < \mathbf{P} \wedge \mathbf{Q}$ *iff* $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
$(<\vee)$ $\boldsymbol{\Gamma} < \mathbf{P} \vee \mathbf{Q}$ *iff* $\boldsymbol{\Gamma} \leq \mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \mathbf{Q}$ *or* $\boldsymbol{\Gamma} \leq \{\mathbf{P}, \mathbf{Q}\}$
$(<\neg\neg)$ $\boldsymbol{\Gamma} < \neg\neg\mathbf{P}$ *iff* $\boldsymbol{\Gamma} \leq \mathbf{P}$
$(<\neg\wedge)$ $\boldsymbol{\Gamma} < \neg(\mathbf{P} \wedge \mathbf{Q})$ *iff* $\boldsymbol{\Gamma} \leq \neg\mathbf{P}$ *or* $\boldsymbol{\Gamma} \leq \neg\mathbf{Q}$ *or* $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$
$(<\neg\vee)$ $\boldsymbol{\Gamma} < \neg(\mathbf{P} \vee \mathbf{Q})$ *iff* $\boldsymbol{\Gamma} \leq \{\neg\mathbf{P}, \neg\mathbf{Q}\}$

Lemma 2 *The following structural principles hold for normal unilateral propositions (theorem 7):*

1. $P \not< P$
2. *If* $\Gamma_1 < P$, $\Gamma_2 < P$, *..., then* $\Gamma_1, \Gamma_2, \ldots < P$.
3. *If* $\Gamma \leq P$ *and* $Q \prec P$ *for all* $Q \in \Gamma$*, then* $\Gamma < P$.
4. *If* $\Gamma, P \leq P$*, then* $\Gamma \leq P$.
5. *If* $\Gamma_1 \leq P$, $\Gamma_2 \leq P$, *..., then* $\Gamma_1 \cup \Gamma_2 \cup \ldots \leq P$.
6. *If* $\Gamma < P$ *and* $\Delta, P < Q$*, then* $\Gamma, \Delta < Q$.
7. *If* $\Gamma_1 \leq P_1$, $\Gamma_2 \leq P_2$, *..., and* $P_1, P_2, \ldots \leq Q$ *then* $\Gamma_1, \Gamma_2, \ldots \leq Q$
8. *If* $P \preceq Q$ *and* $Q \prec R$ *then* $P \prec R$
9. *If* $P \prec Q$ *and* $Q \preceq R$ *then* $P \prec R$
10. *If* $P \preceq Q$ *and* $Q \preceq R$ *then* $P \preceq R$

Lemma 3 *Normality of unilateral propositions is preserved under* $\wedge$, $\vee$, *and* $\uparrow$ *(theorem 8).*

We state without proof the following straightforward facts about ground-theoretic equivalence:

Lemma 4 $\approx$ *has the following properties:*

1. $\approx$ *is an equivalence relation—for all* $\mathbf{P}, \mathbf{Q}, \mathbf{R}$ *we have (i)* $\mathbf{P} \approx \mathbf{P}$*, (ii) if* $\mathbf{P} \approx \mathbf{Q}$ *then* $\mathbf{Q} \approx \mathbf{P}$*, and (iii) if* $\mathbf{P} \approx \mathbf{Q}$ *and* $\mathbf{Q} \approx \mathbf{R}$ *then* $\mathbf{Q} \approx \mathbf{R}$
2. $\approx$ *is preserved under conjunction, disjunction, and double negation—for all* $\mathbf{P}, \mathbf{Q}, \mathbf{R}$*, if* $\mathbf{P} \approx \mathbf{Q}$*, then (i)* $\mathbf{P} \wedge \mathbf{R} \approx \mathbf{Q} \wedge \mathbf{R}$*, (ii)* $\mathbf{P} \vee \mathbf{R} \approx \mathbf{Q} \vee \mathbf{R}$*, and (iii)* $\neg\neg\mathbf{P} \approx \neg\neg\mathbf{Q}$
3. $\approx$ *satisfies the DeMorgan rules—for all* $\mathbf{P}, \mathbf{Q}$*: (i)* $\neg(\mathbf{P} \wedge \mathbf{Q}) \approx \neg\mathbf{P} \vee \neg\mathbf{Q}$ *and (ii)* $\neg(\mathbf{P} \vee \mathbf{Q}) \approx \neg\mathbf{P} \wedge \neg\mathbf{Q}$

We now establish some easy *sufficient* identity conditions for unilateral contents:

Lemma 5 *Let* P, Q *be raisable unilateral propositions.*

(i) *If the mode-space is intermediate:* $P \wedge Q = Q \wedge P$ *and* $P \vee Q = Q \vee P$.
(ii) *If the mode-space is extensional:* $P \wedge P = P \vee P$.
(iii) *If the mode-space is extensional and* P *closed:* $P \vee P = \uparrow P$

Proof For (i): We show that the operation of fusion on the modes is commutative. Let m and n be derivative modes, and suppose $m = V(\gamma)$ and $n = V(\delta)$. By the definition of fusion, $m \sqcup n = V(\gamma^\frown\delta)$ and $n \sqcup m = V(\delta^\frown\gamma)$. Since the mode-space is assumed to be intermediate, V maps two sequences to the same mode if they correspond to the same multi-set, and since $\gamma^\frown\delta$ and $\delta^\frown\gamma$ correspond to the same multi-set, it follows that $m \sqcup n = n \sqcup m$. From this, the result follows immediately by definition of $\wedge$ and $\vee$.

For (ii): By application of the definitons, $P \wedge P = \uparrow P \sqcup \uparrow P \subseteq \uparrow P \cup \uparrow P \cup (\uparrow P \sqcup \uparrow P) = \uparrow P + \uparrow P = P \vee P$. It remains to show that $\uparrow P \subseteq \uparrow P \sqcup \uparrow P$. This follows from the idempotence of $\sqcup$ as defined on derivative modes in extensional mode-spaces. Let m be a derivative mode and assume $m = V(\gamma)$. Then $m \sqcup m = V(\gamma^\frown\gamma)$ and since the same set underlies γ and $\gamma^\frown\gamma$, by extensionality of the mode-space, $m = V(\gamma) = V(\gamma^\frown\gamma) = m \sqcup m$.

For (iii): $\uparrow P \subseteq P \vee P$ is immediate by definition of $\vee$. It remains to show that $\uparrow P \sqcup \uparrow P \subseteq \uparrow P$. So let $m \in \uparrow P \sqcup \uparrow P$. Then for some m_1, m_2: $m = m_1 \sqcup m_2$

and $m_1 \in\uparrow P$ and $m_2 \in\uparrow P$. By assumption, P is closed. It is straightforward to show that $\uparrow P$ is then also closed, and so $\uparrow P$ contains some mode with ground-set $|m_1| \cup |m_2| = |m_1 \sqcup m_2| = |m|$. Since in an extensional mode-space, no two modes have the same ground-set, it follows that $m \in\uparrow P$ and thus $P \vee P \subseteq\uparrow P$. □

We may also establish some substantive *necessary* conditions for certain kinds of unilateral propositions to be identical:

Lemma 6 *Let $P,\ Q,\ R,\ S$ be normal raisable unilateral propositions.*

(i) $\uparrow P =\uparrow Q$ *implies* $P = Q$
(ii) $P \wedge Q =\uparrow R$ *implies* $P = Q = R$
(iii) $P \vee Q =\uparrow R$ *implies* ($P = R$ *and* $Q \leq R$) *or* ($Q = R$ *and* $P \leq R$)
(iv) $P \wedge Q = R \wedge S$ *implies* $\{P,\ Q\} = \{R,\ S\}$
(v) $P \wedge Q = R \vee S$ *implies* $P = Q$ *and* (($R = P$ *and* $S \leq R$) *or* ($S = P$ *and* $R \leq S$))
(vi) $P \vee Q = R \vee S$ *implies that either*

a. $\{P,\ Q\} = \{R,\ S\}$, *or*
b. $P = R$ *and* $Q \leq P$ *and* $S \leq R$, *or*
c. $P = S$ *and* $Q \leq P$ *and* $R \leq S$, *or*
d. $Q = R$ *and* $P \leq Q$ *and* $S \leq R$, *or*
e. $Q = S$ *and* $P \leq Q$ *and* $R \leq S$

Proof By use of the equivalences noted at the beginning of this section and the transitivity and antisymmetry of $\leq$ and $\preceq$.

For (i): Since $P <\uparrow P$, it follows from the antecedent that $P <\uparrow Q$ and hence $P \leq Q$. Likewise since $Q <\uparrow Q$, it follows that $Q <\uparrow P$ and hence $Q \leq P$. By antisymmetry of $\leq$, $P = Q$.

For (ii): Since $P, Q < P \wedge Q$, it follows from the antecedent that $P, Q <\uparrow R$ and hence $P, Q \leq R$. So $P \preceq R$ and $Q \preceq R$. Moreover, $R <\uparrow R$ so $R < P \wedge Q$, so $R \leq \{P, Q\}$. It follows that $R \leq P$ and $R \leq Q$, so $R \preceq P$ and $R \preceq Q$, and hence by antisymmetry of $\preceq$ that $R = P$ and $R = Q$.

For (iii): In similar fashion as before, it follows from the antecedent that $P \leq R$ and $Q \leq R$, as well as that either (a) $R \leq P$ or (b) $R \leq Q$. If (a), then by antisymmetry of $\leq$ we have $P = R$, and if (b), we have $Q = R$.

For (iv): From the antecedent it is straightforward to show (1) that either (1a) $P \preceq R$ and $Q \preceq S$ or (1b) $P \preceq S$ and $Q \preceq R$ and (2) that either (2a) $R \preceq P$ and $S \preceq Q$ or (2b) $R \preceq Q$ and $S \preceq P$. If (1a) and (2a), then by antisymmetry of $\preceq$, $P = R$ and $Q = S$ follows. If (1b) and (2b), then it follows that $P = S$ and $Q = R$. If (1b) and (2a), we have $P \preceq S \preceq Q \preceq R \preceq P$ and so $P = Q = R = S$. Likewise if (1a) and (2b), we have $P \preceq R \preceq Q \preceq S \preceq P$ and so again $P = Q = R = S$. So in all four cases, $\{P, Q\} = \{R, S\}$.

For (v): From the antecedent it follows that $R \leq P$ and $R \leq Q$ and $S \leq P$ and $S \leq Q$. Moreover, either (a) $P, Q \leq R$ or (b) $P, Q \leq S$ or (c) $P, Q \leq \{R, S\}$. If (a), then $P \preceq R$ and so $P = R$. Similarly $Q \preceq R$, and so $Q = R$. Hence $S \leq P = Q = R$, establishing the consequent. If (b), then $P \preceq S$, so $P = S$, and $Q \preceq S$, so $Q = S$. Hence $R \leq P = Q = S$, again establishing the consequent.

Finally, if (c), then it easy to show that either $P \preceq R$ and $Q \preceq S$ or $P \preceq S$ and $Q \preceq R$. In the first case, $P = R \leq Q = S \leq P$, so $P = Q = R = S$. In the second case, $P = S \leq Q = R \leq P$, so again $P = Q = R = S$. Either way, the consequent is again established.

For (vi): From the antecedent it follows (1) that (1a) $P \leq R$ or (1b) $P \leq S$, and (2) that (2a) $Q \leq R$ or (2b) $Q \leq S$, and (3) that (3a) $R \leq P$ or (3b) $R \leq Q$, and (4) that (4a) $S \leq P$ or (4b) $S \leq Q$. It can now be shown that under each of the 16 possible combinations, one of the conditions a.–e. obtain. We shall restrict ourselves to an illustrative four cases; the remaining ones follow the same pattern. Suppose first that (1a), (2a), (3a), and (4a) obtain. Then by (1a) and (3a), $P = R$. By (2a), $Q \leq R = P$. By (4a), $S \leq P = R$. So case b. above obtains. Suppose now that instead of (4a), (4b) obtains. We still have $P = R$ and $Q \leq P$ as before. By (4b), $S \leq Q$, and since $Q \leq P$, we obtain $S \leq P = R$, as required for case b. For a different sort of case, suppose (1a), (2b), (3a), and (4b) obtain. Then still $P = R$, and by (2b), $Q \leq S$, as well as by (4b) $S \leq Q$, so $S = Q$. It follows that case a. above obtains. Finally, suppose that instead of (3a) and (4b), we have (3b) and (4a). Then $P \leq R$, $Q \leq S$, $R \leq Q$, and $S \leq P$. That is, $P \leq R \leq Q \leq S \leq P$, and thus $P = R = Q = S$. Again it follows that case a. obtains. □

Finally, in an extensional mode-space, binary weak full ground (among unilateral contents) can be characterized in terms of disjunction and identity:

Lemma 7 *For P, Q normal raisable unilateral propositions in an extensional mode-space: $P \leq Q$ iff $P \vee Q = Q \vee Q$.*

Proof For the right-to-left-direction, assume $P \vee Q = Q \vee Q$. Then $P < P \vee Q$ so $P < Q \vee Q$, so $P \leq Q$. For the left-to-right direction, assume $P \leq Q$. Then either $P = Q$ or $P < Q$. If $P = Q$, then evidently $P \vee Q = Q \vee Q$. So suppose $P < Q$. Given extensionality, to show that $P \vee Q = Q \vee Q$ it suffices to show for arbitrary Γ that $\Gamma < P \vee Q$ iff $\Gamma < Q \vee Q$. But if $\Gamma < Q \vee Q$, then $\Gamma \leq Q$, so $\Gamma < P \vee Q$. If $\Gamma < P \vee Q$, then either (a) $\Gamma \leq P$, or (b) $\Gamma \leq Q$, or (c) $\Gamma \leq \{P, Q\}$, so $\Gamma_P \leq P$ and $\Gamma_Q \leq Q$ for some Γ_P, Γ_Q with $\Gamma = \Gamma_P \cup \Gamma_Q$. If (a), then since $P < Q$ and $Q < Q \vee Q$, it is easy to show that $\Gamma < Q \vee Q$. if (b), then since $Q < Q \vee Q$, it is equally straightforward that $\Gamma < Q \vee Q$. If (c), then by the previous reasoning, $\Gamma_P \leq Q$ and $\Gamma_Q \leq Q$ and so $\Gamma_P \cup \Gamma_Q = \Gamma < Q \vee Q$. □

A.3 Adequacy of the intermediate system

For φ an equivalence in $\mathscr{L}_{\approx}$, we write $\vdash_{\mathfrak{I}} \varphi$ to say that φ is derivable within $\mathfrak{I}$, and we write $\models_{\mathfrak{I}} \varphi$ to say that φ is valid in the class of models based on complete, constrained, intermediate mode-spaces. We show first that $\mathfrak{I}$ is sound with respect to that class of models.

Theorem 8 (Soundness of the Intermediate System) *$\models_{\mathfrak{I}} \varphi$ whenever $\vdash_{\mathfrak{I}} \varphi$.*

Proof The soundness of (Comm.$\vee$) and (Comm.$\wedge$) is immediate from Lemma 5(i). The soundness of (Reflexivity), (Symmetry), (Transitivity), the preservation rules and the DeMorgan rules is immediate from Lemma 4. □

To prove completeness, we construct a canonical intermediate mode-space and model, and show that every equivalence that is true in this model is derivable within $\mathfrak{I}$. The canonical mode-space is defined as follows.

Definition 8 The *canonical intermediate mode-space* for $\mathscr{L}_{\approx}$ is the pair $\langle M_I, V_I\rangle$, where

- $M_0 := \{A \in \mathscr{L}_{\approx} : A \text{ is a literal}\}$
- $C_n := \wp(M_n)\backslash\{\emptyset\}$
- $M_{n+1} := \{m : m \in M_0$ or m is a non-empty multi-set of members of $C_n\}$
- $C := \bigcup_{n\in N} C_n$
- $M_I := \{m : m \in M_0$ or m is a non-empty multi-set of members of $C\}$
- $V_I(\gamma) = \Gamma$ if γ is a non-empty countable sequence of members of C and Γ is the underlying multi-set.
- $V_I(\gamma)$ is undefined otherwise.

We now establish that $\langle M_I, V_I\rangle$ is indeed a complete, constrained, intermediate mode-space. We show first that it is a mode-space.

Lemma 9 *$\langle M_I, V_I\rangle$ is a mode-space.*

Proof First, M_I is non-empty, since M_I includes all the literals in $\mathscr{L}_{\approx}$.

Second, V_I is a non-empty partial function that maps non-empty sequences of subsets of M_I to members of M_I. For since there are non-empty countable sequences of members of C and corresponding underlying multi-sets, V_I is non-empty. Since the members of C are subsets of M_I, the arguments of V_I are sequences of subsets of M_I. Since any multi-set underlying a non-empty countable sequence of members of C is a non-empty multi-set of members of C, the values of V_I are subsets are members of M_I.

Third, the domain of V_I is closed under non-empty subsequences and countable concatenations, for non-empty subsequences and countable concatenations of non-empty countable sequences of members of C are themselves such sequences.

Finally, $V_I(\gamma_1^\frown\gamma_2^\frown\ldots) = V_I(\delta_1^\frown\delta_2^\frown\ldots)$ whenever $V_I(\gamma_1) = V_I(\delta_1)$, $V_I(\gamma_2) = V_I(\delta_2), \ldots$ For the multi-set underlying $\gamma_1^\frown\gamma_2^\frown\ldots$ is determined by which items occur how many times in γ_1, and in γ_2, and $\ldots$, which is to say that it is determined by the multi-sets underlying $\gamma_1, \gamma_2, \ldots$. Since the multi-sets corresponding underlying $\gamma_1, \gamma_2, \ldots$ are the same as those underlying $\delta_1, \delta_2, \ldots$ whenever $V_I(\gamma_1) = V_I(\delta_1)$, $V_I(\gamma_2) = V_I(\delta_2), \ldots$, the result is then immediate. □

Next, we show that the recursive construction of the canonical mode-space is cumulative in the following sense:

Lemma 10 *In the construction of $\langle M_I, V_I\rangle$, for all n, $M_n \subseteq M_{n+1}$ and $C_n \subseteq C_{n+1}$.*

Proof Suppose $m \in M_0$. Then by definition, $m \in M_1$. Suppose $P \in C_0$. Then $\emptyset \subset P \subseteq M_0 \subseteq M_1$ and hence $P \in C_1$. Now assume that the claim holds up to n. Suppose $m \in M_{n+1}$. Then either $m \in M_0$, in which case $m \in M_{n+2}$ follows by definition, or m is a non-empty multi-set of members of C_n. By IH, $C_n \subseteq C_{n+1}$, hence m is a non-empty multi-set of members of C_{n+1}. By definition of M_{n+2}, it follows that $m \in M_{n+2}$.

Suppose finally that $P \in C_{n+1}$. Then just as before, $\emptyset \subset P \subseteq M_{n+1} \subseteq M_{n+2}$, hence $P \in C_{n+2}$. □

Lemma 11 *The mode-space $\langle M_I, V_I \rangle$ is complete, constrained, and intermediate.*

Proof *Complete*: We show first that $P \sqcup Q$ and $P + Q$ are raisable whenever P and Q are derivative, raisable contents. Since P and Q are derivative contents, $P \in C_n$ for some $n > 0$ and $Q \in C_m$ for some $m > 0$. Now let $j = max(m, n)$. By Lemma 10, P and Q are both in C_j and therefore non-empty subsets of M_j. But then it follows that $P \sqcup Q \in \mathscr{R}$. For suppose $m_1 \in P$ and $m_2 \in Q$, so $m_1, m_2 \in M_j$. Then since every M_n with $n > 0$ is closed under fusion, $m_1 \sqcup m_2 \in M_j$. It follows that $P \sqcup Q \subseteq M_j$, and thus that $\{P \sqcup Q\} \subseteq C_j$. Hence V is defined for $\langle P \sqcup Q \rangle$, and so $P \sqcup Q$ is raisable. Moreover, together with the fact that C_j is closed under union, it also follows from this result that $P + Q$ is raisable. Finally, we show that $\uparrow P$ is raisable if P is. Firstly, $\{V\langle P \rangle\} \in C_{n+1}$, and hence $\{V\langle P \rangle\}$ is raisable. If $P \cap M^D$ is empty, then $\uparrow P = \{V\langle P \rangle\}$, so $\uparrow P$ is raisable. If $P \cap M^D$ is non-empty, then $P \cap M^D \in C_n$, and so $P \cap M^D$ is raisable. Then by the above result, so is $\{V\langle P \rangle\} + (P \cap M^D) = \uparrow P$.

Constrained: Since $V(\gamma)$, when defined, is the multi-set underlying γ, and since γ and δ determine the same multi-set only if they determine the same set, it is immediate that $V(\gamma) = V(\delta)$ only if γ and δ determine the same set.

Intermediate: Since $V(\gamma)$ is the multi-set underlying γ, and $V(\delta)$ is the multi-set underlying δ, if the same multi-set underlies γ and δ, then $V(\gamma) = V(\delta)$. □

The canonical intermediate model of $\mathscr{L}_{\approx}$ is now defined as follows:

Definition 9 The *canonical intermediate model $\mathscr{M}_I$* of $\mathscr{L}_{\approx}$ is $\langle M_I, V_I, [\cdot]_I \rangle$, where for sentences $A, B \in \mathscr{L}_{\approx}$

- $[A]_I = \langle \{A\}, \{\neg A\} \rangle$ if is atomic
- $[\neg A]_I = \neg [A]_I$
- $[A \wedge B]_I = [A]_I \wedge [B]_I$
- $[A \vee B]_I = [A]_I \vee [B]_I$

Equivalences that are true in the canonical intermediate model will also be called canonical.

Note that all contents assigned by $\mathscr{M}_I$ to some sentence are normal. For it is easy to see that all contents assigned to atomic sentences are normal, and so by the result that normality is preserved under truth-functional operations, it follows that all assigned contents are.

We now establish a correlation between the syntactic complexity of the formulas of the propositional language $\mathscr{L}_B$ and the level at which their contents are constructed in the recursive definition of the mode-space.

Definition 10 For all $P \in C$, let $rank(P)$ be the lowest n with $P \in C_n$.

Lemma 12 *For all* $P \in \mathscr{R}$ *and* $Q, R \in \mathscr{R} \cap \mathscr{C}^D$*:*

(i) $rank(Q \sqcup R) = max\{rank(Q), rank(R)\}$
(ii) $rank(Q + R) = max\{rank(Q), rank(R)\}$
(iii) $rank(\uparrow P) = rank(P) + 1$

Proof For (i): Since each M_n is closed under fusion of modes, if $Q, R \in C_n$, then $Q \sqcup R \in C_n$, so the rank of $Q \sqcup R$ cannot be higher than the maximum rank of Q and R. Since each M_n is closed under non-empty submulti-sets, the rank of $Q \sqcup R$ also cannot be lower than the maximum rank of Q and R.

For (ii): In addition to the observations under (i), note that each C_n is closed under unions and non-empty subsets to see that the rank of $Q + R$ can be neither higher nor lower than the maximum rank of Q and R.

For (iii): Assume $rank(P) = n$. Firstly, by construction, $\{V\langle P\rangle\} = \{\lfloor P \rfloor\} \in C_{n+1}$.[58] Moreover, $\{\lfloor P \rfloor\} \notin C_n$. For suppose otherwise. Then $\lfloor P \rfloor \in M_n$. Now if $n = 0$, it follows that $\lfloor P \rfloor$ is a literal in $\mathscr{L}_{\approx}$, which it is not. If $n > 0$, it follows that $P \in C_{n-1}$, contrary to the supposition that $rank(P) = n$. So $rank(\{V\langle P\rangle\}) = n+1$. Now clearly, if $P \cap M^D$ is non-empty, $rank(P \cap M^D) \leq rank(P) = n$, and thus by part (ii), since $\uparrow P = \{V\langle P\rangle\} + (P \cap M^D)$, $rank(\uparrow P) = n + 1$. If $P \cap M^D$ is empty, $\uparrow P = \{V\langle P\rangle\}$, so again $rank(\uparrow P) = n + 1$. □

Definition 11 We define in a simultaneous induction the *positive degree* $pdeg(A)$ and the *negative degree* $ndeg(A)$ of a formula $A \in \mathscr{L}_B$.

- $pdeg(A) = ndeg(A) = 0$ if A is atomic
- $pdeg(\neg A) = ndeg(A)$
- $ndeg(\neg A) = pdeg(A) + 1$
- $pdeg(A \wedge B) = pdeg(A \vee B) = max\{pdeg(A), pdeg(B)\} + 1$
- $ndeg(A \wedge B) = ndeg(A \vee B) = max\{ndeg(A), ndeg(B)\} + 1$

The positive degree of a formula $A \in \mathscr{L}_B$ will sometimes also simply be called A's *degree*, and denoted $deg(A)$. The degree of an *equivalence* $A \approx B$ is $max\{deg(A), deg(B)\}$.[59]

Lemma 13 *In* $\mathscr{M}_I$*, for all* $A \in \mathscr{L}_B$*,* $pdeg(A) = rank([A]_I^+)$ *and* $ndeg(A) = rank([A]_I^-)$.

Proof (For readability, I drop the subscript I.) By induction on the complexity of A. Suppose first that A is atomic. Then $pdeg(A) = ndeg(A) = 0$, and $[A]^+ = \{A\} \in C_0$, so $rank([A]^+) = 0$, and $[A]^- = \{\neg A\} \in C_0$, so $rank([A]^-) = 0$. Suppose now the

58 I write $\lfloor P, Q \ldots \rfloor$ for the multi-set including exactly P, Q, …, each exactly as many times as it is listed.

59 We do this because the positive degree plays a somewhat more central role in the proofs to follow than does the negative degree. The asymmetry mirrors an asymmetry in the role or centrality of the positive content compared to the negative content of a formula. Since and insofar as ground is defined only in terms of positive content, negative content plays a somewhat lesser role. But negative content is still essential for the compositional definition of content, since the positive content of a negation is defined by appeal to the negative content of the negated formula. In an analogous way, the notion of the negative degree of a formula is essential since the positive degree of a negation is defined in terms of the negative degree of the negated formula. Thanks here to a referee for pressing me to clarify this.

thesis holds for A and B (IH). Then it also holds for $A \wedge B$, $A \vee B$, and $\neg A$. I give the proof for $pdeg(A \wedge B)$, the other cases are similar.

$$
\begin{array}{lll}
rank([A \wedge B]^+) & = rank(\uparrow[A]^+ \sqcup \uparrow[B]^+) & \text{by df. } \wedge \\
 & = max\{rank(\uparrow[A]^+), rank(\uparrow[B]^+)\} & \text{by Lemma 12(i)} \\
 & = max\{rank([A]^+) + 1, rank([B]^+) + 1\} & \text{by Lemma 12(iii)} \\
 & = max\{rank([A]^+), rank([B]^+)\} + 1 & \\
 & = max\{pdeg(A), pdeg(B)\} & \text{by IH} \\
 & = pdeg(A \wedge B) & \text{by df. } pdeg
\end{array}
$$

□

For our purposes, the most important bit is the immediate corollary that only equivalences between formulas of equal degree are canonical:

Corollary 14 *For all sentences $A, B \in \mathscr{L}_{\approx}$, if $\mathscr{M}_I \models A \approx B$, then $deg(A) = deg(B)$.*

We are now in a position to prove completeness.

Theorem 15 (Completeness of the intermediate system) $\vdash_{\mathfrak{I}} \varphi$ *whenever* $\models_{\mathfrak{I}} \varphi$.

Proof As indicated earlier, we prove this by showing that every canonical equivalence is derivable. So assume that φ is canonical. Suppose first that $deg(\varphi) = 0$. Then φ has one of these forms, with A and B atomic:

(i) $A \approx B$
(ii) $A \approx \neg B$
(iii) $\neg A \approx B$
(iv) $\neg A \approx \neg B$

Cases (ii) and (iii) cannot obtain, for $[A]^+ = \{A\} \neq \{\neg B\} = [B]^- = [\neg B]^+$, and likewise $[\neg A]^+ = [A]^- = \{\neg A\} \neq \{B\} = [B]^+$. Equivalences of forms (i) and (iv) are canonical only if $A = B$, so they take the forms $A \approx A$ and $\neg A \approx \neg A$, respectively. But equivalences of these forms are derivable by (Reflexivity).

So suppose that equivalences of degree $\leq n$ are derivable if canonical, and suppose φ is of degree $n + 1$. Then φ is an equivalence between two formulas of degree $n + 1$. Each of them can be either a conjunction, or a disjunction, or a negated conjunction or disjunction, or a double negation. However, given (Symmetry) we need not separately consider, say, the case of $A \wedge B \approx C \vee D$ and that of $A \vee B \approx C \wedge D$. Moreover, given the DeMorgan equivalences, we also need not consider equivalences with a negated conjunction or a negated disjunction, since these cases may be reduced using the DeMorgan rules to cases of disjunctions or conjunctions of negations. So the cases we need to consider are these:

(i) $A \wedge B \approx C \wedge D$
(ii) $A \vee B \approx C \vee D$
(iii) $A \wedge B \approx C \vee D$
(iv) $A \wedge B \approx \neg\neg C$
(v) $A \vee B \approx \neg\neg C$

Case (i): By Lemma 6(iv), if $A \wedge B \approx C \wedge D$ is canonical, then so are either (a) $A \approx C$ and $B \approx D$ or (b) $A \approx D$ and $B \approx C$. These equivalences are at most degree n, and so by IH, they are derivable. If (a), then by (Preservation $\wedge$), $A \wedge B \approx C \wedge B$ and $C \wedge B \approx C \wedge D$ are derivable. By (Transitivity), $A \wedge B \approx C \wedge D$ is derivable. If (b), then by (Preservation $\wedge$) we obtain $A \wedge B \approx D \wedge B$ and $B \wedge D \approx C \wedge D$. By (Commutativity $\wedge$) and (Transitivity), we may again derive $A \wedge B \approx C \wedge D$.

Case (ii): By Lemma 6(vi), there are five ways for $A \vee B \approx C \vee D$ to be canonical. The first is that as in case (i), either (a) $A \approx C$ and $B \approx D$ or (b) $A \approx D$ and $B \approx C$ are canonical. By IH, these will be derivable, and similarly as before, using (Preservation $\vee$) and (Commutativity $\vee$) in place of the corresponding rules for conjunctions, we may derive $A \vee B \approx C \vee D$. The other four cases exhibit a common structure, so I shall confine myself to treating one of them, which is that $A \approx C$, $B \leq A$, and $D \leq C$ are canonical. Assuming that this is not also an instance of the first case, it follows that $B \approx D$ is not canonical. We can now show that both $[B]^+ \leq [D]^+$ and $[D]^+ \leq [B]^+$, which entails that $B \approx D$ is canonical, contrary to assumption. For since $[A]^+ = [C]^+$, for any $m \in [C]^+$, $m \sqcup \lfloor [B]^+ \rfloor \in [A \vee B]^+ = [C \vee D]^+$. By construction of the canonical model, no mode in $[C]^+$ contains any content more than finitely many times, so $m \sqcup \lfloor [B]^+ \rfloor$ is always distinct from m. Since $[C]^+$ moreover includes only finitely many modes, $m \sqcup \lfloor [B]^+ \rfloor \notin [C]^+$. It follows that $\lfloor [B]^+ \rfloor \in [D]^+$ and therefore $[B]^+ \leq [D]^+$. Similarly, for any $m \in [A]^+$, $m \sqcup \lfloor [D]^+ \rfloor \in [C \vee D]^+ = [A \vee B]^+$. By analogous reasoning as before, $\lfloor [D]^+ \rfloor \in [B]^+$ and hence $[D]^+ \leq [B]^+$.

The remaining cases (iii)–(v) cannot obtain. For cases (iii) and (iv), it suffices to note that both $[\neg\neg C]^+$ and $[C \vee D]^+$ always include the mode corresponding to the multi-set including only $[C]^+$, and exactly once, whereas every mode in a conjunction corresponds to a multi-set which either contains at least two elements, or contains one element at least twice.

For case (v), by Lemma 6(iii), if $A \vee B \approx \neg\neg C$ is canonical, so is either $A \approx C$ or $B \approx C$. So suppose $[A]^+ = [C]^+$; the other case is analogous. Then whenever $m \in [C]^+$, $[A \vee B]^+$ also includes $m \sqcup \lfloor [B]^+ \rfloor$. As before, no mode in $[C]^+$ contains any content more than finitely many times, so $m \sqcup \lfloor [B]^+ \rfloor$ is always distinct from m, and since $[C]^+$ includes only finitely many modes, $[A \vee B]^+$ and $[C]^+$ are distinct. □

A.4 Adequacy of the extensional system

For φ an equivalence in $\mathscr{L}_{\approx}$, we write $\vdash_{\mathfrak{E}} \varphi$ to say that φ is derivable within $\mathfrak{E}$, and we write $\models_{\mathfrak{E}} \varphi$ to say that φ is valid in the class of models based on complete, constrained, extensional mode-spaces. We show first that $\mathfrak{E}$ is sound with respect to that class of models.

Theorem 16 (Soundness of the extensional system)
For every equivalence $\varphi \in \mathscr{L}_{\approx}$, if $\vdash_{\mathfrak{E}} \varphi$, then $\models_{\mathfrak{E}} \varphi$.

Proof Given the previous soundness result in theorem 8 and the fact that every extensional mode-space is also intermediate, it suffices to establish soundness for the additional rules in $\mathfrak{E}$, i.e. (Collapse $\wedge/\vee$), (Collapse $\vee/\neg\neg$), (Introduction $\leq\wedge$), and (Introduction $\leq\vee$). The soundness of (Collapse $\wedge/\vee$), (Collapse $\vee/\neg\neg$) is immediate from lemma 5(ii)-(iii). The soundness of (Introduction $\leq\wedge$) and (Introduction $\leq\vee$) is straightforward given Lemma 7 and the principles ($<\wedge$) and ($<\vee$) in lemma 1. □

The completeness proof proceeds in close analogy to that for the semi-extensional system. We first define the canonical extensional mode-space, simply replacing any reference to multi-sets in the definition of the canonical semi-extensional mode-space by reference to the corresponding set.

Definition 12 The *canonical extensional mode-space* for $\mathscr{L}_{\approx}$ is the pair $\langle M_E, V_E \rangle$, where

- $M_0 := \{A \in \mathscr{L}_{\approx} : A \text{ is a literal}\}$
- $C_n := \wp(M_n) \backslash \{\emptyset\}$
- $M_{n+1} := \{m : m \in M_0 \text{ or } m \text{ is a non-empty set of members of } C_n\}$
- $C := \bigcup_{n \in N} C_n$
- $M_E := \{m : m \in M_0 \text{ or } m \text{ is a non-empty set of members of } C\}$
- $V_E(\gamma) = \Gamma$ if γ is a non-empty countable sequence of members of C and Γ is the underlying set.
- $V_E(\gamma)$ is undefined otherwise.

By straightforward adjustments to the earlier proof, it may be shown that $\langle M_E, V_E \rangle$ is a mode-space of the desired kind.

Lemma 17 *$\langle M_E, V_E \rangle$ is a complete, constrained, and extensional mode-space.*

The canonical extensional model of $\mathscr{L}_{\approx}$ is defined in the obvious way:

Definition 13 The *canonical extensional model* $\mathscr{M}_E$ of $\mathscr{L}_{\approx}$ is $\langle M_E, V_E, [\cdot]_E \rangle$, where for sentences $A, B \in \mathscr{L}_{\approx}$

- $[A]_E = \langle \{A\}, \{\neg A\} \rangle$ if is atomic
- $[\neg A]_E = \neg [A]_E$
- $[A \wedge B]_E = [A]_E \wedge [B]_E$
- $[A \vee B]_E = [A]_E \vee [B]_E$

The lemmata concerning the correspondence of the degree of syntactic complexity of an $\mathscr{L}_{\approx}$-sentence to the rank in the hierarchy of propositions in the construction of the mode-space unproblematically carries over to the extensional setting, so we again obtain the desired corollary:

Corollary 18 *For all sentences $A, B \in \mathscr{L}_{\approx}$, if $\mathscr{M}_E \models A \approx B$, then $deg(A) = deg(B)$.*

In preparation of the completeness proof, it helps to first prove the following lemma.

Lemma 19 *If sentences A, $B \in \mathscr{L}_B$ are degree $\leq n$, and if every equivalence up to and including degree n is derivable if canonical, then $B \leq A$ is also derivable if canonical.*

Proof Assume the antecedent. By definition of $\leq$, $B \leq A$ is canonical just in case either $B \approx A$ is canonical or $[B]^+ < [A]^+$. By assumption, if $B \approx A$ is canonical, then it is derivable. But then by (Preservation $\vee$), so is $A \vee B \approx A \vee A$, which is $B \leq A$. So suppose that $[B]^+ < [A]^+$. Then $deg(A) > 0$, and so A takes one of these forms

(a) $D \vee E$
(b) $D \wedge E$
(c) $\neg\neg D$
(d) $\neg(D \vee E)$
(e) $\neg(D \wedge E)$

where D and E are degree $< n$.

If (a), and thus $[B]^+ < [D \vee E]^+ = [D]^+ \vee [E]^+$, it follows that either $[B]^+ \leq [D]^+$ or $[B]^+ \leq [E]^+$ and hence that either (i) $B \leq D$ is canonical or (ii) $B \leq E$ is canonical. Since D and E are degree $< n$, $D \vee D$ and $E \vee E$ are degree $\leq n$, so by corollary 18, the equivalences $B \leq D$ and $B \leq E$ are degree $\leq n$. So by assumption, if (i), then $B \leq D$ is derivable, and if (ii), then $B \leq E$ is derivable. Suppose (i). Then by (Introduction $\leq\vee$), $B \leq D \vee E$ is derivable, which is $B \leq A$. Suppose (ii). Then by (Introduction $\leq\vee$), $B \leq E \vee D$ is derivable. Using (Commutativity $\vee$), (Transitivity), and (Preservation $\vee$), we may derive from this $B \leq D \vee E$, that is $B \leq A$.

If (b), and thus $\{[B]^+\} \in [D \wedge E]^+ = [D]^+ \wedge [E]^+$, it follows that $[B]^+ \leq [D]^+$ and $[B]^+ \leq [E]^+$ and hence that both $B \leq D$ and $B \leq E$ are canonical. As before, these equivalences are degree $\leq n$ and thus derivable. By (Introduction $\leq\wedge$), so is $B \leq D \wedge E = A$.

The remaining cases can be reduced to the previous ones using the DeMorgan identities and Lemma 5. For illustration, suppose that case (c) obtains. Then $[B]^+ < [\neg\neg D]^+$. But $[\neg\neg D]^+ = [D \vee D]^+ = [D]^+ \vee [D]^+$. By the reasoning in case (a), $B \leq D \vee D$ is derivable. Using the derivable equivalence of $D \vee D$ to $\neg\neg D$, we may derive $B \leq \neg\neg D$, i.e. $B \leq A$. □

Theorem 20 (Completeness of the Extensional System)
For every equivalence $\varphi \in \mathscr{L}_{\approx}$, if $\models_{\mathfrak{E}} \varphi$, then $\vdash_{\mathfrak{E}} \varphi$.

Proof We show by induction on the degree of equivalences that every canonical equivalence is derivable. Suppose φ is a canonical equivalence. The case of $deg(\varphi) = 0$ is exactly as in the semi-extensional case.

Now assume all canonical equivalences of degree $\leq n$ are derivable and suppose φ is of degree $n + 1$. So φ is an equivalence between two formulas of degree $n + 1$. Each of them can be either a conjunction, or a disjunction, or a negated conjunction or disjunction, or a double negation. The last three cases can be reduced to the first two in the same way we did in the proof of Lemma 19. So we only have three kinds of equivalences of degree $n + 1$ to consider, namely instances of the following forms, where A, B, C, and D are each of some degree $\leq n$:

(i) $A \wedge B \approx C \wedge D$
(ii) $A \vee B \approx C \vee D$
(iii) $A \wedge B \approx C \vee D$

Case (i): By Lemma 6(iv), if $A \wedge B \approx C \wedge D$ is canonical, then so are either (a) both $A \approx C$ and $B \approx D$, or (b) both $A \approx D$ and $B \approx C$. So suppose (a). The equivalences $A \approx C$ and $B \approx D$ are both at most degree n, so by IH, $A \approx C$ and $B \approx D$ are derivable. Using (Preservation $\wedge$), $A \wedge B \approx C \wedge B$ and $C \wedge B \approx C \wedge D$ are derivable. Using (Transitivity), $A \wedge B \approx C \wedge D$ is derivable. Now suppose (b)

$A \approx D$ and $B \approx C$ are canonical. Then these are at most degree n and thus derivable. Using (Preservation $\wedge$), $A \wedge B \approx D \wedge B$ and $B \wedge D \approx C \wedge D$ are derivable. Using (Commutativity $\wedge$) and (Transitivity), $A \wedge B \approx C \wedge D$ is derivable.

Case (ii): By Lemma 6(vi), if $A \vee B \approx C \vee D$ is canonical, there are five ways this can come about. One is that $A \wedge B \approx C \wedge D$ is canonical, in which case as before, either $A \approx C$ and $B \approx D$ are canonical, or $A \approx D$ and $B \approx C$ are canonical. These will then be derivable, and much as in case (i) but using (Commutativity $\vee$) instead of (Commutativity $\wedge$), $A \vee B \approx C \vee D$ is derivable from them. A second way in which $A \vee B \approx C \vee D$ can be canonical is by $A \approx C$, $B \leq A$, and $D \leq A$ being canonical; the remaining cases are analogous and will be omitted. Then by IH and lemma 19, $A \approx C$, $B \leq A$ and $D \leq A$ are all derivable. From these, using mainly (Commutativity $\vee$) and (Preservation $\vee$), we may then derive $A \vee B \approx C \vee D$.

Case (iii): By Lemma 6(v), if $A \wedge B \approx C \vee D$ is canonical, then so is $A \approx B$, which, by IH, is derivable. But then also $[A \wedge B]^+ = [A]^+ \wedge [B]^+ = [A]^+ \wedge [A]^+ = [A]^+ \vee [A]^+ = [A \vee A]^+$, so $A \vee A \approx C \vee D$ is also canonical, and by case (ii) derivable. From these, using mainly (Collapse $\wedge/\vee$) and (Preservation $\wedge$), we may derive $A \wedge B \approx C \vee D$. □

B Comparison of deductive systems

Theorem 21 *For every equivalence $\varphi \in \mathscr{L}_{\approx}$, if $\vdash_{\mathfrak{R}} \varphi$ then $\vdash_{\mathfrak{J}} \varphi$.*

Proof Call a theorem φ of $\mathfrak{R}$ *unproblematic* if the theorem produced by applying the rule (Pres.$\neg$) to φ can also be derived within $\mathfrak{J}$. We show by an induction on the length of derivations that all theorems of $\mathfrak{R}$ are unproblematic. From this it follows straightforwardly that all theorems of $\mathfrak{R}$ are theorems of $\mathfrak{J}$. Consider first the case of a derivation D of length 1. There are three case:

1. D consists in an application of (Reflexivity). Then the result of applying (Preservation $\neg$) can also be achieved simply by an application of (Reflexivity).
2. D consists in an application of (Commutativity $\vee$), so the theorem established by D is $A \vee B \approx B \vee A$. Application of (Preservation $\neg$) yields $\neg(A \vee B) \approx \neg(B \vee A)$. This can be derived within $\mathfrak{J}$ from $A \vee B \approx B \vee A$ by application of the DeMorgan rules and (Commutativity $\wedge$): $(\neg(A \vee B) \approx \neg A \wedge \neg B \approx \neg B \wedge \neg A \approx \neg(B \vee A))$
3. D consists in an application of (Commutativity $\wedge$). Analogous to the previous case.

Suppose then that derivations up to length n produce only unproblematic theorems, and suppose D has length $n + 1$. The cases in which the final step in D consists in the application of one of the premise-less rules just discussed are exactly as before. The remaining cases are five, according as the final step in D is an application of

1. (Symmetry) Then application of (Preservation $\neg$) produces $\neg B \approx \neg A$. By IH, $\neg A \approx \neg B$ can be derived within $\mathfrak{J}$, and thus by (Symmetry), so can $\neg B \approx \neg A$.

2. (Transitivity) Then application of (Preservation $\neg$) produces $\neg A \approx \neg C$. By IH, $\neg A \approx \neg B$ and $\neg B \approx \neg C$ can be derived within $\mathfrak{I}$, and thus by (Transitivity), so can $\neg A \approx \neg C$.
3. (Preservation $\vee$) Then application of (Preservation $\neg$) produces $\neg(A \vee C) \approx \neg(B \vee C)$. By IH, $\neg A \approx \neg B$ can be derived within $\mathfrak{I}$. By (Preservation $\wedge$), $\neg A \wedge \neg C \approx \neg B \wedge \neg C$ can then be derived, and by DeMorgan, so can $\neg(A \vee C) \approx \neg(B \vee C)$.
4. (Preservation $\wedge$) Analogous to the previous case.
5. (Preservation $\neg$) Then application of (Preservation $\neg$) produces $\neg\neg A \approx \neg\neg B$, which can be derived within $\mathfrak{I}$ from $A \approx B$ by (Preservation $\neg\neg$). □

References

Angell, R. B. (1977). Three systems of first-degree entailment. *The Journal of Symbolic Logic*, *47*, 147.
Audi, P. (2012). A clarification and defense of the notion of grounding. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding* (pp. 101–121). Cambridge: Cambridge University Press.
Batchelor, R. (2010). Grounds and consequences. *Grazer Philosophische Studien*, *80*, 65–77.
Bolzano, B. (1837). Wissenschaftslehre. *Scientia*.
Bolzano, B. (2014). *Theory of science*. Oxford: Oxford University Press.
Correia, F. (2010). Grounding and truth-functions. *Logique et Analyse*, *53*(211), 251–279.
Correia, F. (2014). Logical grounds. *The Review of Symbolic Logic*, *7*(1), 31–59.
Correia, F. (2016). On the logic of factual equivalence. *The Review of Symbolic Logic*, *9*(1), 103–122.
Correia, F. (2017). An impure logic of representational grounding. *Journal of Philosophical Logic*, *46*(5), 507–538.
Correia, F. (2018). The logic of relative fundamentality. *Synthese*. https://doi.org/10.1007/s11229-018-1709-8
deRosset, L. (2013). What is weak ground? *Essays in Philosophy*, *14*(1), 7–18.
deRosset, L. (2014). On weak ground. *The Review of Symbolic Logic*, *7*(4), 713–744.
Fine, K. (2010). Some puzzles of ground. *Notre Dame Journal of Formal Logic*, *51*(1), 97–118.
Fine, K. (2012a). Guide to ground. In B. Schnieder & F. Correia (Eds.), *Metaphysical grounding* (pp. 37–80). Cambridge: Cambridge University Press.
Fine, K. (2012b). The pure logic of ground. *The Review of Symbolic Logic*, *5*(1), 1–25.
Fine, K. (2016). Angellic content. *Journal of Philosophical Logic*, *45*(2), 199–226.
Fine, K. (2017a). A theory of truthmaker content I: Conjunction, disjunction and negation. *Journal of Philosophical Logic*, *46*(6), 625–674.
Fine, K. (2017b). A theory of truthmaker content II: Subject-matter, common content, remainder and ground. *Journal of Philosophical Logic*, *46*(6), 675–702.
Krämer, S. (2013). A simpler puzzle of ground. *Thought: A Journal of Philosophy*, *2*(2), 85–89.
Krämer, S. (2018). Towards a theory of ground-theoretic content. *Synthese*, *195*(2), 785–814.
Krämer, S., & Roski, S. (2015). A note on the logic of worldly ground. *Thought: A Journal of Philosophy*, *4*(1), 59–68.
Krämer, S., & Roski, S. (2017). Difference-making grounds. *Philosophical Studies*, *174*(5), 1191–1215.
Litland, J. E. (2016). Pure logic of many-many ground. *Journal of Philosophical Logic*, *45*(5), 531–577.
McDaniel, K. (2015). Propositions: Individuation and invirtuation. *Australasian Journal of Philosophy*, *93*(4), 757–768. https://doi.org/10.1080/00048402.2015.1035291.
Poggiolesi, F. (2016a). A critical overview of the most recent logics of grounding. In F. Boccuni & A. Sereni (Eds.), *Objectivity, realism, and proof : FilMat studies in the philosophy of mathematics* (pp. 291–309). Cham: Springer.
Poggiolesi, F. (2016b). On defining the notion of complete and immediate formal grounding. *Synthese*, *193*(10), 3147–3167.
Poggiolesi, F. (2018). On constructing a logic for the notion of complete and immediate formal grounding. *Synthese*, *195*(3), 1231–1254.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffmann (Eds.), *Modality: Metaphysics, logic, and epistemology* (pp. 109–136). Oxford: Oxford University Press.
Roski, S. (2017). *Bolzano's conception of grounding*. Frankfurt: Klostermann.
Schnieder, B. (2011). A logic for 'because'. *The Review of Symbolic Logic*, *4*(3), 445–465.

Philos Stud (2017) 174:1191–1215
DOI 10.1007/s11098-016-0749-5

Difference-making grounds

Stephan Krämer[1] · Stefan Roski[1]

Published online: 30 July 2016

Abstract We define a notion of difference-making for partial grounds of a fact in rough analogy to existing notions of difference-making for causes of an event. Using orthodox assumptions about ground, we show that it induces a non-trivial division with examples of partial grounds on both sides. We then demonstrate the theoretical fruitfulness of the notion by applying it to the analysis of a certain kind of putative counter-example to the transitivity of ground recently described by Jonathan Schaffer. First, we show that our conceptual apparatus of difference-making enables us to give a much clearer description than Schaffer does of what makes the relevant instances of transitivity appear problematic. Second, we suggest that difference-making is best seen as a mark of good grounding-based explanations rather than a necessary condition on grounding, and argue that this enables us to deal with the counter-example in a satisfactory way. Along the way, we show that Schaffer's own proposal for salvaging a form of transitivity by moving to a contrastive conception of ground is unsuccessful. We conclude by sketching some natural strategies for extending our proposal to a more comprehensive account of grounding-based explanations.

Keywords Grounding · Causation · Explanation · Difference-making · Transitivity

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

Stefan Roski
stefan.roski@gmail.com

[1] Department of Philosophy, Universität Hamburg, Von Melle-Park 6, 20146 Hamburg, Germany

1 Introduction

There is a familiar idea that grounding is in certain respects analogous to causation. In particular, it has been suggested that grounding and causation are alike in that they are both *explanation-backing* relations.[1] For causation, the idea would be roughly as follows. Sometimes, an event *c* causes another *e*. In at least some of these cases, it is correct to say that *e* occurred *because c* occurred. Moreover, the truth of this because-claim is owed at least in part to the causal relationship between *c* and *e*. And at least in some cases, a true because-claim whose truth is thus owed to an instance of causation may be used to give an adequate *causal explanation* of why the event *e* occurred.[2]

For grounding, a parallel idea might run as follows. Sometimes, a fact *f* grounds a fact *g*. In at least some of these cases, it is correct to say that *g* obtains *because f* obtains. Moreover, the truth of this because-claim is owed at least in part to the grounding relationship between *f* and *g*. And at least in some cases, a true because-claim whose truth is thus owed to an instance of grounding may be used to give an adequate *grounding explanation* of why the fact *g* obtains.

We may say that a cause of an event *e* is *causal-explanatorily relevant to e* iff it underlies an adequate causal explanation in the way described. And we may say that a ground of a fact *g* is *grounding-explanatorily relevant to g* iff it underlies an adequate grounding explanation in the way described.

It is sometimes suggested, with respect to causal explanation, that only a certain *elite subclass* of the causes of a given event are causal-explanatorily relevant, and that the members of this subclass may be singled out by a criterion of *difference-making*. Not every causal influence on a given event, however minor and remote, can play a role in explaining why the event occurred, but only those that make a difference to its occurrence.

This idea is central in particular to Michael Strevens' much-discussed theory of causal explanation.[3] A somewhat less controversial, if also less informative version of the claim is that there is a significant distinction between two kinds of causes, naturally described in terms of difference-making, which has some important role to play in the theory of causation and causal explanation. It might be held, for example, that although non-difference-making causes are capable of figuring in explanations *per se*, they will generally not figure in *good* explanations.[4] On this kind of view, the notion of difference-making still seems to be of significant theoretical interest in the study of causal explanation.

[1] A picture roughly like that to be described is proposed by Schaffer (2016) and Audi (2012a, b). See also Schnieder (2015, 2010).

[2] In the debate about causal explanation, a view like this is endorsed, for example, by Strevens (2008), Woodward (2003) and Ruben (2012).

[3] See Strevens (2008), in particular ch. 2. Criteria of difference-making play an important role in the debate on causation in general—sometimes not quite with the same role they have in Strevens; cf. e.g. Lewis (1973, 160f), and List and Menzies (2009). We will briefly come back to this in Sect. 7 below.

[4] This is one way to understand Lewis's view on the matter; cf. Lewis (1986a).

Given the analogy between grounding and causation, a number of questions suggest themselves. Is there a natural way to explicate the notion of difference-making in the application to grounding? If so, is the resulting notion theoretically fruitful? In particular, is it plausible to suppose that it relates to grounding explanations in something like the way in which causal difference-making has been claimed to relate to causal explanation? We argue in this paper that the answer to all three questions is yes.

We begin by introducing a natural abstract characterization of a notion of difference-making for grounds in rough analogy to Strevens' work on causal difference-making (Sect. 2). Next, we establish some basic observations about that notion. Using orthodox assumptions about ground, we show that there are instances of difference-making ground as well as instances of non-difference-making ground, and that difference-making partial grounding is not transitive (Sect. 3). The structure of one of the examples we use to show this is strongly reminiscent of some of the putative counter-examples to the transitivity of grounding described by Schaffer (2012). We show how our conceptual apparatus of difference-making enables us to give a much clearer description than Schaffer can offer of what makes the relevant instances of transitivity appear problematic, thereby establishing a first, significant theoretical payoff of the notion (Sect. 4). We then turn to the relation between difference-making and explanation, and suggest that by taking our notion of difference-making to be a mark of good grounding explanations rather than a necessary condition on grounding, the apparent counter-examples to the transitivity of ground can be dealt with in a satisfactory way (Sect. 5). This marks a second theoretical payoff, which is especially significant because, as we show in Sect. 6, Schaffer's own proposal for avoiding the counter-examples by moving to a contrastive conception of ground is unsuccessful. Finally, we sketch some natural avenues for further developments of our account, for instance to accommodate intuitions about the *proportionality* of explanantia with respect to their explananda (Sect. 7).

2 A notion of difference-making for ground

We begin by making explicit some basic assumptions about grounding that will be in place throughout the paper and that represent orthodoxy in the current debate.[5] We take grounding to be a kind of non-causal priority that is conveyed by certain uses of 'because', 'in virtue of', and cognate phrases that are widespread in philosophical discourse. Standard examples include the claim that a given object is coloured because it is red, or that a given object is red or round because it is round. We shall assume that grounding is a relation obtaining among facts.[6] We can then

[5] For reasons of space, we have to assume a basic familiarity with the notion of grounding. For general introductory overviews see, e.g., Correia and Schnieder (2012) as well as Trogdon (2013).

[6] Some authors prefer to think of ground as expressed by a sentential connective, and deny that there is, strictly speaking, a relation of grounding obtaining between facts. We make this assumption purely for ease of expression; our arguments as well as the idea of grounding as explanation-backing may be easily transposed to the alternative setting. For discussion of the matter, see (Correia and Schnieder 2012, 10ff).

distinguish between the relation of partial ground, which is a binary relation between some fact *f* and another fact *g*, and that of full ground, which is a relation between some plurality of facts $f_1, f_2, \ldots$ and another fact *g*. We sometimes also write $P \prec Q$ to say that the fact that *P* partially grounds the fact that *Q*, and $\Gamma < Q$ to say that the fact that P_1, the fact that P_2, ...together fully ground the fact that *Q*, where $\Gamma = \{$the fact that P_1, the fact that $P_2, \ldots\}$.[7] As is standard, we assume that for a fact to partially ground another is for it to be a member of some full ground of that fact.

Our task is to devise a notion of difference-making that applies to the grounds of a given fact. To see how this might be done, it seems both natural and methodologically advisable to take a cue from the philosophy of science, where the analogous question with respect to the causes of a given event has been the subject of extensive discussion. A particularly helpful suggestion is contained in recent work by Strevens (2008), who describes a general *template* that any account of causal difference-making instantiates:

> All such accounts have a common form. To determine whether a causal influence *c* makes a difference to an explanandum *e*, a comparison is made between two scenarios: the actual scenario, in which *c* is present, and a nonactual scenario in which *c* is not present. The facility with which e occurs in each scenario is evaluated. If it varies, then *c* is classified as a difference-maker. (Strevens 2008, 55)

Indeed, this template seems to us to be general enough to also apply, with minor modifications, in the realm of grounding. We have found the most fruitful way of instantiating Strevens' pattern to be as follows. A scenario is a collection of facts and/or mere states of affairs (the non-obtaining cousins of facts).[8] A fact is considered present in a scenario just in case it is a member of the scenario. Although talk of the facility with which a fact *g* obtains in a scenario sounds like a gradual matter, we shall, in the first instance, interpret it in terms of the binary distinction of whether some part of the scenario is a full ground of *g*. For the time being, the parts of a scenario may here simply be identified with its subsets; depending on subtle details in the theory of ground, there may be reasons also to allow other kinds of scenarios as parts of a scenario.[9]

To determine whether a ground *f* makes a difference to a groundee *g*, we now compare scenarios in which *f* is present with ones in which *f* is not present. A natural idea is to consider any actual scenario *S* in which *f* is present, and compare it with

[7] The symbolism is taken over from Fine (2012b). Note, though, that our understanding of the symbolism differs from Fine's in two respects. First, Fine prefers not to think of ground as a relation between facts and hence does not use the symbolism to abbreviate fact-talk. Second, Fine uses $\prec$ for what he calls strict partial ground, which is defined in terms of his notion of weak ground. The notion we express by $\prec$, which is more common in the current debate, is what Fine calls partial strict ground and writes $\prec^*$. Under Fine's semantics, partial strict ground is strictly stronger than strict partial ground (cf. Fine 2012b, 4).

[8] Although we prefer to allow for scenarios including mere states of affairs, this is not necessary for our purposes, so readers with ontological qualms about such entities need not be concerned.

[9] The matter is taken up in the appendix.

the result $S' = S \backslash \{f\}$ of removing f from S.[10] We then compare the facility with which g obtains in both S and S'. That is, we ask whether g is fully grounded by a subset of S, and whether g is fully grounded by a subset of S'. Since every subset of S' is a subset of S, either the answer is the same with respect to both S and S', or the answer is positive with respect to S and negative with respect to S'. In the latter case, the facility with which g obtains varies between S and S', and we conclude that f is a difference-making ground of g. On the other hand, if such variation does not occur for any scenario S, we conclude that f is a non-difference-making ground of g. We therefore propose the following definition of difference-making partial grounds:[11]

(Df. $\prec_D$) The fact that P is a difference-making partial ground of the fact that Q ($P \prec_D Q$) $=_{\text{df.}}$
for some scenario S which contains a full ground of Q, $S \backslash \{$the fact that $P\}$ does not contain a full ground of Q.

It seems to us that a partial ground that does not satisfy our definition of difference-making could indeed quite naturally be described as not making a difference.[12] For suppose $P \prec Q$, but $P \not\prec_D Q$. Then for *any* full ground Γ of Q that includes the fact that P, removing that fact from Γ results in a scenario *still* containing a full ground of Q. In this way, the fact that P is seen to be at best an *idle wheel* in any scenario fully grounding Q. Another way to make the point is by asking what we can *add* to the fact that P so as to obtain a full ground of the fact that Q. Given that $P \not\prec_D Q$, any collection of facts that will do the job *already* contains a full ground of Q. So the fact that P really does not get us any closer to the fact that Q.[13]

---

[10] We are then not strictly speaking comparing an actual scenario with a non-actual one, as Strevens would have us do. We could do so, however, by *replacing f* with some non-obtaining state of affairs. But since this produces exactly the same results as our simpler procedure, we stick to the latter. Note also that Strevens speaks of *the* actual scenario, whereas we have many, namely any collection of facts. We see no attractive way of amending our procedure to appeal to just one actual scenario. Finally, many people's first idea for cashing out talk of comparison of actual scenarios and non-actual versions of them will be in terms of possible worlds and counterfactuals. But this is by no means mandatory; indeed, Strevens' own account also does not explicate his talk of comparison of actual and nonactual scenarios in terms of counterfactuals and possible worlds (cf. Strevens 2008, 111ff). For our purposes, the present way of cashing Strevens' idea out is much more fruitful. We shall have some use for counterfactuals in considering alternative possibilities later on, though (cf. Sect. 4).

[11] There is an obvious strengthening of this notion of a difference-making partial ground, on which it is required that for *all* full grounds that include the fact that P, removing that fact yields a collection that does not contain a full ground anymore. However, given the common assumption that if $\Gamma < P$ and $\Delta < P$, then $\Gamma \cup \Delta < P$, this would imply that any fact which has several full grounds—which, at least assuming transitivity, is true of the vast majority of facts—has no difference-making partial grounds. So the resulting notion of difference-making would not be very useful.

[12] Note that (Df. $\prec_D$) has the consequence that any difference-making partial ground is a partial ground, as one would have hoped. For suppose S contains a full ground of Q, but $S \backslash \{$the fact that $P\}$ does not. Let $\Gamma \subseteq S$ be such that $\Gamma < Q$. Then Γ is not a subset of $S \backslash \{$the fact that $P\}$. It follows that the fact that P is a member of Γ, and hence that $P \prec Q$.

[13] In Krämer (2016), a similar notion of something getting us closer to the truth of a proposition is employed in formulating a criterion of evidential *relevance*. Given that the notions of relevance and difference-making seem to be very closely related, this may provide some additional motivation for our approach to difference-making, and hints at the possibility of a unified account of difference-making and relevance.

Note that for all we have said so far, there may not actually *be* any instances of non-difference-making partial ground. And indeed, given what we have just said about what such putative grounds would be like, one may have the intuition that these features would precisely disqualify the relevant facts as candidate grounds. As we shall see in the next section, however, given only fairly orthodox, if not uncontroversial assumptions about grounding, it can be shown that examples of both difference-making and non-difference-making partial grounding exist.

3 Difference-making and transitivity

In general, whether one should take there to be instances both of difference-making grounding and of non-difference-making grounding depends on one's preferred theory of ground. However, it would take a highly unorthodox view of ground to deny that there are instances of *difference-making* partial grounding. For take any instance in which a single fact that P is a *full* ground of a fact that Q. Then since removing the fact that P from the collection {the fact that P} leaves an empty set, it is clear that the fact that P is a difference-making partial ground of the fact that Q.[14] For example, it is standardly supposed that a typical true disjunction $P \vee Q$ is grounded by each of its true disjuncts.[15] So if it is a fact that P, then $P < P \vee Q$ and hence $P \prec P \vee Q$. But $\emptyset \not< P \vee Q$, so the fact that P is a difference-making partial ground of the fact that $P \vee Q$. For a second example, take any true conjunction $P \wedge Q$ where the fact that P and the fact that Q are distinct and ground-theoretically independent facts: neither helps ground the other in any way. Then on any standard view of the logic of ground, $P \prec P \wedge Q$ because $P, Q < P \wedge Q$. But $Q \not< P \wedge Q$, for otherwise it follows that the fact that Q is, or grounds, the fact that P. So the fact that P is a difference-making partial ground of the fact that $P \wedge Q$.

It would be easier to deny that there are instances of non-difference-making partial ground. Nevertheless, there are quite strong reasons for thinking that there are such instances, and indeed that there is a *systematic* way of producing examples of non-difference-making partial ground. For given any two suitably[16] independent facts that P and that Q, plausible assumptions concerning the logic of ground allow us to argue that

[14] Fine (2012a, 47f) suggests that some facts may be *zero-grounded*, which is supposed to amount to being grounded by the empty set of facts, and distinguished from being ungrounded. If so, then it may be that not every case in which a single fact fully grounds another gives rise to a case of difference-making partial grounding. However, zero-grounding, if there is such a thing, is supposed to be a feature of only a rather special and rare sort of fact, so our general point is not threatened. For simplicity, we tacitly exclude the possibility of zero-grounding in what follows.

[15] When we speak of standard assumptions in the (propositional) logic of ground, we mean assumptions that are explicitly endorsed in both Correia (2010) and Fine (2012a), which are so far the only reasonably developed systems for the propositional logic of ground. We give more precise versions of our informal arguments here by reference to these systems in the appendix.

[16] The exact condition of independence required varies slightly with the details of the logic of ground assumed. Roughy speaking, it is sufficient to choose facts with disparate subject matters, such as that this ball is red and that that chair is brown. Details are given in the appendix.

(1) $P \prec Q \vee (P \wedge Q)$
(2) $P \nprec_D Q \vee (P \wedge Q)$

For (1), note that as before, $P \prec P \wedge Q$, and by the disjunction principle, $P \wedge Q < Q \vee (P \wedge Q)$. Using the following, relatively weak form of the claim that grounding is *transitive*, we obtain (1).

(T_1) If $P \prec Q$ and $Q < R$, then $P \prec R$

An informal argument for (2) may now be given as follows. Let $\Gamma \cup \{$the fact that $P\} < Q \vee (P \wedge Q)$. Note that $\Gamma \cup \{$the fact that $P\}$ must 'contain enough' to ground Q. Since the fact that P, by the assumption of ground-theoretic independence, cannot be of help in this, $\Gamma \backslash \{$the fact that $P\}$ must contain a scenario that fully grounds Q. But by the disjunction principle and the transitivity principle

(T_2) If $\Gamma < Q$ and $Q < R$, then $\Gamma < R$

it follows that this scenario will then be a full ground of $Q \vee (P \wedge Q)$. In this way, the fact that P may be shown to be an idle member of any full ground of $Q \vee (P \wedge Q)$, and thus to be a non-difference-making partial ground of $Q \vee (P \wedge Q)$.

We would now like to describe a concrete example of non-difference-making partial ground, which is roughly analogous in structure to the kind of case just discussed, and which is moreover strongly reminiscent of a case described by Jonathan Schaffer (2012, 126).[17] Consider a ball b which is everywhere red, except for one tiny green spot s, constituting less than 1 % of b's surface area. Call the exact overall distribution of colours over its surface area d. We also appeal to a property of being largely-red, which we understand to be a disjunctive colour distribution property whose disjuncts are those partial or complete exact distributions of colours over the surface area of b which render at least 99 % of b's surface red. Note that d is one of those disjuncts. Then

(4) b is green in spot s $\prec$ b has colour distribution d.
(5) b has colour distribution d $<$ b is largely-red.
(6) b is green in spot s $\prec$ b is largely-red.

(4) is plausible since part of what it takes for b to have d is for b to be green in spot s; the facts that b is green in spot s and that b is red everywhere else jointly fully ground the fact that b has colour distribution d. This is roughly analogous to the grounding of a conjunction by its conjuncts. (5) is plausible since one way to be largely-red is to have colour distribution d. This is in effect an instance of a disjunction being grounded by its true disjuncts. (6) follows from (4) and (5) by the transitivity principle (T_1).

Assuming that (6) is indeed true, it appears to be an instance of non-difference-making partial ground. To see this, we may ask what must be added to the fact that b is green in spot s to obtain a full ground of the fact that b is largely-red. Anything

[17] We have in mind Schaffer's case of the dented sphere. His particular example suffers from some special problems not affecting the example we describe below, which is why we prefer to focus on our case. Schaffer describes two more putative counter-examples to transitivity, which we do not discuss here. For criticism of the examples, see e.g. Litland (2013).

of that sort will have to ensure on its own that 99 % of the surface area of *b* is red, and thus will on its own constitute a full ground of the fact that *b* is largely-red. (It may help to imagine you are painting the ball. If you start with painting the spot *s* green, then the amount of red colour you have to apply after that in order to make the ball largely-red would have been sufficient to ensure that *b* is largely-red even if you had not already painted *s* green.)

In both of our examples, the premises to the transitivity-based inference are examples of difference-making partial ground. In the first, logical example, the premises were that $P \prec P \wedge Q$, which was already shown to be an instance of difference-making, and that $Q < Q \vee (P \wedge Q)$, which is a case of full ground, and therefore automatically of difference-making ground. In the second example, consider first premise (4). We observed above that the facts that *b* is green in spot *s* and that *b* is red everywhere else jointly fully ground the fact that *b* has colour distribution *d*. But clearly, the fact that *b* is red everywhere else is not on its own a full ground of the fact that *b* has *d*. Premise (5) is again an example of full ground. As a result, both of our examples show that on the proposed notion of difference-making, the relation of difference-making partial ground is *not transitive*.

It might thus be suggested that rather than taking the examples to establish the existence of non-difference-making partial ground, we should take them to establish the *non-transitivity* of partial ground. And indeed, Schaffer offers his analogous case as a counter-example to the transitivity of partial ground (cf. Schaffer 2012, p. 127f). And it must be admitted that especially in the second example, the because-claim corresponding to the putative non-difference-making instance of grounding sounds at best suspect:

(7) *b* is largely-red partly because *b* is green in spot *s*.

So it might be claimed that we have actually just rediscovered Schaffer's reasons for doubting the transitivity of ground, and perhaps identified an interesting structural feature of partial grounds, namely that they are all difference-making in our sense (and hence that (3) and (6) are not cases of grounding, after all).

We propose a way to defend the transitivity of ground, and thus the existence of non-difference-making grounds, against this objection in Sect. 5 below. But before that, we want to highlight a way in which our understanding of difference-making in application to grounds is theoretically fruitful which does not depend on transitivity.

4 Difference-making and subversive double agents

Consider again our second potential example of a non-difference-making partial ground:

(6) *b* is green in spot $s \prec b$ is largely-red

It should be conceded on all sides that this is at best quite a peculiar and somehow second-rate instance of grounding, and prior to deciding whether to count it a case of grounding, we should get as clear as possible about what *makes* the case so peculiar. Only then will we have a clear understanding of the reasons we may have for asserting or denying (6).

Schaffer offers a helpful account of the peculiarity of his analogue for (6) (cf. Schaffer 2012, p. 127). Transposed to our case, it reads as follows:

> (i) The presence of the green spot makes no difference to the largely-redness of the ball. (ii) The ball would be largely-red either way. (iii) The presence of the green spot in no way helps support the largely-redness of the ball, but (iv) is if anything a *threat* to the largely-redness of the ball. (v) The ball is largely-red *despite* the green spot, not because of it.

Remark (i) is, we agree, a very natural thing to say. However, it would be good to have an explication of just how we are to understand the talk of difference-making here. Perhaps we may read Schaffer as offering (ii) as an explication of (i).[18] However, it is clear that the kind of counterfactual dependence at issue in (ii) is not in general a requirement for grounding. For instance, the fact that snow is white grounds the fact that snow is white or grass is green. But it is not the case that if snow were not white, grass would not be green, and so the disjunctive fact that snow is white or grass is green would obtain either way – whether snow is white or not.

This particular failure of counterfactual dependence is akin to failures of counterfactual dependence for causation in cases of causal overdetermination. Indeed, we might say that the disjunctive fact that snow is white or grass is green is ground-theoretically overdetermined since both constituent disjuncts obtain. Since (6) and the grounding relationships by which it is mediated do not seem to involve overdetermination, one might think that they also should not give rise to a failure of counterfactual dependence.[19]

But there are many clear cases of grounding without counterfactual dependence in which no overdetermination is involved. Thus, if it is true that P, then $P < P \vee \neg P$. Still, it is not true that if P were false, it would be false that $P \vee \neg P$. At the same time, there is clearly no overdetermination involved here. We need not even choose a logical truth as the groundee. Suppose John always comes by bus or by bike. Suppose today he came by bus, but if he had not come by bus, he would have come by bike. Then John came by bus < John came by bus or by bike. But there is neither counterfactual dependence nor overdetermination involved.

It might still be, of course, that there is a range of special cases in which grounding can only plausibly be asserted provided that counterfactual dependence holds. But it is hard to see why (6) should belong to such a special subclass. For the

[18] Although this reading is not mandatory, it seems quite natural. As we have already mentioned, it is quite common to explicate difference-making in terms of counterfactual dependence in something like this way. (For a classic statement of this intuition in the case of causal difference-making, see Lewis (1986b, 161–162).) More or less every way of understanding the counterfactual in (ii) gives quite a plausible claim, moreover. For example, it is true that if it were not the case that the ball is green in spot *s*, the ball would still be largely-red.

[19] Schaffer's remark in his (2016, 31) that in the case of a disjunction being grounded by its true disjuncts 'one loses counterfactual dependence due to grounding overdetermination' is suggestive of the idea that failures of counterfactual dependence, at least with respect to the grounding of disjunctions, always result from grounding overdetermination.

putative grounding relationship is mediated exactly by the kind of disjunct-disjunction relationship, which appears to be a source of failures of counterfactual dependence. We conclude that (ii) is of highly questionable relevance to the matter of the plausibility of (6). Consequently, if we are to make anything much of the observation (i) of the lack of difference-making, this notion will have to be understood in some other way.

We would like to suggest that (i) is best understood in terms of our conception of difference-making. Of course, we cannot interpret (i) as the claim that (6) is an instance of non-difference-making *grounding*. This would require that (6) is true, which is precisely what is at issue. But it is easy to find a suitably neutral version of this claim. We begin by taking the *transitive closure* of ground and write $<_T$ and $\prec_T$ for the resulting relations; informally, we shall speak of *t*-grounds.[20] If grounding is already transitive, but only then, $<_T$ and $\prec_T$ will coincide with $<$ and $\prec$. Now given the assumptions (4) and (5), anyone can agree to the following variant of (6):

(6_T) b is green in spot $s \prec_T b$ is largely-red

We may now define a notion of difference-making for $\prec_T$ in exact parallel to our previous definition:

(Df. $\prec_{TD}$) The fact that P is a difference-making partial t-ground of the fact that Q ($P \prec_{TD} Q$) $=_{\text{df.}}$
for some scenario S which contains a full t-ground of Q, $S \backslash \{$the fact that $P\}$ does not contain a full t-ground of Q.

(6_T) then turns out to be an example of non-difference-making partial t-grounding. We would like to suggest that this observation nicely captures the intuitive peculiarity of (6) aimed at in Schaffer's remark (i).

Let us turn to remark (iii), that the presence of the green spot in no way helps support the largely-redness of the ball. This again seems a very natural thing to say, and to bring out a further important feature of the case. Nevertheless, it stands in need of explication. In particular, talk of something's helping support a fact is sometimes simply tantamount to an assertion of partial ground. But then (iii) would amount to nothing more than a flat-out denial of (6) rather than give a reason for such a denial. Using our conceptual apparatus, we can give a different, non-question-begging interpretation of (iii). For as we have seen, (6_T) reports a non-difference-making instance of t-grounding: the fact that b is green in spot s is an idle wheel in any full t-ground of the fact that b is largely-red. This, we claim, captures a good sense in which the presence of the green spot in no way helps support the largely-redness of b, and it does not beg the question of whether (6) is true or false.

The remark (iv) that the green spot is if anything a *threat* to the largely-redness of the ball b highlights another very important and striking peculiarity of (6). Once more, it cries out for explication. Again using the fact that the presence of the green spot is, if not a ground, still a t-ground, we propose to explicate it in terms of the t-ground-theoretic notions related to difference-making. First, call a partial t-ground

[20] More precisely, we let $<_T$ be the closure of $<$ under the principle Cut: If $\Gamma < P$ and $P, \Delta < Q$, then $\Gamma, \Delta < Q$. We then let $\prec_T$ be the partial cousin of $<_T$: $P \prec_T Q$ iff $\Gamma <_T Q$ for some Γ with $P \in \Gamma$.

of the fact that Q a *double-agent* iff it would have been a partial t-ground of the fact that $\neg Q$, if Q had not been the case:[21,22]

(Df. double agent) P is a *double agent* wrt $Q =_{\text{df.}} P \prec_T Q$, and $\neg Q \mathrel{\Box\!\!\rightarrow} (P \prec_T \neg Q)$

If b had not been largely-red, the fact that it is green in spot s would have been a partial t-ground of the fact that b is not largely-red. So in its role as a t-ground, it pulls equally, as it were, on the side of b's being largely-red and the side of b's not being largely-red, simply coming down on whichever side obtains.

It seems, however, that the fact actually pulls more strongly on the side of b's not being largely-red. We can capture this idea by combining the notion of a double agent with that of difference-making. For if b had not been largely-red, the fact that b is green in spot s would have been a *difference-making* partial t-ground of the fact that b is not largely-red. In such a case, we shall call a double agent *subversive*.[23]

(Df. subversive) P is a *subversive* double agent wrt $Q, =_{\text{df.}}$ P is a double agent wrt Q, $P \not\prec_{TD} Q$, and $\neg Q \mathrel{\Box\!\!\rightarrow} (P \prec_{TD} \neg Q)$

The observation that the fact that b is green in spot s is a subversive double agent wrt the fact that b is largely-red also seems to nicely capture Schaffer's final remark that the ball is largely-red *despite* the green spot, not because of it; or at least, they capture those parts of the remark, which are neutral over the truth of (6). We thus see that our results of the previous section, recast in a form that is neutral over the transitivity of $\prec$, enables us to give a clear, non-question-begging description of what is so strange about (6) as a putative case of ground. This marks a first serious theoretical payoff of our notion of difference-making, and its strengthening in the notion of a subversive double agent.

In the next section we indicate a further payoff of the conceptual apparatus we have developed by using it to accommodate the counter-intuitive ring of (6) and its kin in a way consistent with the transitivity of ground.

[21] An alternative way to capture the two-faced nature of the relevant kind of ground is by appeal to a suitable *non-factive* understanding of ground (cf. Fine 2012a, 48ff). Writing $\prec_{T0}$ for non-factive partial t-ground, we would then count a fact P a double agent wrt Q iff $P \prec_T Q$ and $P \prec_{T0} \neg Q$. This option promises to yield more satisfactory results in the case of necessarily obtaining groundees Q, for which the present proposal counts any partial ground a double agent, assuming the orthodox view that counterfactuals with impossible antecedents are vacuously true. However, since the existence of a clear non-factive understanding of ground may be doubted and since the counterfactual serves our present purposes well enough, we here stick to the counterfactual version.

[22] Again, the logic of ground would allow us to systematically produce examples of double agents given suitably independent P, Q. For we then have $P \prec_T (P \wedge Q) \vee (\neg P \wedge Q)$. Now if the groundee had been false, Q would have been false, and we may assume that P would still have been true. But then we would get that $P \prec_T \neg(\neg P \wedge Q)$, and since $\neg(\neg P \wedge Q) \prec_T \neg((P \wedge Q) \vee (\neg P \wedge Q))$, by transitivity, $P \prec_T \neg((P \wedge Q) \vee (\neg P \wedge Q))$, as required for P's being a double agent.

[23] Given suitably independent P, Q, a logical example is given by the true four-way disjunction $R := (P \wedge Q) \vee (\neg P \wedge Q) \vee (\neg P \wedge \neg Q) \vee Q$. Crucially, if R were false, P would be required to render false the third disjunct of R, and thereby would come out a difference-making ground of $\neg R$.

5 Difference-making and explanatory relevance

We have observed above that if ground is transitive, then

(6) *b* is green in spot $s \prec b$ is largely-red

is a true claim of partial grounding, and one in which the partial ground mentioned fails to be a difference-maker with respect to the fact being grounded. And we have suggested in the previous section that at least a large part of the intuitive reasons for rejecting (6) is captured by the observation that the putative partial ground is a *subversive double agent* with respect to the fact being grounded.

We still face the question whether (6) is true or false, and relatedly, whether grounding is transitive or not. We know of no *decisive* considerations with respect to either question, but we incline towards the orthodox view that grounding is indeed transitive and that (6) is accordingly true. Our aim in this section is to use the conceptual apparatus we have put into place thus far to show this view to be defensible.

The general strategy we shall employ may readily be guessed from our introductory remarks at the beginning of this paper. As we have mentioned, it is a familiar idea that a criterion of difference-making for causes may serve to single out a certain elite subclass of the causes of a given event as causal-explanatorily relevant. We suggest that the criterion of difference-making for (t-)grounds should be accorded an analogous role: it serves not to separate mere t-grounds from real grounds, but rather to separate mere grounds from grounding-explanatorily relevant grounds.

In defence of this suggestion, we will use certain general observations about explanations to argue that one normally cannot adequately explain why a given fact obtains by citing only non-difference-making t-grounds—and especially t-grounds which are subversive double agents with respect to the fact to be explained. This result then allows us to explain the odd ring of (6) without conceding (6) to be false. It can be explained as resulting from a general tendency, firstly, to read talk of grounding in terms of 'because', and secondly, to evaluate because-claims as attempts at an explanation, causing us to mistake constraints on good explanations for necessary conditions on grounding.

Following Bromberger (1965) and more recently Schnieder (2015, 183f), we urge that a sharp distinction be made between *explaining why* something is the case and *telling why* something is the case. If it is true that *P* because *Q*, then to successfully tell someone why *P*, it is sufficient to inform the person that *P* because *Q*. But to successfully *explain* to the person why *P*, it is also required that the information one provides be sufficient to enable the person to *understand* why *P*, to solve the *epistemic predicament* that gave rise to their *need* for an explanation why *P*.

The act of explaining why something is the case is thus always aimed at resolving some sort of (real or supposed) epistemic predicament. The ability of a ground to figure in a successful grounding explanation is therefore dependent upon its capability to help remove the relevant epistemic predicament. It seems highly plausible that, except perhaps in some exceptional circumstances, only difference-making grounds have this capability. For the special ground-theoretic profiles of non-difference-making grounds and especially subversive double agents make it very hard to see how such grounds could remove any sort of relevant epistemic predicament. The distinctive feature of a non-difference-making ground *f* is that it does not bring us closer to a full ground of the relevant groundee *g*: anything we could add to *f* to obtain a full ground of *g* already contains such a ground. In light of this, it does not seem as though there could be any sort of puzzlement over how it comes about that the fact *g* obtains which is even partially resolved by the mention of fact *f*. Assuming *f* to be a subversive double agent makes the point even more dramatic. For what puzzlement over the question what makes a given fact that *P* obtain could possibly be removed by the mention of a fact that pulls more strongly on the side of $\neg P$ than it does on the side of *P*? In any ordinary context, the information that some ball has a tiny green spot will not remove the epistemic predicament of someone asking why it is almost everywhere red—if anything it will make it worse. We conclude that difference-making, and a fortiori failure to be a subversive double agent, is a plausible minimal requirement on the grounding-explanatory relevance of a ground.[24,25]

These considerations provide clear grounds for rejecting (or criticizing) grounding explanations given by the assertion of

(7) *b* is largely-red because *b* is green in spot *s*

without requiring us to view (7) or (6) as expressing falsehoods. The odd ring of (6) and (7) may accordingly be explained in a way consistent with their truth, and thus with the transitivity of ground, albeit not with the transitivity of grounding-explanation. Note, moreover, that there is precedent for our view in the debate on causal explanation, where the transitivity of causal explanatory relevance is frequently rejected independently of the transitivity of causation.[26]

[24] Note that this is our preferred view, but it is not fully mandatory for the arguments in this section. In particular, it would suffice if difference-making was merely a minimal condition for *good* explanations. The odd ring of the because-claims under discussion could then be blamed on their expressing a particularly bad explanation due to being utterly uninformative and misleading.

[25] It might be objected that many grounding claims sound plausible even though they, too, are badly suited to remove any relevant epistemic predicament. For instance, it does not sound implausible—or at least not as implausible as (6)—to say that $P, Q < P \wedge Q$. But would anyone who is puzzled over why $P \wedge Q$ obtains be helped by pointing out that this is because *P* and *Q* obtain? In response, we wish to highlight a significant disanalogy. The problem in this last case is one of triviality. Any typical epistemic predicament with respect to $P \wedge Q$ will extend to *P* and/or *Q*, so the envisaged explanation appeals to facts for which the hearer is also in need of an explanation. The problem in our case of non-difference-making grounding, however, is of a quite different structure. For partially explaining the ball's being largely-red by citing its green spot is bad *even if* the hearer is not also puzzled over the green spot. Here the problem with the tie appealed to in the explanation is not that it is obvious and thus uninformative, but rather that it has the wrong strength and/or direction, as it were.

[26] See Owens (1992, 16ff), Heslow (1981).

6 Against Schaffer's contrastive solution

Schaffer proposes a different way to salvage the transitivity of grounding against the apparent counter-example (6). His proposal can be seen as involving two parts, which we describe in turn. The first part is an account of grounding as *contrastive*, coupled with an adjusted notion of transitivity that is applicable to the contrastive conception. On this account, claims of partial grounding take the following canonical form

($\prec_{\text{contrast}}$) The fact that P rather than P^* grounds the fact that Q rather than Q^*.
where

> [t]he fact that [P] and the fact that [Q] are required to be obtaining facts, but the fact that [P^*] is required to be a non-obtaining alternative to the fact that [P], and the fact that [Q^*] is required to be a non-obtaining alternative to the fact that [Q]. (Schaffer 2012, 130)

Roughly speaking, what is grounded is never simply a given fact, but that this fact, rather than some given alternative, obtains. And likewise what grounds something is never simply a given fact, but that this fact, rather than some given alternative, obtains. One might therefore say that the relata of ground are not facts but *differences* – namely, the differences between the relevant obtaining facts and their non-obtaining alternatives. Given this picture of grounding as relating differences, the obvious interpretation of the claim that grounding is transitive is (cf. Schaffer 2012, 132):[27]

(T_{contrast}) If the fact that P rather than P^* grounds the fact that Q rather than Q^*, and the fact that Q rather than Q^* grounds the fact that R rather than R^*, then the fact that P rather than P^* grounds the fact that R rather than R^*

The second part of Schaffer's proposal is the claim that problem cases such as that of the largely-red ball, once recast in explicitly contrastive form, no longer pose a threat to the transitivity of ground as explicated in (T-contrast). For to bring the case into contrastive form, we need to choose suitable alternatives to the three facts involved. And the idea is that however we choose alternatives, either the premises to the transitivity-based inference turn out false, or the conclusion turns out acceptable (cf. Schaffer 2012, 136f).

We may illustrate the difficulties arising for the choice of alternatives as follows. Consider first the fact that our ball b is green in spot s. A natural alternative to this fact is that the ball is red in spot s. So let us consider as the first relatum in our grounding chain the difference between b being green in spot s and b being red in spot s. Next, we need to choose an alternative to the second fact involved, viz. the fact that b has colour distribution d. The difference between this fact and the alternative to be chosen must be grounded by the previous difference. The obvious choice is then that b is red all over. So our first contrastive grounding claim reads:

[27] Note that this picture also yields a very strong and very literal connection between grounding and difference-making, in that it portrays ground simply as the making of one difference by others.

(4*) The fact that *b* is green rather than red in spot *s* grounds the fact that *b* has colour distribution *d* rather than being red all over.

Finally, we need to choose an alternative to the third fact involved, viz. that *b* is largely-red. But now we are in trouble, for both *b* having colour distribution *d* and *b* being red all over are ways for *b* to be largely-red. So that *b* is one rather than the other cannot plausibly be taken to ground that *b* is largely-red rather than something else.

(5*) The fact that *b* has colour distribution *d* rather than being red all over grounds the fact that *b* is largely-red rather than ???

So if we choose our first two alternatives in the way we did, we do not obtain a threat to transitivity.

Unfortunately, there are other possible choices of alternatives which seem much more troublesome. Call d' a possible colour distribution for *b* which is like *d* except in that it has *b* red in spot *s*, but green in half of its total surface area. Then:

(5′) The fact that *b* has colour distribution *d* rather than d' grounds the fact that *b* is largely-red rather than half green.

Leaving the first choice of alternative as it was, we obtain the following contrastive variant on (4):

(4′) The fact that *b* is green rather than red in spot *s* grounds the fact that *b* has colour distribution *d* rather than d'.

This is plausible. Consider the overall difference between the colour distributions *d* and d'. Part of that difference is that *d* has *b* green in spot *s*, whereas d' has *b* red in spot *s*. And *that* part of the overall difference is fully accounted for by the fact that *b* is green rather than red in spot *s*. Therefore the fact that *b* is green rather than red partially grounds the fact that *b* has colour distribution *d* rather than d'.

We now have a pair of contrastive grounding claims, namely (4′) and (5′), to which the contrastive transitivity principle ($T_{contrast}$) applies. We then obtain

(6′) The fact that *b* is green rather than red in spot *s* grounds the fact that *b* is largely-red rather than half green.

But if anything, this sounds even worse than (6). So moving to a contrastive conception of grounding does nothing to alleviate the challenge to transitivity from cases like the largely-red ball.

Of course, Schaffer could try to devise some criterion that prohibits the choice of contrasts we have used to obtain (6′). It is, however, far from obvious how this might be done. Indeed, Schaffer himself readily admits that he is in no possession of objective criteria for the selection of contrasts, and even exhibits a slight pessimism as to whether such criteria may be found at all.[28] Meanwhile, given that Schaffer himself favours a view according to which grounding is a relation that backs explanation, our way to explain the odd ring of (6) and (7) while retaining the transitivity of ground should certainly not be inherently objectionable to him.

[28] Cf. Schaffer (2016, 68).

7 Going further: proportionality and cohesion

There are various directions in which the present account of difference-making for grounding and its relation to grounding explanation may be further developed. We'd like to conclude this paper by tentatively exploring one of them.

Note that the notion of difference-making we have proposed only allows for two choices with respect to the partial grounds of a given fact: either a given partial ground makes a difference to the obtaining of the fact, or it does not. Yet, with the notions of being a double-agent and a subversive double-agent, we are in the position to also make rough *qualitative* distinctions within the realm of non-difference-making partial grounds. Double-agents and subversive double-agents can be seen as two increasingly 'severe' forms of non-difference-making, which, as we have suggested, may in turn correspond to increasingly severe forms of explanatory irrelevance. In order to account for these forms of non-difference-making, we have used counterfactuals (though there might be a more elegant solution using a non-factive understanding of grounding). A natural question is whether an appeal to counterfactuals would also allow us to make helpful qualitative distinctions within the realm of difference-making grounds that in turn may afford useful distinctions with respect to the explanatory relevance of difference-making partial grounds.

A very straightforward idea that suggests itself is to combine our notion of difference-making with a certain form of counterfactual dependence. The idea is best illustrated by an example. Consider the following case. Call an object *signal-coloured* iff it is either red, yellow, or green. Suppose that some street-sign s has some fully determinate shade of red, C. Now, that s has C fully grounds s's having the determinable colour red. The latter, in turn is a full ground of the fact that s is signal-coloured. Since every full ground is a difference-maker, it follows by transitivity that

(8) s has $C \prec_{TD}$ s is signal-coloured, and
(9) s is red $\prec_{TD}$ s is signal-coloured.

Arguably, (8) and (9) differ in an important respect. Roughly speaking, even though the fact that s is C makes a difference to s's being signal-coloured in our sense, it does not seem to be the most relevant bit of information in that respect. For, if s would not have this *particular* shade of red, but a slightly different one, it would still be signal-coloured (a bit more technical: in the closest non-C worlds, s will still have some determinate shade of red). On the other hand, if s would not have been red, but some other colour, it might not have been signal-coloured (a bit more technically: among the closest non-red worlds, there are some worlds in which s is beige, blue, etc.). We can capture this difference between (8) and (9) by calling the former a case of *weak* difference-making and the latter a case of *strong* difference-making. Strong difference-making (in symbols $\prec^{!}_{TD}$) can be defined as follows:

(Df. Strong) $P \prec^{!}_{TD} Q =_{\text{df.}} P \prec_{TD} Q$ and $\neg(\neg P \Box\!\!\rightarrow Q)$.

The distinction between weak and strong difference-makers can naturally be related to the notion of explanatory relevance. At least in many cases, weak difference-makers seem to be too specific to satisfactorily explain why a given fact obtains.

That the sign *s* is signal-coloured is not explained by the fact that it has *this particular shade* of red, but rather by the fact that it has some shade *of red*. Since any other shade of red among the array of different specific shades of red incompatible with *C* would have done as well, information about the specific shade seems explanatorily irrelevant with respect to the fact that *s* is signal-coloured.[29]

Notably, our case instantiates a pattern that is widely discussed in the debate on causal difference-making and causal explanation.[30] In particular, on Strevens' account of causal-explanatory relevance, it is a requirement for a causal influence on a given event to be explanatorily relevant that it is characterized in a sufficiently unspecific way. To pick an illustrative example from Strevens (2008, 96), suppose that a given cannonball is shot at a window and shatters it. The specific weight of the ball, say 10.208 kg, is among the causal influences that bring about the window's shattering. However, thus Strevens', the ball's *specific* weight is explanatorily irrelevant with respect to the window's shattering. All else being equal, *any* ball with a weight of more than 5kg would have broken the window. Hence, thus his suggestion, what makes a difference that is explanatorily relevant to the window's shattering is not the ball's having some determinate mass, but rather its having the determinable property of weighing more than 5kg.[31]

Just as in Strevens' causal case, the strong difference-maker in our case seems to have an explanatory advantage over the weak one because it instantiates a more general pattern and affords, in that sense, a greater unification than weak difference-makers. With respect to the question of why *s* is signal-coloured, the answer that *s* is red is invariant over a variety of mutually incompatible determinate shades of red.

It may be, however, that the aforementioned answer is not the most general answer one can give. The reason is that the class of strong difference-makers may not be exhausted by *s*'s being red; *s*'s being either red or green may also count as a strong difference-maker for *s*'s being signal-coloured. Clearly, the closest worlds in which *s* is neither red nor green include worlds in which *s* is not signal-coloured. So assuming that the fact that *s* is either red or green is a ground of the fact that *s* is signal-coloured, it is also strong difference-maker.[32] It seems, however, that an explanation of *s*'s being signal-coloured in terms of its being either red or green is also unsatisfactory, or at least less satisfactory than the one in terms of *s*'s being red. Hence, being a strong difference-maker may only be a necessary condition for being explanatorily relevant.

[29] We can, at this point, stay agnostic with respect to the issue whether being a strong difference-maker merely determines pragmatic acceptability of certain because-claims, or whether it captures an objective criterion of explanatory relevance.

[30] See e.g. Strevens (2008, §3.5), Sartorio (2005, 75), and cf. Yablo (1992, §4).

[31] Of course the specific weight *is* causally responsible for specific features of the window's breaking (for *how* it breaks exactly) in a way the determinable weight is not. Strevens assumes, however, a scenario wherein the explanandum in question is not the window's breaking-in-a-highly-specific-way, but rather its simply breaking. This seems plausible: in any ordinary sense of 'explanation', there is a wide array of explanations of the occurrence of some given event, where highly specific details of how the event occurred are simply beside the question. Compare on this also Schaffer (2012, 135).

[32] Whether the ground-theoretic assumption holds depends on some subtle details of one's theory of ground and of how the property of being signal-coloured is conceived. For the sake of argument, however, we grant the ground-theoretic assumption to our opponent.

So, one may wish to hold that while weak difference-makers, like *s*'s specific colour, may be too informative, some strong difference-makers, like the disjunctive property of having one of the two colours red and green, seem too uninformative.[33] The natural idea is then to say that for a ground to be explanatorily relevant, it must be unspecific enough to be a strong difference-maker, but among the strong difference-makers, it must be maximally informative and specific. At least in many cases, it seems to us, an appropriate standard for comparative informativeness is given by saying that the fact that P is more informative than the fact that Q iff $P <_T Q$. Our considerations then suggest that to be explanatorily relevant, a partial ground of a fact must be a cohesive difference-making ground of the fact, where this is understood thus:

(Df. Cohesive) P is a cohesive difference-making ground of $Q =_{\text{df.}} P \prec^{!}_{TD} Q$, and there is no P' such that P' is a strong difference-maker wrt Q and $P' <_T P$.

A cohesive difference-making ground of a given fact is a maximally informative strong difference-maker with respect to the fact. Specifying cohesive difference-makers seems to be a plausible desideratum for good grounding-explanations that mirrors comparable desiderata that have been proposed for causal explanations.[34] Of course, in order to substantiate this proposal, it would be desirable to corroborate it with more data than the example we have considered, and also to sketch its relation to other, extant theories of explanatory relevance. To explore these issues must, however, be deferred to future research.

8 Conclusion

Sometimes, the grounds of a given fact give rise to good grounding-explanations of the fact, but sometimes it appears that they do not. Likewise, sometimes the causes of a given event give rise to good causal explanations of the event's occurrence, and sometimes it appears that they do not. To delineate the class of explanatorily relevant causes of an event, some authors have appealed to a notion of difference-making, and proposed explications of that notion designed for the application to causes. In this paper, we have explored the potential for a parallel move in the case of grounding. We have firstly, described a notion of difference-making that applies to the partial grounds of a fact, and we have made a strong case that there are grounds on both sides of the division induced by the notion. We have also introduced a number of significant subdivisions within the realms of both difference-making and non-difference-making partial grounds. We have then made a preliminary case that these

[33] See Strevens (2008, 101ff). We adopt the terminology of a *cohesive* difference-maker from him.

[34] Note that in the debate on the metaphysics of causation, notions of difference-making that follow an idea analogous to the one we have been developing in this section have been suggested as necessary conditions for being a causal influence (often under the label 'proportionality'); cf. e.g. Yablo (1992, 273ff), and Sartorio (2005). For reasons outlined earlier (Sect. 5), we prefer to view cohesive difference-making as a condition (or a good-making feature) for *explanations*, rather than as a condition for the obtaining of the respective explanation-backing relation. A view congenial to ours is that of Weslake (2013, §6) who defends the idea that proportionality is a dimension of explanatory value.

divisions among the grounds of a fact correspond to qualitative differences between candidate grounding-explanations of the fact. In addition, we have applied our conceptual apparatus to the analysis of certain putative counter-examples to the transitivity of grounding, and we have argued that these may naturally be accommodated within our framework as mere failures of transitivity for good grounding explanations. This result is significant especially because, as we have also shown, Schaffer's alternative strategy to accommodate the counter-examples within a contrastive account of grounding fails to remove the threat to transitivity.

A number of questions remain open. The contention that our notions of difference-making relate in the way suggested to explanatory relevance has to be checked against a wide range of data. Undoubtedly, many other natural subdivisions among the difference-makers and non-difference-makers of the grounds of a fact that are similar to ours may be made, and their relations to each other and to explanation examined. And of course, the theory of the various kinds of grounds we have distinguished, if they are found fruitful, still needs to be developed. What we hope to have shown, in addition to the more local results stated above, is that these open questions constitute worthwhile avenues for further research.

Acknowledgments We would like to thank our fellow phlox research-group members Michael Clark and Martin Lipman as well as Neil McDonnell for their comments on an earlier version of the paper. We'd also like to thank Ansten Klev and an anonymous referee of this journal for their comments. Stephan Krämer's work on this paper has been funded by the Deutsche Forschungsgemeinschaft (Grant KR 4516/1-1); Stefan Roski's work has been funded by the Behörde für Wissenschaft und Forschung Hamburg (Grant: Welt der Gründe). Both of us gratefully acknowledge the support.

Appendix 1: Logical cases of non-difference-making

We have claimed above that plausible assumptions concerning the logic of ground allow us to identify a *systematic* way of producing instances of non-difference-making ground given a pair of suitably independent facts to start with. This appendix develops our above, rough and informal argument for this claim in more detail.

The required logical assumptions concern the pure and the propositional logic of ground. Unfortunately, there is currently no standard, fully worked out propositional logic of ground. Rather, what we have is: (I) a system of natural inference rules that are plausible relative to a very *fine-grained* conception of ground, proposed by Kit Fine, but with as yet no adequate semantics, and acknowledged as incomplete with respect to any plausible understanding of ground (cf. Fine 2012a, 67); (II) a natural truthmaker-semantics, also proposed in Fine (2012a), which yields a logic adequate to a much more *coarse-grained* conception of ground; (III) a logical system proposed by Correia (2010), addressing again a coarse-grained conception of ground, proven sound and complete for a corresponding algebraic semantics.[35] An

[35] The fineness of grain of a conception of ground is a matter of the conditions under which sentences are substitutable *salva veritate* in the scope of a grounding-operator. More coarse-grained conceptions allow more substitutions, more fine-grained conceptions less.

additional difficulty is that the languages of these logical systems are expressively too weak to state anything like the quantified claim that for every Γ, if $\Gamma, P < Q$, then $\Gamma^0 < Q$ for some $\Gamma^0 \subseteq \Gamma$ with $P \notin \Gamma^0$.

Nevertheless, we believe a useful case for our claim can be made. To this end, we shall do two things. *Firstly*, we discuss a sharpening of our above informal argument in the context of the truthmaker-semantics for ground. We suggest that in this context, it is best to work with a slightly more refined conception of what it is for Γ to *contain* a full ground of some fact than the simple set-theoretic conception employed above. It can then be shown that for suitably independent, true P and Q, the informal argument given above goes through. It would take a lot of work to reconstruct the argument within the logic of (III), so we choose not to do so here. However, it is known that the logic obtained on the semantics in (II) is very close to that in (III) (cf. Fine ms, 11), and it is clear that a version of our argument can be given for Correia's system, too. *Secondly*, we discuss our informal argument in the context of the system (I). We show that given an additional rule that turns out valid on both of the only two semantics known to us that validate the others of Fine's rules, our argument goes through given very modest assumptions of independence of P and Q. Without the additional rule, a move parallel to that made before of refining the relevant notion of containment will secure our result.

Appendix 1.1: The coarse-grained framework

In his truthmaker-semantics for ground, Fine associates with each sentence A a set of states that verify A and a set of states that falsify A. States are here thought of as *obtaining*. Since no sentence is both true and false and no sentence is neither, exactly one of the two sets of states associated with a given sentence is empty.[36] For our purposes, it is easiest to reason directly about the sets of states and forget about the sentences to which they are assigned. We may then think of a non-empty set of states as a true proposition which is verified by exactly its members. Following Fine, we assume that (i) whenever there are states $s, t, u, \ldots$, there is also their fusion $\bigsqcup\{s, t, u, \ldots\} = s \sqcup t \sqcup u \sqcup \ldots$,[37] and that (ii) any proposition P is closed under non-empty fusions, so that $\bigsqcup P' \in P$ whenever $\emptyset \subset P' \subseteq P$.

The full and partial notions of (strict) ground are defined by Fine in terms of prior notions of full and partial *weak* ground. The definitions are as follows:[38]

(Df. $\leq$) $\quad P_1, P_2, \ldots \leq Q =_{\text{df.}} s_1 \sqcup s_2 \sqcup \ldots \in Q$ whenever $s_1 \in P_1, s_1 \in P_2, \ldots$

[36] Fine in Fine (2012a) actually speaks of *facts* rather than states. Since we have used 'fact' for the relata of the grounding relation facts, to avoid terminological confusion, we use 'state' instead. (Note that in other versions of his truthmaker semantics, Fine also appeals to states, and then typically allows non-actual and indeed often impossible states alongside the actual, obtaining ones. So we should perhaps emphasize again that we here restrict 'state' to actual, obtaining states.)

[37] Fusion is assumed to be associative, so that $\bigsqcup P_0 \sqcup \bigsqcup P_1 \sqcup \ldots = \bigsqcup(P_0 \cup P_1 \cup \ldots)$. Use of this fact will often be tacit in what follows.

[38] With respect to (Df. $\prec$), it should be noted that this definition of $\prec$ is faithful to our above understanding of partial ground as applying to the parts of a strict ground, and thereby deviates from Fine's definition of $\prec$ by the condition that $P \preceq Q$ and $Q \npreceq P$.

(Df. $\preceq$) $P \preceq Q =_{\text{df.}} \Gamma, P \leq Q$ for some set of propositions Γ
(Df. $<$) $P_1, P_2, \ldots < Q =_{\text{df.}} P_1, P_2, \ldots \leq Q$ and $Q \not\preceq P_i$ for all i
(Df. $\prec$) $P \prec Q =_{\text{df.}} \Gamma, P < Q$ for some set of propositions Γ

Now say that a state s is part of a state t iff for some state s', $t = s \sqcup s'$, and say that two states s and t overlap iff they share some (non-null[39]) part.

The conjunction $P \wedge Q$ of two propositions is the set $\{s \sqcup t : s \in P \text{ and } t \in Q\}$, and the disjunction $P \vee Q$ is the set $P \cup Q \cup (P \wedge Q)$. Now consider any truths P, Q which are independent in the sense that $\bigsqcup P$ and $\bigsqcup Q$ do not overlap. We may now show that

(A) $P \prec Q \vee (P \wedge Q)$

This is immediate by definition given that $P, Q < Q \vee (P \wedge Q)$, which may be established as follows. Suppose $s \in P$ and $t \in Q$. Then $s \sqcup t \in P \wedge Q$, and hence $s \sqcup t \in Q \vee (P \wedge Q)$. So $P, Q \leq Q \vee (P \wedge Q)$. Now consider $\bigsqcup(Q \vee (P \wedge Q)) = \bigsqcup P \sqcup \bigsqcup Q$. Since $\bigsqcup P$ and $\bigsqcup Q$ do not overlap, $\bigsqcup P \sqcup \bigsqcup Q$ is not a part of $\bigsqcup P$ or $\bigsqcup Q$. But it is easy to verify that $R \preceq P$ only if $\bigsqcup R$ is part of $\bigsqcup P$, so it follows that $Q \vee (P \wedge Q) \not\preceq P$ and that $Q \vee (P \wedge Q) \not\preceq Q$. Hence $P, Q < Q \vee (P \wedge Q)$.

Now suppose that some scenario S which includes the fact that P contains a scenario which is a strict full ground of $Q \vee (P \wedge Q)$. To show that P is not a difference-maker, we need to show that $S \backslash \{$the fact that $P\}$ still contains a strict full ground of P. If this fails for any scenario, it fails for a scenario which *is* a strict full ground of $Q \vee (P \wedge Q)$, so we may restrict attention to such scenarios. Suppose, therefore, that

(H) $\Gamma, P < Q \vee (P \wedge Q)$

Writing $\bigwedge \Gamma$ for the conjunction of all the members of Γ,[40] it follows that every verifier of $\bigwedge \Gamma \wedge P$ is a verifier of $Q \vee (P \wedge Q)$. Since any such verifier has a part that verifies P, by the assumption of independence, no such verifier is a verifier of Q, and hence it must be a verifier of $P \wedge Q$. But since no verifier of P is part of a verifier of Q, it follows that every verifier of $\bigwedge(\Gamma \backslash \{P\})$ must have a part that verifies Q.

But does it follow that $\Gamma \backslash \{P\}$ contains a strict full ground of $Q \vee (P \wedge Q)$? If a scenario contains another just in case the latter is a subset of the former, it does not.[41] In a rough approximation, the problem is that although some subset of $\Gamma \backslash \{P\}$

[39] Since fusion is defined for the empty set of states, we always have a minimal state which is part of every state, called the *nullstate*. The relation of overlap must therefore be taken to require sharing of non-null parts, otherwise it trivializes.

[40] More formally, let Γ be indexed by an index set I, so that $\Gamma = \{P_i : i \in I\}$. Then $\bigwedge \Gamma = \{\bigsqcup\{f(i) : i \in I\} : f \in \Pi\{P_i : i \in I\}\}$.

[41] Here is a counter-example. Suppose that $P = \{s\}$, $Q = \{t\}$, where t and s do not overlap. Then $Q \vee (P \wedge Q) = \{t, s \sqcup t\}$. Now suppose $R = \{u\}$ with u strictly between t and $s \sqcup t$. Note that $s \sqcup u = s \sqcup t$. Now consider $\Gamma = \{R\}$. Then $\Gamma, P \leq Q \vee (P \wedge Q)$. But as before, $Q \vee (P \wedge Q) \not\preceq P$, and likewise $Q \vee (P \wedge Q) \not\preceq R$, since u is a proper part of $s \sqcup t$. So $\Gamma, P < Q \vee (P \wedge Q)$. But $R \not\subseteq Q \vee (P \wedge Q)$, so $R \not\leq Q \vee (P \wedge Q)$, and a fortiori, $R \not< Q \vee (P \wedge Q)$.

must be strong enough to establish Q, it may be that every such subset is, as it were, *strictly between* Q and $P \wedge Q$, and thereby fail to ground $Q \vee (P \wedge Q)$.[42]

However, intuitively, what this shows is not that P may after all be a difference-maker with respect to $Q \vee (P \wedge Q)$, but that the set-theoretic interpretation of the containment of one scenario in another is unsatisfactory. For in the intuitively relevant sense, since every verifier of $\bigwedge(\Gamma \backslash \{P\})$ contains a verifier of Q, $\Gamma \backslash \{P\}$ contains $\{Q\}$, and thereby a strict full ground of $Q \vee (P \wedge Q)$. So what we should do is refine the conception of containment appealed to in the definition of difference-making. We shall therefore say that a scenario Γ *contains* a scenario Γ^0 just in case every verifier of $\bigwedge \Gamma$ has a part that verifies $\bigwedge \Gamma^0$. Then since $\Gamma \backslash \{P\}$ contains $\{Q\}$ whenever $\Gamma, P < Q \vee (P \wedge Q)$, P comes out a non-difference-making partial ground of $Q \vee (P \wedge Q)$.

Appendix 1.2: The fine-grained framework

We use the following rules:[43]

Sub($<$/ $\prec$)	From $\Gamma, A < B$ infer $A \prec B$
Sub($<$/$\leq$)	From $\Gamma < B$ infer $\Gamma \leq B$
Sub($\leq$/$\preceq$)	From $\Gamma, A \leq B$ infer $A \preceq B$
Sub($\prec$/$\preceq$)	From $A \prec B$ infer $A \preceq B$
Trans($<$/$<$)	From $A < B$ and $B < C$ infer $A < C$
Trans($\leq$/$<$)	From $\Gamma \leq A$ and $A < B$, infer $\Gamma < B$
Trans($\prec^{sp}$/$\preceq$)	From $A \prec^{sp} B$ and $B \preceq C$ infer $A \prec^{sp} C$
Irr($\prec^{sp}$)	From $A \prec^{sp} A$, infer $\bot$
Rev Sub($\leq$/$<$)	From $A_1, A_2, \ldots \leq B$ and $A_1 \prec^{sp} B$, $A_2 \prec^{sp} B$, …, infer $A_1, A_2, \ldots < B$
($\wedge$I)	From A and B, infer $A, B < A \wedge B$
($\vee$I)	From A, infer $A < A \vee B$ or $B < A \vee B$
($\vee$E)	From $\Gamma < A \vee B$, infer that either $\Gamma \leq A$, or $\Gamma \leq B$, or for some Γ_A, Γ_B with $\Gamma = \Gamma_A \cup \Gamma_B$: $\Gamma_A \leq A$ and $\Gamma_B \leq B$
($\wedge$E)	From $\Gamma < A \wedge B$ infer that for some Γ_A, Γ_B with $\Gamma = \Gamma_A \cup \Gamma_B$: $\Gamma_A \leq A$ and $\Gamma_B \leq B$

We make the following assumptions.

[42] Note that if we were to assume that every proposition P is *convex* in the sense that $u \in P$ whenever $s, t \in P$, s is part of u, and u is part of t, this case cannot obtain. This is significant given that according to Fine (ms, 11), grounding defined as above on convex propositions coincides with ground as per the logic of Correia (2010). So it would appear that with respect to Correia's logic, the argument for P being a non-difference-making partial ground of $Q \vee (P \wedge Q)$ indeed goes through in its original form. For criticism of the convexity constraint and the corresponding principle in Correia's logic, see Krämer and Roski (2015) and Correia (2016).

[43] We write $\prec^{sp}$ for the Finean notion of strict partial ground, defined as non-mutual weak partial ground. It should be emphasized that Fine puts forth all these rules with a reading of $\prec$ as strict partial ground in mind. However, in the cases in which we substitute the notion of partial strict ground, i.e. Sub($<$/$\prec$) and Sub($\prec$/$\preceq$), it is clear that they retain all of their plausibility under this reinterpretation.

(A1) P
(A2) Q
(A3) $P \npreceq Q$

Using the rules ($\wedge$I), ($\vee$I), Trans(<), and Sub(</$\prec$) we obtain

(1) $P \prec Q \vee (P \wedge Q)$

Now suppose

(S) $\Gamma, P < Q \vee (P \wedge Q)$

By ($\vee$E), we may infer from (S) that one of the following three claims holds:

(i) $\Gamma, P \leq Q$
(ii) $\Gamma, P \leq P \wedge Q$
(iii) For some Δ_1, Δ_2 with $\Delta = \Delta_1 \cup \Delta_2$: $\Delta_1 \leq Q$ and $\Delta_2 \leq P \wedge Q$

But from (i), it follows by Sub($\leq$/$\preceq$) that $P \preceq Q$, contrary to our assumption.

Now suppose (iii), and let $\Delta_1 \leq Q$. Note that $P \notin \Delta_1$, for otherwise again $P \preceq Q$, contrary to our assumption. So Δ_1 is a subset of $\Gamma \cup \{P\}$, not including P, with $\Delta_1 \leq Q$. Using that $Q < Q \vee (P \wedge Q)$ as well as Trans($\leq$/<), we obtain that $\Delta_1 < Q \vee (P \wedge Q)$. So Δ_1 is the required witness for the claim that P is a non-difference-making partial ground of $Q \vee (P \wedge Q)$ for the case that (iii) holds.

Suppose finally that (ii) holds. Then by Sub($\leq$/$\preceq$), we have $R \preceq P \wedge Q$ for all $R \in \Gamma, P$. Since any relationship of weak partial ground is either mutual or a relationship of strict partial ground, it follows by Rev Sub($\leq$/<) that one of the following two conditions holds:

(a) $\Gamma, P < P \wedge Q$
(b) $P \wedge Q \preceq R$ for some $R \in \Gamma, P$

If (a), then by ($\wedge$E) we obtain that for some subset Δ_1 of Γ, P, we have $\Delta_1 \leq Q$. Note that $P \notin \Delta_1$ for otherwise $P \preceq Q$, contrary to our assumption. So Δ_1 is a subset of $\Gamma \cup \{P\}$, not including P, with $\Delta_1 \leq Q$, and thereby as before, $\Delta_1 < Q \vee (P \wedge Q)$. So Δ_1 is the required witness for the claim that P is a non-difference-making partial ground of $Q \vee (P \wedge Q)$ for the case that (a) holds.

If (b) holds, then we have for some $R \in \Gamma \cup \{P\}$ both $P \wedge Q \preceq R$ and $R \preceq P \wedge Q$. It is tempting to infer from this that $P \wedge Q$ and R are mutual weak *full* grounds, as per the rule:

Rev Sub($\preceq$/$\leq$) From $A \preceq B$ and $B \preceq A$, infer $A \leq B$

And on the only proposals for a *semantics* for a logic of ground incorporating Fine's rules—in as yet unpublished work by Krämer (ms) and Correia (ms)—this inference indeed comes out valid. But then $R \leq P \wedge Q$, and since $P \wedge Q < Q \vee (P \wedge Q)$ we obtain $R < Q \vee (P \wedge Q)$. Now it can be shown that $R \neq P$. For otherwise we obtain $P \leq P \wedge Q$. But $P \wedge Q \npreceq P$, for otherwise we obtain by Trans($\prec^{sp}$/$\preceq$) that $P \prec^{sp} P$, and thus $\perp$ by Irr($\prec^{sp}$). It follows by Rev Sub($\leq$/<) that $P < P \wedge Q$. But then by the elimination rule ($\wedge$E), $P \leq Q$ and hence $P \preceq Q$, contrary to our assumption. So

$P \notin \{R\}$, and hence $\{R\}$ is our required witness for the claim that P is a non-difference-making partial ground of $Q \vee (P \wedge Q)$ for the case that (iii) holds.

Now it may be that there could also be a plausible semantics for the kind of logic of ground Fine proposes which delivers natural counter-models to Rev Sub($\preceq / \leq$). In that case R would again, roughly speaking, have to lie between Q and $P \wedge Q$ in strength. But then similar means to those proposed above would give a natural way to salvage our case by refining the conception of containment to fit the relevant sense of 'between'.

References

Audi, P. (2012a). A clarification and defense of the notion of grounding. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding* (pp. 101–121). Cambridge, UK: Cambridge University Press.

Audi, P. (2012b). Grounding: Toward a theory of the in-virtue-of relation. *The Journal of Philosophy*, *109*(12), 685–711.

Bromberger, S. (1965). An approach to explanation. In R. J. Butler (Ed.), *Analytical philosophy, vol. 2*, pp. 72–105

Correia, F. (2010). Grounding and truth-functions. *Logique et Analyse*, *53*(211), 251–279.

Correia, F. (2016). On the logic of factual equivalence. *The Review of Symbolic Logic*, *9*(01), 103–122.

Correia, F. (ms). An impure logic of representational grounding. unpublished manuscript.

Correia, F., & Schnieder, B. (2012). Grounding: An opinionated introduction. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding: Understanding the structure of reality* (pp. 1–36). Cambridge: Cambridge University Press.

Fine, K. (2012a). Guide to ground. In B. Schnieder & F. Correia (Eds.), *Metaphysical grounding* (pp. 37–80). Cambridge: Cambridge University Press.

Fine, K. (2012b). The pure logic of ground. *The Review of Symbolic Logic*, *5*(1), 1–25.

Fine, K. (ms). A theory of truth-conditional content II. unpublished manuscript.

Heslow, G. (1981). The transitivity of causation. *Analysis*, *41*(3), 130–133.

Krämer, S. (2016). A hyperintensional criterion of irrelevance. *Synthese* pp. 1–14. doi:10.1007/s11229-016-1078-0.

Krämer, S. (ms). Towards a theory of ground-theoretic content. unpublished manuscript

Krämer, S., & Roski, S. (2015). A note on the logic of worldly ground. *Thought: A Journal of Philosophy*, *4*(1), 59–68.

Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.

Lewis, D. (1986a). Causal explanation. In D. Lewis (Ed.), *Philosophical papers* (Vol. II, pp. 214–240). Oxford, New York: Oxford University Press.

Lewis, D. (1986b). Causation. In D. Lewis (Ed.), *Philosophical papers* (Vol. II, pp. 159–213). Oxford, New York: Oxford University Press.

List, C., & Menzies, P. (2009). Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy*, *106*(9), 475–502.

Litland, J. E. (2013). On some counterexamples to the transitivity of grounding. *Essays in Philosophy*, *14*(1), 19–32.

Owens, D. (1992). *Causes and coincidences*. Cambridge, UK: Cambridge University Press.

Ruben, D. H. (2012). *Explaining explanation (updated and extended second edition)*. Boulder: Paradigm Publishers.

Sartorio, C. (2005). Causes as difference-makers. *Philosophical Studies*, *123*, 71–96.

Schaffer, J. (2012). Grounding, transitivity, and contrastivity. In B. Schnieder & F. Correia (Eds.), *Metaphysical grounding* (pp. 122–138). Cambridge: Cambridge University Press.

Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies*, *173*(1), 49–100.

Schnieder, B. (2010). A puzzle about 'because'. *Logique et Analyse*, *53*(211), 317–343.

Schnieder, B. (2015). The asymmetry of 'because'. *Grazer Philosophische Studien*, *91*, 175–208.

Strevens, M. (2008). *Depth*. Cambridge, MA: Haward University Press.

Trogdon, K. (2013). An introduction to grounding. In M. Hoeltje, B. Schnieder, & A. Steinberg (Eds.), *Varieties of dependence* (pp. 97–122). Munich: Philosophia.

Weslake, B. (2013). Proportionality, contrast, and explanation. *Australasian Journal of Philosophy*, *91*(4), 785–797.

Woodward, J. (2003). *Making things happen*. Oxford, New York: Oxford University Press.

Yablo, S. (1992). Mental causation. *The Philosophical Review*, *101*(2), 245–280.

The Whole Truth

Stephan Krämer

Abstract We often care not just whether an account of some subject matter is correct, but also whether it is complete—the whole truth, as we might say. This chapter criticizes extant intensional explications of the notion of a whole truth by showing that they yield implausible results in an important range of cases. The difficulty is traced to the inability of an intensional framework to adequately capture constraints of relevance imposed by an intuitive understanding of the whole truth. I go on to develop and defend a novel account of what it is for a truth P to be the whole truth with respect to a subject matter: roughly speaking, it is for every fact pertaining to the subject matter to be relevant to making P true, or equivalently, for P to relevantly entail every truth pertaining to the subject matter. The proposal is formally spelled out within the framework of truthmaker semantics as developed by Kit Fine in a series of recent publications. As part of this, a novel, truthmaker-based semantics for the totality operator '. . . and that's it' is sketched and argued to be superior to previous intensional accounts.

1 Introduction

In both scientific and everyday contexts, we often care not just whether a given account of some subject matter is *correct*—whether it describes its subject matter accurately—but also whether it is *complete*—whether it describes its subject matter exhaustively. In this chapter, I propose a novel account of what it is for a proposition to be both correct and complete—*the whole truth*, as one might say—with respect to a given subject matter. The most distinctive feature of the account is that it imposes requirements of *relevance*. Roughly speaking, on the view to be developed, for P to be a whole truth with respect to a given subject matter is for every fact pertaining

Stephan Krämer
University of Hamburg, Department of Philosophy, Überseering 35, 22297 Hamburg, Germany
e-mail: stephan.kraemer@uni-hamburg.de

to that subject matter to be *relevant* to making P true, or equivalently, for P to *relevantly* entail every truth pertaining to the subject matter.

Section 2 clarifies, by means of examples, the notion of a whole truth that I aim to capture, and uses the examples to motivate two preliminary informal characterizations of that notion. I then turn to the task of making these characterizations formally precise. Section 3 presents a natural formalization within the intensional framework of possible worlds semantics, and argues that it yields implausible results in an important range of cases. The difficulty is traced to the inability of an intensional framework to adequately capture the relevance constraints imposed by an intuitive understanding of the whole truth. This diagnosis is then used to rule out as unpromising a number of otherwise natural seeming strategies to refine the possible worlds analysis. Section 4 introduces the hyperintensional framework of truthmaker semantics, as recently developed by Kit Fine (esp. [7, 8]) and uses it to develop alternative formalizations of our informal characterizations of the whole truth. It is then argued in section 5 that this proposal avoids the difficulties faced by the previous accounts. Section 6 shows how to extend the truthmaker-based account by so-called totality operators to capture an important weakening of the previous notion of a whole truth. The operators are also given a hyperintensional, relevantist treatment, which is again argued to be superior to extant intensional accounts. Section 7 concludes.

2 Some examples, and two informal characterizations

The notion of a complete truth has application in a wide range of contexts. A first, very large class of examples concerns the interpretation of answers to wh-questions. Thus, consider the following question-answer pair:

Sara: What did you have for breakfast?
Jack: I had eggs, bacon, and orange juice.

At least in typical contexts, the natural interpretation of Jack's answer to Sara's question is as *exhaustive*. So interpreted, Jack's answer is incompatible with his also having had coffee for breakfast. Correspondingly, Jack's answer is wholly appropriate only if it is a *complete* truth with respect to the subject matter of what Jack had for breakfast. Thus, the evaluation of Jack's answer as appropriate or otherwise involves at least a tacit application of a concept of a complete truth. Parallel comments apply for most answers to wh-questions.[1]

[1] This observation and the concept of a complete, or *exhaustive* answer to a question constitute an important topic in the theory of questions within linguistics. The classical reference is [10]. It should be pointed out that for some questions, like 'Where can I buy an international newspaper?', incomplete answers—often called 'mention-some' answers in contrast to complete, 'mention-all' answers—do seem entirely felicitous. It is controversial whether this is due purely to pragmatic factors, or whether there are semantic differences involved. That question is immaterial for our purposes though.

There are also a number of more specialized contexts in which the notion of a whole truth is of particular significance. Perhaps the most notorious of these is the context of legal testimony. It often matters a great deal not just whether a witness's account is true, but also whether it is in some relevant sense exhaustive—as reflected in the oath to tell the truth, *the whole truth*, and nothing but the truth.[2] In addition, there are several more *theoretical* contexts in which the notion of a complete truth has application. For instance, in metaphysics, a number of important hypotheses may be construed as claims concerning what *kind* of truth might be complete. Thus, physicalism may be understood, in a first approximation, as the claim that a complete physical description of the world is by itself already a complete description of the world *full stop*.[3] A second possible application in metaphysics concerns the grounds of universal quantification. It has been argued that any true universal quantification $\forall x\, Fx$ is fully grounded by the totality of its instances Fa, Fb, … together with the totality claim that $a, b, \ldots$ are all the things that exist.[4] That latter claim might plausibly be understood as the claim that the conjunction of the proposition that a exists, and the proposition that b exists, and … is a complete truth with respect to the question of what individuals exist.

Another class of examples, from logic and epistemology, concerns defeasible forms of reasoning. Since birds can normally fly, from the premise that Tweety is a bird one may defeasibly infer that Tweety can fly. In making this inference, we might say, one tacitly conjectures that the premise constitutes a *complete truth* with respect to those of Tweety's properties that are relevant to Tweety's being able to fly—and so in particular, that Tweety is not a penguin, and does not have broken wings, etc.[5] A related issue shows up in the context of Bayesian epistemology. Assume that Bob's credence in the proposition that Tweety can fly given that Tweety is a bird is high—say, 0.9. Then according to standard Bayesianism, when Bob learns that Tweety is a bird, he should update by conditionalizing on this newly obtained evidence and set his credence in the proposition that Tweety can fly to 0.9. However, this is so only under the assumption that Tweety's being a bird is *all* that Bob learns in the situation under discussion—and so in particular that Bob does not also learn that Tweety is a penguin (and that penguins can't fly).[6] Whether Bob's setting his

---

[2] Exactly how the phrase 'the whole truth' is to be understood in this context is admittedly a delicate matter. Since the witness will typically not have knowledge of every aspect of the situation under discussion, the requirement to tell the whole truth is presumably not that of giving an accurate and complete account of the situation. Rather, the requirement might be to accurately and completely report one's pertinent *knowledge* about the situation. By doing so, one then implicitly specifies a whole truth with respect to the matter of what pertinent knowledge one possesses. The logic of the notion of all one knows (with respect to a certain subject matter) is studied within a standard intensional framework by [13].

[3] This has been stressed by [12, p. 529].

[4] For a statement of the view in the context of the contemporary theory of ground, see [4, pp. 60ff].

[5] In the theory of defeasible reasoning, this assumption is often called the closed-world assumption. Similar forms of completeness-assumptions may be seen to motivate John McCarthy's influential system for defeasible reasoning, the logic of circumscription ([17, 18]).

[6] In Bayesian epistemology, this is captured in the requirement of total evidence, to the effect that one must always update on one's total evidence.

credence to 0.9 is rational thus depends on whether the truth that Bob learned that Tweety is a bird is a complete truth with respect to the question of what Bob learned in the relevant situation.[7]

What, then, does it mean for a truth to be complete with respect to a subject matter? There are two natural and complementary strategies we may pursue in trying to clarify this. The first strategy is to try to say how the truth must relate to *other truths* in order to be complete. I shall call this the *horizontal* strategy. The obvious suggestion is that it must *entail every other truth* that pertains to the subject matter in question.[8] For suppose there is a truth about the relevant subject matter that is not entailed by a candidate whole truth P. Then P is compatible with a false hypothesis pertaining to that subject matter, and this would seem to constitute a way in which P falls short of completeness. If, on the other hand, P does entail every truth pertaining to the subject matter, then it is prima facie plausible to conclude that P is complete with respect to that subject matter.

The second strategy is to try to say how the truth must relate to *the world* in order to be complete. I shall call this the *vertical* strategy. Here a natural suggestion is that a candidate whole truth P must *report every fact*—construed as a worldly item—that pertains to the relevant subject matter.[9] Thus, whenever it pertains to the subject matter that a certain fact obtains, then P must state that that fact obtains. If it does not, then this would seem to constitute a way in which P falls short of completeness. If, on the other hand, P does report every fact pertaining to the relevant subject matter, then it seems prima facie plausible to conclude that P is complete with respect to that subject matter.

It might be objected, certainly with respect to the first, and possibly with respect to the second suggestion, that it is too demanding. For note that one of the truths with respect to Jack's breakfast, for example, is (we may assume) that Jack did not eat a crocodile for breakfast. So by our first suggestion, a complete truth with respect

[7] Another area in which the notion of a complete truth with respect to some subject matter plays a prominent role is of course that of mathematical logic, when we ask whether a given formal theory constitutes a complete description of the mathematical structures that it aims to capture. I do think that at least some of the notions of completeness in play here belong to the same family as the notions considered here, and that there is considerable interest in applying the account I will develop to the mathematical cases. Nevertheless, I shall set these cases aside for the purposes of this chapter, both because they introduce too many additional complications and because for this area, there are perfectly clear and precise formal explications of completeness available, so it is less clear that there is any need for a novel account.

[8] Note that 'pertaining' should here be understood as 'wholly pertaining'. If a given truth P pertains in part to the relevant subject matter m and in part to some other subject matter m', then a complete truth about m may unproblematically fail to entail P simply because it is quiet on m'.

[9] For reasons analogous to those given in the previous footnote, 'pertaining' should be read as 'wholly pertaining'. I shall henceforth take this reading for granted.—My talk of a proposition's *reporting* a fact is perhaps in need of clarification. Ordinarily, we would think of the reporting of a fact as a linguistic activity carried out by speakers. In those cases, I assume, the fact that a speaker reports a certain fact f in presenting as true a proposition P is due in part to a certain relevant relationship between P and f, for which I also use the word 'report'. I assume, moreover, that we may extend the range of that relation to all propositions, independently of whether they are, or indeed can be, expressed by a speaker of some given language.

to the subject matter of what Jack had for breakfast would have to entail this truth. But at least on one natural understanding of completeness, this is not so. On that understanding, it is sufficient, roughly speaking, if a truth entails what Jack *did* eat for breakfast, and it need not also entail what Jack did *not* eat.[10] Similarly, it might be held that among the facts pertaining to the subject matter of what Jack had for breakfast is the fact that he did not eat a crocodile.[11] But intuitively, in one good sense of 'complete', a complete truth with respect to that subject matter need not report that fact: it need only report the facts concerning what Jack did eat, and may omit mention of facts concerning what he didn't eat.

The issue arises also for some of the theoretical applications of the notion of a complete truth mentioned above. The complete physical description of the world does not entail the truth that there are no demons (construed as non-physical objects). But that does not, it seems, prevent it from being complete in the, or at least a, metaphysically important sense of 'complete'. Again, this sense seems only to require a complete truth to entail every truth about what there is, and what it is like, not the truths about what there is *not*, and what things are *not* like (cf. [12, pp. 529f]).

Fortunately, it appears that there is a systematic way of bridging the gap between a complete truth in this weaker sense and a complete truth in the stronger sense. For if P is a complete truth in the weaker sense, then firstly, it is true to say that P holds, *and that's it*. Secondly, this strengthened truth appears then to entail all the missing negative truths. For instance, if we add to the complete physical description of the world the claim that *that's it*, the result does seem to rule out the existence of demons. This suggests that we can accommodate the objection against the above horizontal and vertical proposals by extending them with an account of the *totality operator* 'that's it'.[12] I shall take up this task in section 6 below. Until then, I shall focus on the stronger notion of completeness.

There is a second way in which the suggested characterizations of the notion of a complete truth may seem too demanding, namely in that they require a complete truth to be maximally specific: to describe its subject matter in full detail. But returning to Jack's breakfast, there is a truth and a fact pertaining to that subject matter stating *exactly* how much orange juice Jack had. But clearly, Jack's account of his breakfast does not have to entail that truth, or report that fact, in order to meet the contextually relevant standard of completeness. Strictly speaking, therefore, our account of the whole truth should be relativized to a contextually relevant level of specificity. Since making such relativization explicit would mainly be a distraction in the discussion to follow, I shall leave it implicit throughout, and merely occa-

[10] In the linguistic literature on questions, a distinction is standardly made between strongly exhaustive and weakly exhaustive answers, which mirrors the distinction adverted to in the text. Thus, a strongly exhaustive answer to the question who came to the party must also specify who did not come, whereas a weakly exhaustive one need only say who did come.

[11] Unlike the existence of negative truths, the existence of negative facts is controversial. The complaint is therefore less obviously compelling as applied to the second, vertical proposal.

[12] A suggestion like this is made in [1], and formally elaborated in [12]. The linguistic accounts of questions and (exhaustive) answers sometimes make use of an 'exhaustivization' operator that plays a very similar role. An operator of this sort was first introduced by [10, esp. ch. V].

sionally comment on how the notion of a level of specificity could be implemented formally.

For definiteness, let me state explicitly the two informal characterizations of the notion of a complete, or whole truth suggested above:

(Horizontal) A proposition *P* is a complete truth (wrt subject matter *m*) iff
(i) *P* is true, and (ii) *P* entails every truth (pertaining to *m*).
(Vertical) A proposition *P* is a complete truth (wrt subject matter *m*) iff
(i) *P* is true, and (ii) *P* reports every fact (pertaining to *m*).

Note that I have stated the suggestions as necessary and sufficient conditions for a proposition to be *a* complete truth (with respect to a subject matter). But especially when using 'whole', rather than 'complete', it also seems very natural to use the definite article, referring to *the* whole truth, rather than a whole truth – indeed, I have frequently done so myself in the preceding paragraphs, and I shall continue to do so when it seems stylistically preferable. Use of the definite article suggests uniqueness, of course, in the sense that there can only be one proposition deserving of the title *whole truth* (with respect to a given subject matter). At least for subject-matter restricted notions of completeness, (Horizontal) and (Vertical) clearly do not guarantee uniqueness, since any truth entailing a truth that is complete with respect to a subject matter *m* will also be counted complete with respect to *m*. For the absolute notion of completeness, (Horizontal) implies that any two complete truths entail one another and so are intensionally equivalent. But since I shall be rejecting the identification of intensionally equivalent propositions, this does not by itself yield uniqueness. For the most part of my discussion, I shall focus on completeness understood as setting a threshold – roughly, to be complete, a truth just needs to be strong enough – but I shall comment on occasion on the possibility of specifying a narrower condition that would guarantee uniqueness.

The two characterizations (Horizontal) and (Vertical) are neither obviously equivalent nor obviously incompatible. They jointly entail that a truth entails every truth (pertaining to a given subject matter *m*) iff it states every fact (pertaining to *m*), and relative to that assumption, they are equivalent. Pending clarification of the notion of entailment and the notion of fact-stating, it is unclear whether the assumption should be taken to hold. It seems plausible, however, that the notions of entailment and fact-stating should permit of reasonably natural explications under which the assumption comes out true. And given the independent intuitive appeal of (Horizontal) and (Vertical), it seems desirable to give such an explication of the characterizations.

In order to develop our preliminary informal characterizations into an adequate and precise analysis of the notion of a complete truth, we thus need to clarify the key concepts invoked in (Horizontal) and (Vertical) in an appropriate way. More specifically, we need to come up with appropriate answers to the following questions:[13]

[13] Admittedly, since (Horizontal) and (Vertical) invoke the relations of entailment and reporting only with respect to truths and facts, respectively, the questions below are slightly more general than is perhaps strictly required for the purpose of clarifying these characterizations. But since we are naturally interested not just in which propositions actually are complete truths, but also in

Q.1 What is a proposition?
Q.2 What is it for a proposition to be true?
Q.3 What is it for a proposition to entail another?
Q.4 What is a state of affairs?
Q.5 What is it for a proposition to report a state of affairs?
Q.6 What is it for a state of affairs to obtain, i.e. to be a fact?

Moreover, at least if we wish to give an explicit treatment of subject matter restrictions, we also need to answer these additional questions:

Q.7 What is a subject matter?
Q.8 What is it for a proposition to pertain to a subject matter?
Q.9 What is it for a state of affairs to pertain to a subject matter?

So let us see how this might be done.

3 Against intensional analyses

We may formulate prima facie natural and attractive answers to the above questions by appeal to the notion of a possible world, which for present purposes we may take as primitive. We may then answer Q.1–Q.3 above in the familiar way:

PW.1 A proposition is a set of possible worlds.
PW.2 A proposition is true iff it has the actual world as a member.
PW.3 A proposition P entails a proposition Q iff P is a subset of Q.

Let us set aside for the moment the matter of subject matter restrictions. Given PW.1–PW.3, we obtain the following account of the unrestrictedly whole truth:

PW.UWT A proposition P is an unrestrictedly complete truth iff $P = \{@\}$.

(@ is the actual world). To see this, note that since {@} is a truth, any complete truth must entail, and hence be a subset of, {@}. Since the empty set is a false proposition, {@} itself is the only true proposition satisfying this constraint.

The most obvious way to capture a notion of a state (of affairs[14]) within the possible worlds framework is to identify it with the proposition that the state obtains, i.e. the set of worlds in which the state obtains. For a state to obtain is then simply for the proposition it is identified with to be true. Q.4–Q.6 may then be answered as follows:

PW.4 A state of affairs is a proposition.
PW.5 A proposition P reports a state of affairs s iff P entails s.

---

which propositions would be complete truths under various counterfactual circumstances, we shall eventually have to answer the more general questions anyway.

[14] For brevity, I shall henceforth often simply speak of states.

PW.6 A state of affairs obtains iff it is true.

Evidently, with entailment and state-reporting thus explicated, both (Horizontal) and (Vertical), in their unrestricted versions, come out true under the present approach.

I accept that PW.UWT captures one reasonable sense of 'complete truth', and I accept that there may be natural applications of the notion of a complete truth for which PW.UWT is plausible. But, I maintain, there are also many applications of that notion with respect to which PW.UWT is wholly implausible. The easiest way to see this is to observe that in many contexts in which we may evaluate truths for completeness, necessary truths, and necessary consequences, do not come for free.[15] Let me give two examples.

(1) Plausibly, all of the truths of pure set theory are necessary. So they are all identified, under the possible worlds conception, with the set of all possible worlds. So, under the conception of entailment captured in PW.3, *every* proposition, and hence every truth, entails all of them. Hence, the truth that the empty set has no members, for example, is classified as a complete truth with respect to the subject matter of set theory. But this seems absurd. At least, there would appear to be a good sense of 'complete truth' in which this is not so.

(2) Plausibly, Socrates has his nature non-contingently: if Socrates is essentially F, then in every world in which Socrates exists, Socrates is essentially F. But then under the possible worlds approach, it is implausibly easy to give a complete account of the nature of Socrates. It suffices to state, firstly, that Socrates exists, and secondly, Socrates is essentially F only if the truth that Socrates exists entails that Socrates is essentially F. Again, this seems absurd. Certainly, there seems to be a good sense of 'complete truth' in which this is not so.[16]

Strictly speaking, both examples concern the notion of a *restrictedly* complete truth. But it is clear that this is an inessential feature of the examples. The restriction to subject matter plays no role in how the difficulty arises, it only makes it simpler to specify concrete examples of truths incorrectly classified as complete. Still, it is worth nevertheless to briefly examine how PW.UWT may be extended to cover restricted notions of completeness. Following [15, 14], we may take a subject matter m to be (represented by) an equivalence relation $\sim_m$ on the set of all possible worlds. The idea is that two worlds stand in the equivalence relation $\sim_m$ to one another just in case they are exactly alike as far as the subject matter m is concerned.

We then need to say what it is for a proposition to *pertain* to a subject matter. For a given equivalence relation $\sim_m$ on the set of worlds W and a world $w \in W$, let $[w]_m = \{v \in W : v \sim_m w\}$ be the equivalence class of w under $\sim_m$. Now consider a proposition P, and suppose that $w \in P$. We may then ask whether P is true at *every* world v for

[15] This kind of difficulty is of course familiar from many other applications of the possible worlds framework. It is nevertheless worthwhile to go through the issue in some detail, partly because it will help us identify the source of the difficulty in this application, and because it will provide us with a useful foil in the construction of a more satisfactory account below.

[16] Note that the metaphysical assumptions about essence are inessential. We might simply imagine someone raising the question which truths involving Socrates are necessary given that Socrates exists. Then under the intensional analysis, the answer 'Socrates exists' would come out as expressing the whole truth about that subject matter.

which $w \sim_m v$, i.e. at every world which is exactly like w with respect to m. If this is not so, then we may conclude that P is *not just about* m, but says something about how things are that do not concern the subject matter m. And if P is not just about m, it should not be a requirement on a complete truth with respect to m that it entail P. So we may take a proposition P to pertain to a subject matter m iff $[w]_m \subseteq P$ whenever $w \in P$. (This is equivalent to how Lewis characterizes a proposition's being *entirely about* a subject matter, cf. [14, p. 163].[17]) For explicitness, we state the resulting answers to Q.7–Q.9:

PW.7 A subject matter m is an equivalence relation $\sim_m$ on the worlds.
PW.8 A proposition P pertains to a subject matter m iff for all $w \in P$, $[w]_m \subseteq P$.
PW.9 A fact pertains to a subject matter iff the proposition that the fact obtains pertains to that subject matter.

We then obtain the following account of a restrictedly complete truth:

PW.RWT A proposition P is a complete truth wrt m iff $\{@\} \subseteq P \subseteq [@]_m$.

(PW.RWT reflects my policy, mentioned earlier, of focusing on completeness as a matter of being strong enough, by counting any truth that entails $[@]_m$ a complete truth with respect to m. There is, of course, an obvious way to strengthen the condition so as to yield a notion of P being *the* whole truth with respect m, by requiring $P = [@]_m$.[18])

Applying this apparatus to the example (1), we see that the subject matter of pure set theory ends up being identified with the universal relation on the worlds, since all worlds are alike with respect to pure sets. So there are only two propositions pertaining to that subject matter, namely the impossible proposition that is true at no worlds, and the necessary proposition true at all worlds. As a result, any truth comes out as a complete truth with respect to the subject matter of pure set theory. In the case of (2), the subject matter of the nature of Socrates is represented by a relation in which two worlds stand just in case either both contain Socrates or neither contains Socrates. The strongest truth pertaining to that subject matter is then the set of worlds in which Socrates exists, i.e. the proposition that Socrates exists. Hence a truth is complete with respect to the subject matter of the nature of Socrates iff it entails that Socrates exists.

Where has the possible worlds analysis gone wrong? It seems to me that the explication of state-reporting that the analysis involves is very implausible. In particular, it seems to me that the true proposition that Socrates exists does not, intuitively, (also) report the state that Socrates is human—even though necessarily, the state

[17] The same formal machinery of equivalence relations on the set of worlds could also be used to represent levels of specificity. An equivalence relation on the worlds would then be taken to represent a standard for specificity or level of detail on which exactly the non-equivalent worlds are distinguished.

[18] Formally speaking, the latter notion – essentially that of a truth being identical to an equivalence class of an equivalence relation on the worlds – coincides with the notion studied in [13], there in the context of an informal interpretation of a proposition capturing exactly what is known or believed (about a certain topic) by an epistemic agent.

obtains if the proposition is true. Likewise, it seems to me that the true proposition that no empty set has members does not, intuitively, (also) report the state that the empty set is a member of its singleton—even though necessarily, the state obtains if the proposition is true. What is missing, in spite of the presence of the right sort of *modal* connection between truth and state, is an appropriate connection of *relevance*.

For a proposition P to report a state, I would like to suggest, requires that the state be *wholly relevant* to making P true. And the state that Socrates is human is not wholly relevant to making it true that Socrates exists. It may not be wholly *irrelevant*, admittedly. For clearly, the state that Socrates exists is wholly relevant to making it true that Socrates exists, and it is not implausible to suppose that the state that Socrates is human contains the state that Socrates exists as a part. But then the state that Socrates is human still contains *more* than just the state that Socrates exists, and what it contains beyond this state is intuitively irrelevant to making it true that Socrates exists. Similarly, the state that the empty set is a member of its singleton is not wholly relevant for making it true that no empty set has members. Here it even seems that the state is *wholly irrelevant* to the truth of that proposition.

This diagnosis suggests that a number of otherwise tempting strategies for refining the possible worlds analysis are non-starters. In particular, one might have been tempted to respond to the above difficulties by saying that a truth is complete with respect to a subject matter only if it *logically* entails every truth pertaining to that subject matter. But as is well known, even (classical) logical entailment does not guarantee a connection of relevance between a proposition and what it entails. Still, it is worthwhile to confirm the point by briefly discussing how the proposal might be implemented. One way to do this is to work with a more relaxed conception of a possible world that covers even metaphysically impossible worlds, so long as they are still logically possible.[19] One might then take there to be worlds without any sets, and worlds in which Socrates exists but is not human, since these are not logically inconsistent. But it seems to me that even logical truths and consequences should not in general come for free. For instance, a complete truth about identity should report the fact that if $x = y$ and $y = z$ then $x = z$. For a truth to report this fact, the fact needs to be relevant to the truth. And the fact is not relevant to every truth. It is entirely irrelevant, for example, to the truth that snow is white. So a proposition should not automatically count as reporting the fact that identity is transitive, just because that fact obtains as a matter of logical necessity.

4 A hyperintensional, truthmaker analysis

The discussion of the previous section suggests that in order to obtain an adequate analysis of complete truths, we need to take proper account of when a state is, or is

[19] If one does not wish to do this, one would have to consider sentences rather than propositions as bearers of completeness. This strikes me as unattractive in any case, since intuitively it is the content of a sentence, not the sentence itself, which is evaluated for completeness. But independently of that, the proposal would face the same problem as the one discussed in the main text.

not, *wholly relevant* to making a given proposition true. So what is it for a state to be thus relevant to a proposition? We may distinguish between two approaches to this question, the definitional and the axiomiatic one. Under the definitional approach, we start by trying to analyse the pertinent notion of relevance in different, independently understood terms, and then see how the relation thus defined behaves, and whether it fits our intuitive demands. Under the axiomatic approach, we begin by taking some connection of relevance between states and propositions as basic, and develop its theory as best we can in accordance with our intuitive understanding of the notion.[20] I shall here pursue the latter approach. The basic relevantist connection I shall avail myself of is that of *wholly relevant verification*, which is the key concept underlying the truthmaker conception of content as recently developed by Kit Fine.[21]

Within the truthmaker conception, the usual appeal to a category of possible worlds is replaced by an appeal to a more general category of *states*. The states are assumed to be ordered by a relation of part-whole ($\sqsubseteq$), and it is assumed that given any set of states $T = \{s_1, s_2, \ldots\}$, we may form the mereological fusion $\bigsqcup T = s_1 \sqcup s_2 \sqcup \ldots$ defined as the smallest state of which all of the fused states are parts. In contrast to a world, a state may be *incomplete*: it may leave open the truth-value of many propositions. It may also be *inconsistent* or *impossible*: it may verify both of a pair of incompatible propositions.[22] A notion of a possible world may be recovered as the notion of a maximal consistent state, i.e. a state that contains every state it is compatible with. We shall assume that every consistent state in the state-space is part of a possible world. An obtaining state may then be identified with a *fact*, and the fusion of all facts with *the actual world* @. A proposition P is true iff verified by at least one fact, or equivalently, if @ contains some verifier of P.

A proposition P is identified with a pair (P^+, P^-) of a non-empty set P^+, comprising the verifiers of P, and a non-empty set P^-, comprising the falsifiers of P. Note that there is no requirement that if a given state s verifies P, then any state t containing s as a part must verify P as well; this reflects the intuitive requirement that verification be *wholly* relevant. We do, however, impose modal constraints on how the verifiers and the falsifiers may be related. In particular, we demand that no falsifier be compatible with any verifier—call this *Exclusivity*—and that for every

---

[20] Once we have done that, we may of course, in a second and optional step, proceed to examine whether we might be able to give a reductive definition of our relevance connection under which it behaves as required by our theory.

[21] See esp. [7, 8]. The presentation of truthmaker semantics to follow is heavily indebted to these works.

[22] It should be made clear that in the formal development of the theory, the distinction between possible and impossible states is officially taken as primitive. Two states are then said to be incompatible iff their fusion is impossible, and two propositions are incompatible iff every verifier of the one is incompatible with every verifier of the other. So when I here characterize an impossible state as one that verifies both of a pair of incompatible propositions, this should be taken as merely an informal gloss, relying on an intuitive understanding of what it means for a pair of propositions to be incompatible. Thanks here to an anonymous referee.

proposition, every world contains either a verifier or a falsifier of that proposition—call this *Exhaustivity*.[23]

Operations of conjunction, disjunction and negation on the propositions may then be defined as follows:

$$
\begin{aligned}
(\neg P)^+ &= P^- \\
(\neg P)^- &= P^+ \\
(P \wedge Q)^+ &= \{s \sqcup t : s \in P^+ \text{ and } t \in Q^+\} \\
(P \wedge Q)^- &= P^- \cup Q^- \\
(P \vee Q)^+ &= P^+ \cup Q^+ \\
(P \vee Q)^- &= \{s \sqcup t : s \in P^- \text{ and } t \in Q^-\}
\end{aligned}
$$

We are now ready to answer some of the questions Q.1–Q.9:[24]

TM.1 A proposition is a pair of a set of verifiers and a set of falsifiers.
TM.2 A proposition is true iff at least one of its verifiers obtains.
TM.4 A state is anything that can play the role of truthmaker for a proposition.
TM.6 A state of affairs obtains iff it is part of the actual world.

Let us consider next how the notion of state-reporting may best be explicated within this framework. In order to ensure that a state be wholly relevant to the truth of any proposition reporting it, it suffices to demand that the state be part of some verifier of the proposition. In order to also ensure that if a proposition reports a state, the truth of the proposition entails that the state obtains, we demand that the state be part of *every* verifier of the proposition.

TM.5 A proposition P reports a state s iff s is part of every verifier of P.

In principle, a weaker condition would have sufficed to ensure that the truth of a proposition requires the obtaining of any fact it states, namely that every verifier of the proposition *necessitates* the fact in question.[25] However, this would lead us to count the proposition that Socrates exists or (Socrates is human and Socrates is not human) as reporting the fact that Socrates is human. For that fact is part of the impossible state that Socrates is both human and not human, which verifies the second disjunct, and the fact is moreover necessitated by every verifier of the first disjunct. This seems to be an undesirable result. The non-modal, more wholeheartedly relevantist explication TM.4 provides a more natural and theoretically well-behaved conception of state-reporting.

[23] Intuitively, these constraints ensure bivalence of every proposition at every world, i.e. that at every world every proposition is either true or false but not both. Say that the propositions $P_1, P_2, \ldots$ modally entail Q iff Q is true at every world at which all of $P_1, P_2, \ldots$ are true. It can be shown that the logic of modal entailment for propositions satisfying *Exclusivity* and *Exhaustivity* is classical. (See [7] for discussion, from where I have also borrowed the terminology.)

[24] Note that TM.4 is offered simply as a potentially helpful gloss on what counts as a state. Officially, we are taking the notion of a state as basic, much as we did for the notion of a possible world in the previous section.

[25] If we assume that every consistent state is part of a maximal consistent state, i.e. a world, we may say that a state s necessitates a state t iff every world that contains s contains t. More generally, a state s necessitates a state t iff every state that is compatible with s is compatible with t.

Focusing for the moment on unrestrictedly whole truths, we now have all the material in place required for a full truthmaker-theoretic interpretation of (Vertical): a proposition P is a complete truth iff (a) P is true, so P is verified by some fact and (b) P reports every fact, so every verifier of P contains every fact. Condition (a) is equivalent to the condition that some verifier of P be part of @, and (b) to the condition that @ be part of every verifier of P.[26]

TM.UWT P is an unrestrictedly complete truth iff @ verifies P, and $@ \sqsubseteq s$ whenever s verifies P.

So under this approach, a truth is complete iff it reports the maximal fact that is the actual world, and thus only if the entire actual world is wholly relevant to the verification of the truth. A complete truth is allowed to have further verifiers, as long as these contain the actual world as a proper part. Note that since the actual world is a world, and hence a maximal consistent state, any such additional verifiers will be inconsistent states. One might thus consider singling out for special attention the 'pure' complete truths that are verified only by @. For now we rest content with TM.UWT.[27]

Let us now consider the notion of entailment and the interpretation of (Horizontal), as well as the question of its equivalence with (Vertical). It turns out that there is a very natural notion of entailment we can define within the truthmaker framework under which (Horizontal) comes out equivalent to (Vertical) as interpreted via TM.5. This is the relation that Fine calls *inexact entailment*, which obtains between a proposition P and a proposition Q just in case every verifier of P contains a verifier of Q as part (cf. e.g. [7, p. 669]).[28]

TM.3 A proposition P entails a proposition Q iff every verifier of P contains a verifier of Q.

For suppose that P is a whole truth as per (Vertical). Then P is verified by @, and $@ \sqsubseteq s$ whenever s verifies P. We may then show that P inexactly entails every truth,

[26] Since @ is a fact, if every verifier of P contains every fact, every verifier of P contains @. Since every fact is part of @, if every verifier of P contains @, by transitivity of part, every verifier of P contains every fact.—Given that every set of states T, including the empty set, has a fusion $\bigsqcup T$, which is the least upper bound of the set with respect to $\sqsubseteq$, we may dually also take any set of states T, including the empty set, to its biggest common part $\bigsqcap T$, which is the greatest lower bound of T. Writing P^+ for the set of verifiers of P, our condition might therefore be expressed more succinctly as: $@ \in P^+$, and $\bigsqcap P^+ = @$.

[27] Note that even insisting on purity in the sense described will not yield uniqueness, since there will be many propositions that are verified only by @ but differ with respect to their falsifiers. I shall not here pursue the question whether any of them could be seen as holding claim to the title of *the* unique whole truth.

[28] Strictly speaking, in the passage mentioned, Fine uses the term 'inexact consequence' and defines it for sentences, and by a slightly different condition which is however easily seen to be equivalent to the one I use. The differences are immaterial for our purposes.—It bears emphasis that my claim is simply that in order to obtain a plausible version of (Horizontal), the notion of entailment invoked in this principle should be understood as per TM.3. I do not intend to make any claim about what the 'real' notion of entailment is. In particular, I do not deny that every case of classical logical entailment or of ordinary modal entailment is a bona fide case of entailment.

and so is a whole truth as per (Horizontal). For let Q be any truth. Then @ contains some verifier of Q, call it t. Since every verifier of P contains @ as a part, and @ contains t as a part, it follows that every verifier of P contains t as a part. So P inexactly entails Q, and since Q was arbitrary, P inexactly entails every truth, and hence is a whole truth in the sense of (Horizontal). For the other direction, suppose P is a whole truth as per (Horizontal), so P is true and inexactly entails every truth. Then in particular, P inexactly entails every truth that is verified solely by @. So all of P's verifiers must contain @. Since P is true, at least one of P's verifiers is a fact. Since @ is the only fact that contains @, it follows that one of P's verifiers is @. So P reports every fact and hence is a whole truth as per (Vertical).

It is relatively straightforward to extend this account to accommodate notions of completeness restricted to some subject matter.[29] A subject matter may be thought of in the first instance as the question which of a certain set of states obtain. For example, the subject matter of the colour of my shirt may be thought of as the question which states to the effect that my shirt is of a certain colour obtain. Plausibly, if for each state s_i of some states $s_1, s_2, \ldots$ it pertains to a given subject matter m whether s_i obtains, then it also pertains to m whether $s_1 \sqcup s_2 \sqcup \ldots$ obtains, and it pertains to m whether a state t obtains whenever $t \sqsubseteq s_1 \sqcup s_2 \sqcup \ldots$ Given this assumption, we may simply represent a subject matter by a single state, namely the fusion of all states such that it pertains to the subject matter whether the state obtains.[30]

A state will therefore be taken to pertain to a subject matter iff it is a part of that subject matter. When should we take a proposition to pertain to a subject matter? I propose to take this to require that all the verifiers and all the falsifiers of the proposition be parts of the subject matter. For suppose a proposition P has a verifier (falsifier) that is not part of a given subject matter m. Then there is a way for P to be true (false) which turns on matters foreign to m. This seems to constitute a good sense in which P does not pertain purely or wholly to m. The proposal may also be inferred from Fine's suggestion (cf. [8, p. 676ff]) that the subject matter of a proposition may be identified with the fusion of all its verifiers and all its falsifiers, given the plausible assumption that a proposition's own subject matter should be the smallest subject matter the proposition wholly pertains to. Finally, the claim that this understanding of pertaining to a subject matter is appropriate at least for our purposes receives further confirmation from the fact that it allows us to establish the desired equivalence of (Horizontal) and (Vertical).

First, we make explicit our answers to Q.7–Q.9

TM.7 A subject matter is any state.
TM.8 A state pertains to a subject matter iff it is part of the subject matter.

[29] For a more detailed discussion of the notion of subject matter within truthmaker semantics, see [8, pp. 676ff].

[30] Note that a subject matter will typically be an impossible state, since it will result from fusing many mutually incompatible states. For instance, since both the state of my shirt being red all over and the state of my shirt being blue all over pertain to the subject matter of the colour of my shirt, the state representing it will contain both as part and therefore be impossible.—One might consider imposing a number of constraints on a state if it is to qualify as a subject matter, but we shall not discuss the matter here.

TM.9 A proposition pertains to a subject matter iff it is verified and falsified only by parts of the subject matter.

A restricted counterpart to TM.UWT may then be stated as follows:[31]

TM.RWT *P* is a complete truth with respect to subject matter *m* iff every verifier of *P* contains $@ \sqcap m$, and some verifier of *P* is part of @.

(Here $@ \sqcap m$ denotes the biggest state that is part of both @ and *m*, and hence the biggest fact pertaining to *m*.)

Under the interpretation provided by this account, (Horizontal) and (Vertical) are equivalent. For assume that *P* is a complete truth with respect to subject matter *m* according to (Horizontal). That is, *P* is true and *P* inexactly entails every truth pertaining to *m*. Then in particular, *P* inexactly entails every truth pertaining to *m* that is verified only by $@ \sqcap m$. So all of *P*'s verifiers contain $@ \sqcap m$. Hence *P* is true and reports every fact pertaining to *m*, and so is a complete truth with respect to *m* according to (Vertical). For the converse direction, assume that *P* is true and reports every fact pertaining to *m*. Let *Q* be a truth pertaining to *m*. Then every verifier of *Q* is part of *m*, and since *Q* is true, at least one of them is also part of @, and hence part of $@ \sqcap m$. But then since *P* reports every fact pertaining to *m*, *P* reports $@ \sqcap m$, hence every verifier of *P* contains $@ \sqcap m$ and thus, by the transitivity of part, the mentioned verifier of *P*.[32]

This truthmaker-based analysis of the notion of a complete truth avoids the difficulties discussed above for the possible worlds analysis.[33] Case (1), recall, was that of pure set theory. The problem was that under the possible worlds analysis, every truth comes out as complete truth with respect to the subject matter of pure set theory, since every truth pertaining to that subject matter is necessary, and thereby

[31] Like (Horizonal) and (Vertical) above, TM.RWT only sets a threshold to be passed by a truth in order to count as complete with respect to a subject matter, and thereby counts any truth stronger than a complete one also complete. To get at a notion of an exactly complete truth about a subject matter *m*, which is not allowed to be say true things about matters other than *m*, we might require the truth to be verified only by $@ \sqcap m$. But for the same reason as before, in the case of TM.UWT, this does not exclude variation in the falsifiers, and I shall once more set aside the question of whether, and how, one might uniquely characterize the appropriate falsifiers for *the* unique whole truth with respect to a subject matter.

[32] We noted in the context of the intensional account of the whole truth that the formal machinery used to represent subject matters—i.e. equivalence relations on the set of worlds—could also be used to represent levels of specificity. Under the truthmaker account, we are representing subject matters in a very different way, which will not be of help in capturing levels of specificity. Instead, we can do something similar to what we did in the possible worlds framework by taking the *congruence relations* on a state-space to represent levels of specificity. Roughly, these are *order-preserving* equivalence relations on the state-space.

[33] It is worth pointing out that the analysis does not make essential use of the most distinctive aspect of exact truthmaker semantics, i.e. the exactness requirement on verification: to determine whether a proposition is a whole truth (with respect to some subject matter), it is enough to know what its *inexact* verifiers are, which are exactly those states that contain an exact verifier as part. As we shall see below, though, in order to obtain a satisfactory account of the totality operator 'that's it', and thereby of the weaker sense of completeness described in section 2, we need to appeal to exact truthmaking. Thanks to an anonymous referee for raising this matter.

vacuously entailed by every proposition under the modal construal of entailment. In our diagnosis of the problem, we argued that the state s that the empty set is a member of its singleton is not wholly relevant for making true the proposition P that no empty set has members. So s is an actual part of the subject matter m of pure set theory, so $s \sqsubseteq @ \sqcap m$. But P does not report s, for s is not part of every, or indeed any verifier of P, since it is not wholly relevant to the truth of P. Hence P is true, and even a truth pertaining to the subject matter of set theory, but under the truthmaker approach it is not a complete truth with respect to that subject matter.

Case (2) was about the subject matter of the nature of Socrates, and the problem was that every truth that entails that Socrates exists is automatically classified by the possible worlds analysis as a complete truth with respect to that subject matter. In our diagnosis of the problem, we argued that the state that Socrates is human is not wholly relevant to making it true that Socrates exists. So the truth that Socrates exists does not report the fact that Socrates is human, and since that is a fact pertaining to the subject matter, that truth is correctly classified by the truthmaker approach as not a complete truth with respect to the nature of Socrates.

5 Do we still get too much for free?

The problem for the possible worlds analysis was that it gave away necessary truths and consequences for free, as it were, and that in many contexts it seemed that a putative complete truth should not get them for free. As we saw, the truthmaker analysis does not face the same difficulty. Just because a proposition is necessarily true, or a necessary consequence of a given putative complete truth P, it does not follow that it is inexactly entailed by P, and hence the truthmaker analysis does not give it to P for free. This is fine as far as it goes, one might respond, but it does not mean that there is not a narrower class of truths or consequences that even our account gives away for free, and for which contexts may be found in which a putative complete truth should not get these for free. This section is devoted to answering this concern.

The truthmaker analysis does indeed give away certain things for free. Just as the possible worlds analysis gives away for free all the necessary consequences of a putative complete truth, so the truthmaker analysis gives away for free any *inexact* consequences of a putative complete truth P. And just like the possible worlds analysis gives away for free all the necessary truths, there may also be a special sort of truth that the truthmaker analysis gives away for free. Recall that for any set of states, we may form their fusion. It is normally assumed in the context of the truthmaker framework that this holds even for the empty set of states. The result of fusing no states, as it were, is called the *nullstate*, which is the one state that is part of absolutely every state. Now say that a proposition is *trivial* iff it is verified by (perhaps among other things) the nullstate.[34] Then every proposition inexactly entails every trivial proposition. Hence any putative complete truth is automatically classified as entailing, in the relevant sense, any trivial proposition, and likewise it is automatically classified as reporting the nullstate.

Let me describe a few classes of examples that might at first glance appear problematic for my account. Firstly, there is the class of *conceptual entailments*. Plausibly, that Kant was a bachelor inexactly entails that Kant was unmarried. For presumably, the state of Kant's having been a bachelor is the same state as the state of Kant's having been an unmarried male, which surely contains as a part the state of Kant's having been unmarried. Now, while this particular example of a conceptual entailment is fairly obvious, there may well be conceptual, inexact entailments that are not obvious. In such cases, one might suspect that the failure to make explicit a relevant truth conceptually entailed by a truth P might well intuitively disqualify P from completeness.

Secondly, there are cases of *grounding*. At least under the semantics for grounding described in [5], [4], and [8], any case of (full) grounding is a case of inexact entailment. But truths of grounding can be highly non-obvious, for instance if the truths about a person's mental life are indeed fully grounded in physical truths. And presumably one can ask questions about a person's mental life such that no answer consisting purely of physical truths would intuitively be considered acceptably complete.

Thirdly, there may be *trivial logical truths*. Of course, most logical truths are not trivial under the truthmaker treatment. Indeed, it is easy to see that there is no way to form a trivial truth by application of the usual truth-functional operations to propositions which are neither trivially true nor trivially false. For when the application of any such operation brings new verifiers or falsifiers into play, these are always states properly containing states that were already in play as verifiers or falsifiers of the propositions to which the operation is applied.[35] There may, however, be other, reasonably natural operations on propositions that do allow this. The most plausible candidate I am aware of is the *incremental conditional* $P \rightarrow Q$ of [6, esp. §4].[36] One obtains a verifier of $P \rightarrow Q$ by first considering any function that maps every verifier s of P to a verifier of Q. For s verifying P, one then considers the smallest state t such that $s \sqcup t$ contains $f(s)$. Finally, one takes the fusion of all these smallest states. In the case of $P \rightarrow P$, taking the identity function that maps every verifier of P to itself then yields the nullstate as verifier of $P \rightarrow P$, so $P \rightarrow P$ is trivially true. But this might appear problematic. For to the extent that sometimes it is not admissible to omit $\forall x\, x = x$ from a complete truth, perhaps it is sometimes also not admissible

[34] Note that this is a not an epistemic notion of triviality; there is no obvious and reason for thinking that every proposition that is verified by the nullstate should be immediately recognizable as true by anyone who entertains it.

[35] Contrast this with the situation of the necessary truth in the possible worlds framework, which may be obtained from any proposition P by an application of negation and disjunction to form $P \vee \neg P$.

[36] Note that in that paper, Fine does not call the conditional in question 'incremental'. I am borrowing this apt term from [20], and Fine's exchange with Yablo on [19]; see [9] and [21].

to omit an instance of $P \rightarrow P$, and so a putative complete truth should not get that instance for free.

Finally, there may be *trivial non-factive grounding truths*. On the currently best developed account of iterated grounding claims, due to [16], true non-factive grounding claims are *zero-grounded*: grounded, but by the empty collection of facts or truths.[37] Still assuming a treatment of ground within truthmaker semantics along the lines described by Fine, a truth will turn out zero-grounded iff verified by the nullstate.[38] And surely there are contexts in which truths of non-factive grounding may not simply be omitted from a truth without sacrificing its completeness.

A minimal response one might give to this objection would be as follows: Yes, there are cases in which trivial truths or consequences do not come for free, but the truthmaker approach gives these away for free, and so in these cases the approach does not give the right results. But at least it can handle a lot more cases than its possible-worlds based rival, and so it still constitutes a step in the right direction, even if it does not quite get us where we would ideally want to be.

However, I think that a less concessive response is warranted. True, there are contexts in which trivial truths may not be omitted from a truth without rendering that truth incomplete in the contextually relevant sense. But it seems to me that that sense of completeness is different from the sense of completeness at issue in the cases I used to motivate the truthmaker approach, and that it is different in *kind*, not just in *degree*. There is, on the one hand, a wide range of applications in which the pertinent notion of a complete truth is such that, intuitively, a truth's being complete is a matter of the truth describing both the *entirety* of the relevant portion of reality, and doing so *in full detail*—in other words, a matter of the truth reporting every fact that is part of the relevant portion of reality. It is to these applications that the present account is intended to apply. And in these applications, trivial truths and entailments do come for free, since they can neither help report further parts of reality, nor can they add to the level of detail in which a portion of reality is being described.

In the contexts in which trivial truths and entailments do not come for free, completeness is not, or not purely, a matter of which parts of reality are being described and at what level of detail or specificity. Rather, in these contexts, completeness (also) demands that we answer, roughly speaking, for each of a relevant range of *representations* of reality, whether they represent reality accurately. To meet this demand, we may have to employ several accurate representations of the same parts of reality. So here we have a partly representational notion of completeness, whereas in the other cases we have what we may call a purely *worldly* notion of completeness.

---

[37] The idea that some truths might be grounded in this way is due to [4, p. 47f].

[38] In general, under the truthmaker semantical account, a truth P is (weakly fully) grounded by the collection of truths Γ just in case for every choice T of one verifier each from the members of Γ, the fusion $\bigsqcup T$ of the states in T verifies P (cf. [4, § 1.10],[8, pp. 685ff, 700ff]). In the case where Γ is empty, that means that $\bigsqcup \emptyset$ must verify P, and $\bigsqcup \emptyset$ is just the nullstate.

And it would be a mistake to try and capture these different notions within a single theory; they should be treated separately.[39]

With respect to the above examples, this response commits me to a disjunction: either the examples are not actually cases of inexact entailment, or they invoke a notion of a complete truth outside the scope of my account, targeting representational rather than purely worldly completeness. Which of the disjuncts obtains is, in at least some of the cases, a difficult question. For what it is worth, I incline towards the second disjunct with respect to the conceptual entailments and trivial logical truths. I suspect there might be more to be said for the first disjunct with respect to the grounding cases. But either way, the disjunction can be seen to be generally plausible irrespective of which disjunct one favours in each case. For once one has taken a view on the truthmakers of the relevant propositions under which the entailments in question are genuine inexact entailments, one has already taken a view under which the propositions merely provide different representations of the same portions of reality.

6 Totality operators

We return, finally, to the issue of totality operators and their connection to the notion of a complete truth. In order for Jack to give a complete answer to the question what he had for breakfast, recall, he need not mention that he did not eat a crocodile. He may just say that he had eggs, bacon, coffee, *and that's it*. Indeed, the final totality clause would normally be taken as understood. Similarly, at least in one important sense, the complete physical description of the world may be the whole truth about the world, even though it does not entail that there are no demons. It is sufficient if the result of appending 'and that's it' to the complete physical description of the world entails that truth.

Given the central role that totality operators thus seem to play in the formulation of many complete truths, a theory of the notion of a complete truth would appear incomplete if it lacked an account of these operators. While it is beyond the scope of this chapter to give a fully developed account, I shall attempt in this section to describe the broad outlines of a truthmaker-based treatment of totality operators, and to highlight a few advantages such a treatment seems to me to offer in comparison to possible-worlds based rivals.

§6.1 offers an account of the unrestricted totality operator. §6.2 briefly describes how the account may be extended to accommodate restricted totality operators. §6.3 considers an alternative possible-worlds based rival first proposed in [1] and developed further by [12], and argues that the present account is superior.

---

[39] A similar distinction between a worldly and a representational version of a notion is often made with respect to grounding. Here, too, a truthmaker based approach may be seen as adequate for the worldly, though perhaps not the representational notion; cf. e.g. [2, 3, 8, 11].

6.1 An Unrestricted Totality Operator

Our task is to define a one-place operation Δ on propositions that corresponds to the intuitive understanding of relevant uses of 'that's it', and that bridges the gap between a truth that is complete in the weaker sense exemplified by Jack's description of his breakfast and a truth that is complete in the stronger sense of inexactly entailing every truth pertaining to the relevant subject matter. Within the truthmaker framework, that task divides into two parts, that of specifying the verifiers and that of specifying the falsifiers of ΔP. We begin with the verifiers.

The basic idea is to take there to be a function δ on the states, such that for any given state s, the state δs is the state of s obtaining, *and that's it*. Assuming that there is a coherent notion of an unrestricted *that's it* in application to propositions, it is very plausible that there should be such a function. To see this, note that for any given state s, we may consider a proposition P verified only by s. It seems very plausible that ΔP should then also have just a single verifier. For if there is only one way for P to be true, it is hard to see how there could be several ways for it to be the case that P is true, and that's it. But then we may simply let δs be the state verifying ΔP, where P is verified only by s.[40]

In giving an account of the δ function, we may again pursue either a definitional or an axiomatic approach. Just as I did for the relation of relevant verification before, I shall here adopt the axiomatic approach. So I shall merely lay down some intuitively plausible principles connecting the δ function to the mereological and modal aspects of the state-space, draw out some of their consequences, but leave open the question of the definability of δ in other, independently understood terms.

We may start with three somewhat rough, intuitive ideas about the behaviour of δ that are suggested by its informal reading in terms of 'that's it'. First, since δs is the state to the effect that *s obtains*, and that's it, δs should contain s as a part. Second, since δs is the state to the effect that s obtains, *and that's it*, in some sense, δs should not contain anything more than s. Third, δs should be incompatible with any state it does not contain.

It is easy to see that, were it not for the hedge 'in some sense' in the second principle, the three principles would not be jointly satisfiable except in special cases, namely when s is already incompatible with any state it does not contain. The first and third principle are straightforward and will be taken for granted henceforth:

$\delta.1$ $s \sqsubseteq \delta s$
$\delta.2$ δs is incompatible with any state it does not contain

With respect to the second principle, we need to ask how it may be clarified. A natural thought is that δs should not be allowed to contain anything *positive* beyond what is contained in s, it should only be allowed to go beyond s in a purely *negative* way, namely by precisely ruling out the obtaining of any further positive facts not already

[40] Since there may be several propositions P verified solely by s but differing with respect to their falsifiers, I am also relying here on the assumption that the verifier of ΔP does not vary with the falsifiers in these cases. Again, this seems intuitively highly plausible.

contained in s. Let us take for granted, for the moment, an exclusive and exhaustive distinction among the states between wholly positive states and partly negative states. It is very plausible to suppose that (a) a state is partly negative whenever it has a partly negative part, and (b) that any fusion of wholly positive states is itself wholly positive. We may then define the positive part s^p of a state s as the fusion of all its wholly positive parts. In this terminology, the idea that δs may only go beyond s in a purely negative way may be captured as the claim that s and δs always have the same positive part. In addition, it seems plausible that the output of the δ-function should (a) depend only to the positive part of its input, and (b) differ whenever the positive parts of the input states differ. We thus arrive at two further plausible constraints on δ:

δ.3 $(\delta s)^p = s^p$
δ.4 $\delta s = \delta t$ iff $s^p = t^p$

Note that by setting $t = \delta s$ in δ.4, we can derive from these constraints the S4-principle for δ: $\delta s = \delta\delta s$.

We now turn to the question how the verifiers of ΔP may be obtained from δ and the verifiers of P when P has more than just one verifier. In effect, this is the question of how Δ should be taken to behave in application to a *disjunction*. For a proposition with multiple verifiers $s_1, s_2, \ldots$ is essentially the same as the disjunction of the propositions $P_1, P_2, \ldots$ verified, respectively, only by $s_1, s_2, \ldots$ From a formal point of view, the most natural option would be to take ΔP to be verified by δs whenever s verifies P, and by no other states. That option also seems to fit well with our intuitive judgements. To elicit clear intuitions about the matter, it is best to consider again some examples with an implicitly restricted subject matter. A fairly natural example of an application of ‘that’s it’ to a disjunction is as follows. Suppose we are asking Jack what he has had to drink for lunch—suspecting, perhaps, that he’s been on the booze. He might reply: I had two or three glasses of orange juice—I don’t quite remember—and that’s it! Then the most natural way of interpreting him would be as saying, in effect, that either he had two glasses of orange juice, and that’s it, or he had three glasses of orange juice, and that’s it. This is just as predicted by the suggested account, which I therefore propose we accept:

$$(\Delta P)^+ = \{\delta s : s \text{ verifies } P\}$$

It should be noted, though, that this makes it relatively easy for ΔP to come out true. All one needs to do is formulate a truth that, intuitively, has a sufficiently large subject matter that the entire actual world, or at least its positive part, is relevant to verifying the truth. And since relevant verification does not, on the truthmaker conception, imply any sort of *minimality*, this may be relatively easy. For instance, if we give the positive part of the world a name, say ‘Bill’, then it will presumably be true that Bill exists, and that’s it. Indeed, under a natural account of the truthmakers of existential quantifications, it will be true that something exists, and that’s it. While this is perhaps a somewhat odd or anomalous result, I don’t think it is too objectionable, or even too counter-intuitive. A statement such as that something exists, and that’s it, intuitively strikes one as odd and unhelpful rather than clearly false, it seems to me—and odd and unhelpful it is even under the present account. It is worth emphasizing, moreover, that the truth of ΔP is not sufficient for P to be classified as a *whole truth*. That ΔP is true guarantees only that @ verifies ΔP, whereas in order for P to be a whole truth, *every* verifier of ΔP must contain the actual world. This further condition is of course not met in the cases just mentioned.

We turn now to the question of the falsifiers of ΔP. There are three main constraints on how it may be answered. First, in order for to falsify ΔP, a state must be incompatible with every verifier of ΔP, otherwise it would be possible for ΔP to be both true and false. Second, to guarantee that in every possible world, ΔP is either true or false, we need to make sure that every world contains a falsifier of ΔP if it does not contain a verifier of ΔP. Third, a falsifier of ΔP should be wholly relevant to making ΔP false. Unfortunately, it is not easy to come up with an answer that satisfies all three constraints. We might start by considering the maximally liberal account of the falsifiers of ΔP on which every state incompatible with every verifier of ΔP is considered a falsifier of ΔP.

$$(\Delta P)^- \qquad = \{s : s \text{ is incompatible with } \delta t \text{ whenever } t \text{ verifies } P\}$$

This answer evidently satisfies the first constraint, and it is also easily shown to satisfy the second constraint.[41] The problem is the third constraint, demanding that a falsifier be wholly relevant. Generally speaking, a purely modal condition like the one of incompatibility is not sufficient to guarantee a relevance connection.

Actually, it turns out that the tension of the proposal with the relevance constraint is less dramatic than one might expect. Consider first the case in which all of the verifiers of ΔP are consistent, and hence possible worlds. Then for a state to be incompatible with each of them is equivalent to its not being a part of any of them. And it seems quite plausible that any such state is relevant to making it false that ΔP. Somewhat metaphorically, we might think of any state as saying of itself that it is a fact, and of a state δs as additionally saying of itself that it contains every fact. Thus, any state t not contained in δs presents itself, wholly relevantly, as a counter-example to the claim we take δs to make.[42]

Unfortunately, this defence of the liberal account does not extend to the case in which all verifiers of ΔP are inconsistent and hence every state falsifies ΔP. The perhaps most counterintuitive kind of case we obtain is when a consistent state s falsifies such a proposition ΔP, while s is also part of every verifier of ΔP, and intuitively has nothing to do with the inconsistency of those verifiers. Thus, suppose we are considering the subject matter of the properties of a ball, and P is the proposition that the ball is red all over, green all over, and round. Then the state of the ball being round is a falsifier of ΔP even though intuitively, it plays no part in making it false

---

[41] Suppose w is a world which does not contain a verifier of ΔP. Since w is a world, it is incompatible with every state it does not contain, and so it is incompatible with every verifier of ΔP.

[42] In light of this, one might be tempted to say that we should take the falsifiers of ΔP to be simply those states that are not contained in any verifiers of ΔP. But the problem with this suggestion is that ΔP may be false and yet there may not be any states not contained in any verifiers of ΔP, for instance when ΔP is verified by the fusion of all states. So under this approach, we would be violating the second of the above constraints on the set of falsifiers of ΔP.

that $\varDelta P$. So it cannot be claimed that the liberal account fully captures the intuitive understanding of ‘that’s it’.

We are left with three options. The first is to try to somehow refine the liberal account to rule out the irrelevant falsifiers by appeal to their modal and mereological properties. The second is to assume more structure in the state-space. Specifically, we might take as given a relation of *wholly relevant exclusion* on the states and construct the falsifiers of $\varDelta P$ by choosing a relevant excluder for each verifier of $\varDelta P$ and building their fusion.[43] Both these options introduce a significant degree of additional complexity into the theory. A third option might therefore be to endorse the liberal account, and propose $\varDelta$, so interpreted, merely as a reasonably close approximation to the intuitive interpretation of ‘that’s it’. If the best version of the first two options turns out to yield a significantly more complex and messy theory, there may be good pragmatic justification for the third option.

Before we move on, let me very briefly comment on some aspects of the logical profile of $\varDelta$. Given our constraints on the δ-function, the account of the verifiers of $\varDelta P$, and either the liberal or the exclusionary account of the falsifiers, we can derive the following very plausible principles:[44]

$\varDelta$.1 $\varDelta P$ inexactly entails P
$\varDelta$.2 $\varDelta P = \varDelta\varDelta P$
$\varDelta$.3 $\varDelta(P \vee Q) = \varDelta P \vee \varDelta Q$
$\varDelta$.4 $\varDelta P$, Q modally entail $\varDelta(P \wedge Q)$

I hope to study the logic of $\varDelta$ under the present truthmaker account more fully in future work.

6.2 Restricted Totality Operators

To accommodate totality operators restricted to a particular subject matter, we may work with a binary δ function that takes a subject matter as its second argument, so that $\delta(s, m)$ is the state to the effect that *as far as m is concerned*, s obtains, and that’s it. The unrestricted, unary δ function may then be defined in terms of the binary one by setting m to the maximal subject matter, i.e. the fusion of all states. In developing this approach further, there are a number of choices to be made, and it is not always

[43] Given the relation of relevant exclusion, we might then move to a unilateral conception of propositions as given simply by a set of verifiers, and characterize negation in general in the way that we just proposed to treat the negation of $\varDelta$-propositions. This alternative, exclusionary treatment of negation is discussed in more detail in [7, pp. 634f, 658ff].

[44] $\varDelta$.1 is immediate given the clause for the verifiers of $\varDelta P$ and constraint δ.1. That $(\varDelta P)^+ = (\varDelta\varDelta P)^+$ is immediate given the S4-principle for δ already mentioned above, and that $(\varDelta(P \vee Q))^+ = ((\varDelta P) \vee (\varDelta Q))^+$ is likewise immediate given the verifier clause for $\varDelta$. But under both the liberal and the exclusionary account of falsifiers, sameness of verifiers implies sameness of falsifiers, so we also get the bilateral identities $\varDelta$.2 and $\varDelta$.3. For $\varDelta$.4, note that we may derive from δ.3 and δ.4 that $\delta s = \delta t$ whenever $s \sqsubseteq t \sqsubseteq \delta s$. Combining this with δ.3, we may show that $\delta s \sqcup t = \delta(s \sqcup t)$ whenever $\delta s \sqcup t$ is consistent, from which $\varDelta$.4 is immediate.

obvious in advance of detailed theorizing which of them is best. My aim here is simply to sketch the outlines of one natural approach we might take. While it may not ultimately be the best one, it provides some grounds for thinking that there is nothing in principle standing in the way of developing, within the overall framework here described, a satisfactory theory of restricted totality operators.

We shall assume that binary $\delta(s,m)$ will only be defined when $s \sqsubseteq m$, i.e. when s pertains to m. We may then lay down the following counterparts to (δ.1)–(δ.3) for $s \sqsubseteq m$:

(δ_R.1) $s \sqsubseteq \delta(s,m)$
(δ_R.2) $\delta(s,m)$ is incompatible with any part of m it does not contain
(δ_R.3) $\delta(s,m)^p = s^p$

These straightforwardly entail (δ.1)–(δ.3) under the proposed definition of unary δ.

It is not as obvious what we should say about the conditions for the identity $\delta(s,m) = \delta(t,n)$ to hold. It is very plausible that the identity should only hold if $s^p = t^p$, since whatever the subject matter, δ is not allowed to add anything positive to the state it presents as total. It is also very plausible that if $s^p = t^p$ and $m = n$, then the identity should hold. But there may also be cases in which the identity holds even though $m \neq n$. Thus let s be the state that the ball is red, m the fusion of the ball's being red and the ball's being blue, and n the fusion of m and the ball's being green. Then it seems plausible that $s = \delta(s,m)$, since s is already incompatible with every part of m it does not contain. But likewise it would seem, for the same reason, that $s = \delta(s,n)$, even though $m \neq n$.

Now it may be felt that m is not a good candidate for a subject matter in an intuitive sense. For the only subject matter it could represent, one might think is that of the colour of the ball, and that subject matter seems better represented by n. But it is not immediately obvious how one might characterize those states that represent genuine subject matters, and even if this can be done, it may be desirable to allow $\delta(s,m)$ to be defined even when m is not a subject matter in this narrower sense. If so, then we may still impose the separate necessary and sufficient conditions for the identity of restricted totality-states just proposed:

(δ_R.4a) $\delta(s,m) = \delta(t,n)$ if $n = m$ and $s^p = t^p$
(δ_R.4b) $\delta(s,m) = \delta(t,n)$ only if $s^p = t^p$

Note that (δ_R.4a) implies (δ.4) under the proposed definition of unrestricted δ.

In addition to defining unary δ in terms of binary δ in the way indicated, we may also be able to define the notion of a partly negative state in terms of binary δ, and thereby further reduce the number of primitives assumed in our theory. In informal terms, the idea is to assume that totality-states are the only source of negativity, so that a state is partly negative if and only if it contains a totality-state. However, we need to be careful when making this precise. Since under the present account, every state s equals $\delta(s,s)$, we must not interpret the notion of a totality-state invoked in the informal idea as simply that of a value of δ for some arguments. Instead, let us call a state s a *proper totality-state* iff for some states $t \sqsubseteq m$, $s = \delta(t,m)$ and $t \neq s$. A proper totality-state can therefore be obtained by adding *that's it* to some *other*, strictly smaller state, and it must therefore be considered partly negative. We may then consider defining the set of partly negative states as the set of states containing some proper totality-state as a part.

If the definition is to work as intended, we need to make sure that the set of partly negative states, so defined, satisfies the two assumptions we made above. That is, we need to make sure that any state that contains a partly negative state is itself partly negative, and that any fusion of some wholly positive states is itself wholly positive. The first assumption is a trivial consequence of the proposed definition. The second is not, so we lay down as a further constraint on δ that

(δ_R.5) $s \sqcup t \sqcup \ldots$ contains a proper totality-state only if one of $s, t, \ldots$ does

Let us finally define a subject-matter relative totality operation on propositions in terms of binary δ. Just as we took $\delta(s,m)$ to be defined only when s pertains to m, we shall define $\Delta(P,m)$ only when P pertains to m, i.e. when the fusion of all verifiers and all falsifiers of P is a part of m.[45] The obvious adaptation of the previous verifier-clause, and liberal falsifiers-clause, to the restricted case is then as follows:

$$(\Delta(P,m))^+ = \{\delta(s,m) : s \text{ verifies } P\}$$
$$(\Delta(P,m))^- = \{s \sqsubseteq m : s \text{ is incompatible with } \delta(t,m) \text{ whenever } t \text{ verifies } P\}$$

The falsifier-clause raises the same worries that its unrestricted counterpart does, but as far as I can see it does not raise any new ones.[46]

6.3 Comparison with Intensional Approaches

It is instructive to compare the present, truthmaker-based account of Δ with an alternative, possible-worlds based approach that was first proposed by David Chalmers and Frank Jackson ([1, p. 317f]) and then developed in formal detail by Stephan Leuenberger ([12]). In the possible worlds framework, we need to say at which worlds ΔP is true, given the information at which worlds P is true. The proposal that Chalmers and Jackson make is that a world w verifies ΔP iff (a) w is a P-world, and (b) among the P-worlds, w is *minimal* with respect the relation of *outstripping*, which is a roughly parthood-like relation among the worlds. More precisely, a world w is said to outstrip another world v just in case w contains as a part an intrinsic *duplicate* of v, but v does not contain a duplicate of w. The informal idea is simple and

[45] Note that $\Delta(P,\blacksquare)$ will then always be defined, and we may again define the unary, unrestricted Δ in terms of the binary version by setting the subject matter to the fusion of all states.

[46] The falsifier-clause illustrates the importance of requiring that m contain all the falsifiers as well as all the verifiers of P. Since any proposition is required to have a non-empty set of falsifiers, we need to ensure that m always contains some state incompatible with $\delta(t,m)$ whenever t verifies P if the clause is to be acceptable. Since any falsifier of P is incompatible with every verifier of P, and hence with $\delta(t,m)$ whenever t verifies P, the requirement that m contain the falsifiers of P ensures this, but absent this requirement, we would have no such guarantee.

clear enough: in order for a world w to make true ΔP, w must make true P, and w must be 'as small as possible' consistent with making P true, so that if we were to remove any of w's parts, the result would no longer make P true.

Our own account conforms to a very similar idea: in order for a state s to verify ΔP, s must make P true in the sense of *containing* a verifier of P, and the entire positive part of s must be relevant to verifying P. Here the appeal to the positive part of a state serves a similar purpose as the appeal to duplication in Chalmers and Jackson. Chalmers and Jackson appear to think of worlds on a broadly Lewisian model as a spatio-temporally maximal concrete entity. A world, so construed, makes true a negative truth such as that there are no unicorns firstly, by not containing unicorns, and secondly, by being maximal, and therefore not part of a bigger entity that might include unicorns. Since a world is maximal, it cannot itself be a part of another world, only an intrinsic duplicate of one. In the truthmaker framework, worlds are seen as maximally consistent states of affairs rather than maximal spatio-temporal entities. A world w, so construed, makes true a negative truths such as that there are no unicorns by containing a partly negative part that is wholly relevant to there not being a unicorn. These parts of w of course need not be wholly relevant to verifying P in order that w verify ΔP, hence the need for the restriction to the positive part of w. We can use this observation to define a relation of outstripping on the worlds in the truthmaker framework, by saying that a world w outstrips a world v iff w^p properly contains v^p. It is then straightforward to show, given our constraints on the δ-function, that whenever δs is consistent, δs is the minimal element with respect to outstripping among all the worlds that contain s.[47]

From the truthmaker perspective, the crucial difference between the Chalmers-Jackson account and our own is at which the stage the minimality condition is applied. Let $s, t, \ldots$ be the verifiers of P, and assume for ease of comparison that $\delta s, \delta t, \ldots$ are all consistent and hence possible worlds. To obtain the worlds at which ΔP is true under the Chalmers-Jackson account, we first form, for each state among $s, t, \ldots$ the set of worlds containing that state, then build the union of these sets of worlds to obtain the set of worlds at which P is true, and finally apply the minimality condition to select the minimal elements among these. To obtain the worlds at which ΔP is true under our own account, we first form, for each state among $s, t, \ldots$ the set of worlds containing that state, then apply the minimality condition to obtain the minimal elements of each such set, and finally take the union of all the resulting (singleton) sets.[48]

It seems to me that the second approach yields a more intuitive view of the truth-conditions of ΔP. Consider again the question of Jack's breakfast, and suppose that

[47] More generally, δs is always the unique minimal element with respect to outstripping among all the δ-states that contain s as a part. For suppose that $s \sqsubseteq \delta t$. Then $s^p \sqsubseteq (\delta t)^p$, and since by ($\delta$.3), $(\delta s)^p = s^p$, we have $(\delta s)^p \sqsubseteq (\delta t)^p$. If $(\delta s)^p \sqsubset (\delta t)^p$ then δt outstrips δs. If not, then $(\delta s)^p = (\delta t)^p$, and we may infer by ($\delta$.3) that $s^p = t^p$, and by (δ.4) that $\delta s = \delta t$.

[48] Note that under a possible worlds approach, we have only the information at which worlds P is true to start with, so there is only one stage at which the minimality condition can be applied. It is because the truthmaker approach also takes into account which parts of the P-worlds are wholly relevant to P, i.e. which states are exact truthmakers of P, that it becomes possible to apply the minimality condition at a different and, as I shall argue, more appropriate stage.

Jack had eggs, bacon, and nothing else. Suppose, however, that he gives the following answer: I had eggs, I had bacon, and I either had coffee or I didn't have coffee, and that's it. Then the intuitive account of his statement would seem to be as follows. What Jack said is strue, but it is *not* a strongly whole truth. For it does not entail the truth that Jack did not have coffee. Indeed, it seems to *explicitly* leave open the possibility that he did have coffee. The truthmaker approach is completely in agreement with this intuitive assessment. Let us abbreviate the sentences to which Jack's 'that's it' is applied as $(E \wedge B) \wedge (C \vee \neg C)$. The exact verifiers of that sentence include both a state s verifying $(E \wedge B) \wedge C$ and a state t verifying $(E \wedge B) \wedge \neg C$. As a result, the exact verifiers of $\Delta((E \wedge B) \wedge (C \vee \neg C))$ thus include δs—a minimal world in which Jack had eggs, bacon, and coffee—and δt—a minimal world in which Jack had eggs, bacon, and no coffee. Hence $\Delta((E \wedge B) \wedge (C \vee \neg C))$ does not inexactly or even modally entail the truth that Jack did not have coffee, and hence is not classified as a strongly complete truth.

But suppose Jack had said instead: I had eggs and bacon, and that's it. The intuitive assessment would then be different. Given that he did indeed have eggs, bacon, and nothing else, that statement would intuitively count as a strongly complete truth, as it would entail (among other things) that Jack did not have coffee. Again, this is the result that we get under the truthmaker approach. For none of the exact verifiers of the embedded statement $E \wedge B$ say anything about coffee. Suppose s is such a verifier. Then applying δ to s will yield a minimal world in which Jack had eggs and bacon, and hence a world in which Jack did not have coffee. After all, any state c to the effect that he did have coffee would be a positive state, and since no such state is part of s^p, by our constraint (δ.3), no such state is part of $(\delta s)^p$ either.

But since the two embedded statements $E \wedge B$ and $(E \wedge B) \wedge (C \vee \neg C)$ are logically equivalent, they are true at the same worlds. Under any possible-worlds approach, the content of ΔP can depend only on which worlds P is true at, and hence the statements $\Delta((E \wedge B) \wedge (C \vee \neg C))$ and $\Delta(E \wedge B)$ will in turn be true at exactly the same worlds. So in contrast to the truthmaker approach, no possible-worlds approach can respect the intuitive difference between the two statements. Instead, the Chalmers-Jackson approach counts $\Delta((E \wedge B) \wedge (C \vee \neg C))$ a strongly complete truth, since no world in which Jack had coffee is minimal among the worlds at which $(E \wedge B) \wedge (C \vee \neg C)$ is true. So just as the hyperintensional, truthmaker based account of the notion of a complete truth was seen above to be superior to intensional, possible-worlds based ones, so the hyperintensional, truthmaker based semantics for the totality operator turns out to be superior to intensional, possible-worlds based ones.[49]

[49] It should be noted that while the notion of a strongly complete truth, as pointed out before, is sensitive only to differences in inexact verifiers, the totality operators and thereby the notion of a weakly complete truth is sensitive even to difference concerning only exact verifiers. Suppose Jack had said that he had eggs and bacon, or eggs and bacon and coffee, and that's it—formally, $\Delta((E \wedge B) \vee ((E \wedge B) \wedge C))$. Then his statement is verified both by a minimal world in which Jack had eggs and bacon, and by a minimal world in which Jack had eggs, bacon, and coffee. Intuitively, this is the right result. For as before, Jack's statement explicitly leaves open the possibility that he had coffee, whereas the statement that he had eggs and bacon, and that's it, does not.

It might be objected that the intuitive difference between $\Delta((E \wedge B) \wedge (C \vee \neg C))$ and $\Delta(E \wedge B)$ is perhaps a purely pragmatic difference, not a semantic one, and thus that it is only the content that is communicated, and not the content semantically expressed, that is compatible with Jack's having had coffee. Although I know of no decisive objection to that view, it does not strike me as particularly plausible. If we imagine Jack attempting to cancel the putative implicature, for example by adding 'By saying this, I do not mean leave it open that I had coffee', then this would seem to me to simply warrant the response: 'Well, maybe you did not mean to, but leave it open you did!' However, for the purposes of this chapter, I can allow that there may be an alternative way of accommodating the intuitive data with respect to this example. The truthmaker approach is motivated in the first instance by the considerations of section 3 to the effect that necessary truths and consequences often do not come for free. And it still speaks in favour of the approach that it can handle the present examples naturally and elegantly, without the need to appeal to pragmatic interference, even if the account it offers is not obviously the only possible account.

7 Conclusions

We are typically interested not just in giving true descriptions of a subject matter, but in giving descriptions that are true and complete: the whole truth. In this chapter, I have criticized extant, intensional accounts of the relevant notion of completeness and tried to develop a better, hyperintensional one, utilizing the framework of truthmaker semantics. According to the view I propose, a truth is complete with respect to a given subject matter if and only if every fact pertaining to the subject matter is wholly relevant to making the truth true.

In this strong sense, a truth is complete only if it (relevantly) entails even the negative truths concerning the subject matter in question. In many contexts, it seems appropriate to apply a weaker standard, under which only positive truths have to be implied by the truth in question. The gap between weakly and strongly complete truths may be bridged by means of the totality operator 'that's it'. A truth is weakly complete just in case the result of applying 'and that's it' to the truth yields a strongly complete truth. I have sketched a truthmaker semantics both for an absolute and a subject-matter relative totality operator. While it remains conceptually close in some ways to a previous account due to Chalmers and Jackson, we saw that the move to a hyperintensional framework once more allows in a very natural way to better capture the intuitive understanding of the operator. It remains to work out the semantics in full detail and to determine the logic of the totality operator. But that is a task I have to leave for another occasion.

Acknowledgements For very helpful comments and discussion, I would like to thank Kit Fine, Stephan Leuenberger, Stefan Roski, Josh Schechter, Bruno Whittle, two anonymous reviewers, the editors of this collection, and the audiences at the Glasgow-Melbourne workshop on *Formal Methods* (2017), the workshop on *Hyperintensional Logics and Truthmaker Semantics* in Ghent

(2017), the *Joint Session* in Oxford (2018), and the *GAP.10* (2018) conference. Part of my work on this chapter was funded by the DFG through my Emmy Noether grant on *Relevance* (KR 4516/2-1), and I gratefully acknowledge that support.

References

1. Chalmers, David J. and Frank Jackson. 2001. Conceptual Analysis and Reductive Explanation. *The Philosophical Review 110*:315–361.
2. Correia, Fabrice. 2010. Grounding and Truth-Functions. *Logique et Analyse 53(211)*:251–279.
3. Correia, Fabrice. 2017. An Impure Logic of Representational Grounding. *Journal of Philosophical Logic 46(5)*:1–32.
4. Fine, Kit. 2012a. Guide to Ground. In *Metaphysical Grounding*, eds. Benjamin Schnieder and Fabrice Correia, 37–80. Cambridge: Cambridge University Press.
5. Fine, Kit. 2012b. The Pure Logic of Ground. *The Review of Symbolic Logic 5(1)*:1–25.
6. Fine, Kit. 2014. Truth-Maker Semantics for Intuitionistic Logic. *Journal of Philosophical Logic 43*:549–577.
7. Fine, Kit. 2017a. A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic 46*:625–674.
8. Fine, Kit. 2017b. A Theory of Truthmaker Content II: Subject-matter, Common Content, Remainder and Ground. *Journal of Philosophical Logic 46*:675–702.
9. Fine, Kit. ms. Review of Steve Yablo's Aboutness. Unpublished manuscript.
10. Groenendijk, Jeroen and Martin Stokhof. 1984. *Studies in the Semantics of Questions and the Pragmatics of Answers*. Ph. D. thesis, University of Amsterdam. http://dare.uva.nl/en/record/123669.
11. Krämer, Stephan. 2018. Towards a Theory of Ground-Theoretic Content. *Synthese 195(2)*:785–814.
12. Leuenberger, Stephan. 2014. Total Logic. *The Review of Symbolic Logic 7(3)*:529–547.
13. Levesque, Hector J. 1990. All I know: A Study in Autoepistemic Logic. *Artificial Intelligence 42(2–3)*:263–309.
14. Lewis, David K. 1988a. Relevant Implication. *Theoria 54*:162–174.
15. Lewis, David K. 1988b. Statements Partly About Observation. *Philosophical Papers 17(1)*:1–31.
16. Litland, Jon E. 2017. Grounding Ground. *Oxford Studies in Metaphysics 10*:279–316.
17. McCarthy, John. 1980. Circumscription—A Form of Non-Monotonic Reasoning. *Artifical Intelligence 13*:27–39.
18. McCarthy, John. 1986. Applications of Circumscription to Formalizing Common Sense Reasoning. *Artifical Intelligence 28*:89–116.
19. Yablo, Stephen. 2014. *Aboutness*. Princeton: Princeton University Press.
20. Yablo, Stephen. 2016. Ifs, Ands, and Buts: An Incremental Truthmaker Semantics for Indicative Conditionals. *Analytic Philosophy 57(3)*:175–213.
21. Yablo, Stephen. 2018. Reply to Fine on Aboutness. *Philosophical Studies 175*:1495–1512.

Journal of Philosophical Logic
https://doi.org/10.1007/s10992-022-09695-z

That's It! Hyperintensional Total Logic

Stephan Krämer[1]

Received: 3 March 2022 / Accepted: 19 December 2022

Abstract

Call a truth complete with respect to a subject matter if it entails every truth about that subject matter. One attractive way to formulate a complete truth is to state all the relevant positive truths, and then add: *and that's it*. When the subject matters under consideration are non-contingent, a non-trivial conception of completeness must invoke a hyperintensional conception of entailment, and of the completion operation denoted by 'that's it'. This paper develops two complementary hyperintensional conceptions of completion using the framework of truthmaker semantics and determines the resulting logics of totality.

Keywords Truthmaker semantics · Hyperintensionality · Totality operator · Total logic · Completeness · The whole truth

1 Introduction

Among all the truths concerning a given subject matter we may distinguish between those that constitute merely part of the overall truth concerning that subject matter, and those that exhaust the subject matter: the *complete truth* or truths about that subject matter. To a first approximation, we might take a truth to be *(strongly) complete* (with respect to a given subject matter) iff it entails every truth (pertaining to that subject matter). Totality operators, which may informally be glossed by the phrase '…and that's it', are an important resource in articulating complete truths. For a toy example, consider the subject matter of my breakfast. Suppose I had porridge and tea, and nothing else. In particular, then, I did not have eggs, bacon, coffee, nor roasted crocodile, … In a straightforward sense, these are truths pertaining to the subject matter, so a complete truth about that subject matter would need to entail them. Rather than list them all, it seems, I may say I had porridge and tea, *and that's it*, and achieve the same effect. The advantages of proceeding in this way—roughly

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

[1] Department of Philosophy, University of Hamburg, Überseering 35, Hamburg, 22297, Germany

speaking, by conjoining the *positive* truths about the subject matter, and adding that that's all—do not merely concern efficiency. Let us call a truth P *weakly complete* with respect to a subject matter iff it falls short of completeness only in the way just illustrated, by omitting some negative truths, so that the proposition that P, and that's it, is plausibly a strongly complete truth with respect to that subject matter. It seems that we may learn something significant about a given subject matter by determining what kind of truth is weakly complete with respect to it. An attractive first idea for explicating physicalism, for instance, is as the claim that the conjunction of all physical truths is a complete truth. But if physicalism is true, then, presumably, it is also true that there are no angels, and this truth does not appear to be entailed by the conjunction of all physical truths. The sense in which the latter might be complete, accordingly, is the sense of weak completeness rather than strong.[1]

In previous work [6], I argued that some applications call for a *hyperintensional* understanding of these notions, and proposed an explication of the notion of a whole truth within the framework of truthmaker semantics. The present paper employs the same framework to develop a formal semantics for suitable hyperintensional totality operators, and determines the resulting propositional logics of totality. I begin by describing the main motivations for a hyperintensional approach (Section 2). After introducing the general framework of truthmaker semantics (Section 3), I use it to develop a detailed account of two natural and complementary totality operators (Section 4), describe the resulting logics of totality (Section 5), and compare them to the intensional logic of totality studied by [8] (Section 6). An appendix gives proofs of soundness and completeness.

2 Motivating a Hyperintensional Approach

A strongly complete truth, we said, is a truth that entails every truth. A weakly complete truth still has to entail every positive truth.[2] A totality operator is one that takes any weakly complete truth into a strongly complete one. These notions are of very general application. More or less any question that might be raised can be regarded as determining a subject matter, and thus as determining a collection of all truths, and a collection of all positive truths, pertaining to that subject matter. We can then ask whether a given proposition constitutes a correct and weakly or strongly complete answer to the question according as it entails all elements of either collection. Correspondingly, in any such context, there is application for some form of totality operator turning a weakly complete truth into a strongly complete truth.

Let me give some examples. I have already mentioned the role of totality operators in the formulation of physicalism and related doctrines concerning the *kind* of truth that is weakly complete. Another example in metaphysics concerns the grounds of

[1] On these issues, cf. e.g. [1, pp. 317f, 8, pp. 529f]

[2] We shall see below that there are reasons to regard this condition as necessary but not quite sufficient for weak completeness.

true universal quantifications. Suppose that everything is concrete. On one natural account of the grounds of this truth, it is grounded in the truths that a is concrete, b is concrete, ... together with the truth that a exists, b exists, ..., *and that's it*. Another context in which totality operators may have application is the theory of non-monotonic reasoning, when inferences are to be based on the totality of one's evidence. Suppose I know that tweety is a bird. If that's *all* I know about tweety, I may be entitled to infer that tweety can fly, but I am not entitled to that inference if I also know that tweety is a penguin, and that penguins don't fly. There is also both explicit and implicit use being made of totality operators in natural language, especially in question-answer discourse. Suppose you ask me what's in my backpack, and I say: a bottle of water and an apple. Under the most natural interpretation of my utterance, it excludes the possibility of there being additional items in my backpack. This interpretation thus appears to invoke a tacit totality operator, as in: my backpack contains a bottle of water, an apple, and that's it.[3]

Depending on how we understand the notion of entailment invoked in the characterization of completeness, we obtain different versions of our triad of concepts. If we understand entailment in purely modal terms, for example, we obtain a liberal notion of a complete truth as a truth that necessitates every truth. If we invoke a more demanding conception of entailment, we obtain a likewise more demanding notion of a complete truth. I do not want to claim that there is a uniquely correct or best way to go here; which notion of entailment is the most appropriate one may simply depend on our particular purposes and interests. What I do want to claim is that some legitimate and reasonable purposes and interests call for a hyperintensional conception of completeness and hence of entailment.

Perhaps the clearest and most compelling reason to want an alternative, hyperintensional account of strongly and weakly complete truths is that some version of these notions should apply in an appropriately discriminating way to *non-contingent* subject matters. Obvious examples are various mathematical subject matters. Many statements about mathematical objects are necessarily true if true at all. Objects of pure mathematics are typically held to exist necessarily if at all. Many important questions about their natures are likewise plausibly non-contingent, such as the question whether they are abstract, whether they depend on the mathematical structures they are elements of, and whether they are mind-independent. In the modal sense of entailment, truths about these matters are thus vacuously entailed by any truth, and so any truth whatsoever will be regarded as complete with respect to this subject matter. But there surely is a non-trivial sense in which a correct description of certain mathematical objects may be, or fail to be, complete concerning the questions under discussion. For example, consider the question of the integer solutions of the

[3]There is an extensive body of literature in linguistics about this phenomenon, often invoking so-called *exhaustivity* operators which are somewhat similar to the totality operator studied by Chalmers & Jackson and Leuenberger. Spector [9] provides a helpful overview and comparison of various proposed exhaustivity operators. A proper engagement with this literature is beyond the scope of this paper, but it is worth mentioning that the existing approaches in linguistics are usually set within the intensional framework of possible worlds and thus not suited to address the specific problems that motivate my hyperintensional approach.

equation $x^2 - 4 = 0$. The answer that 2 and -2 are solutions, and that's it, is correct and complete. The answer that 2 is a solution, though necessarily equivalent, is correct but not complete.[4]

In response, one might perhaps consider an interpretation of entailment as *logical* rather than metaphysical necessitation. But similar problems then arise for logically non-contingent subject matters. Even if the necessity of identity is a logical truth, a correct theory of identity may in a very natural sense fail to be complete if it fails to *state* the necessity of identity.[5] More could undoubtedly be said about these issues, but the present remarks may suffice at least to motivate interest in the project of developing a more discriminating, hyperintensional theory of the whole truth, and a corresponding totality operator.

3 Truthmaker Semantics

In truthmaker semantics, a proposition is modelled not by the set of worlds at which it is true, but by the set of *states* that *make* it true, as well as the set of states that make it false.[6] Both notions, that of a state and that of truthmaking, require some comment. Unlike a possible world, a state may be *partial* in the sense that it settles the truth-value of only some propositions. Unlike a possible world, it may also be *impossible* in the sense that it makes an impossibility, even a contradiction true. Like a possible world, however, a state is required to be (relatively[7]) *specific*: it can only make a disjunction of the language true by making one of its disjuncts true. Thus, the truthmakers of the proposition that the ball is red or blue include the state of the ball being red and the state of the ball being blue, but we do not recognize a further disjunctive state of the ball being blue-or-red.[8]

In order to make a given proposition true, a state must be *wholly relevant* to the truth of that proposition. In particular, it is not enough that the state's obtaining *necessitates* the truth of the proposition. For instance, the state of snow being white does not make true the proposition that 2+2=4, since it is wholly irrelevant to the truth of that proposition. And the state of it being cold and rainy is not a truthmaker for the proposition that it is cold, since it is *partially irrelevant* to the truth of that proposition, by virtue of containing the irrelevant part of it being rainy. Analogous remarks apply to falsitymaking. I shall also call truthmakers verifiers, and falsitymakers falsifiers, while stressing that any epistemic connotations of these terms should be thoroughly discarded.

[4]Thanks to an anonymous referee for the example.

[5]These and similar points are developed at greater length in Section 3 of [6].

[6]For a more detailed exposition of the framework, see [4].

[7]We need not understand the specificity requirement in an absolute sense, but may instead take it as relating to the language that we wish to interpret. The requirement is then that each state be specific enough that it makes a disjunction true only if it makes one of its disjuncts true.

[8]As indicated in the previous footnote, we may take there to be such a thing as *the* state of the ball being red as long as the language under discussion does not provide the resources to discriminate between the different shades of red.

The states are taken to form a set S ordered by part-whole ($\sqsubseteq$), and we assume that given any states, we may form their *fusion*, which will also be a state. More formally, recall that a partial order on a set S is any binary relation on S that is transitive, reflexive, and anti-symmetric. We call a partial order on S *complete* iff every subset T of S (including the empty set) has a least upper bound with respect to that order.[9] Then our basic structure is that of a *state-space*:

Definition 1 A *state-space* is a pair $(S, \sqsubseteq)$ where

1. S is a non-empty set
2. $\sqsubseteq$ is a complete partial order on S

If $(S, \sqsubseteq)$ is a state-space and $T = \{t_1, t_2, \ldots\} \subseteq S$, we denote the least upper bound of T w.r.t. $\sqsubseteq$ by $\bigsqcup T$ or $t_1 \sqcup t_2 \sqcup \ldots$ and also call it the *fusion* of (the members of) T. When $T = \emptyset$, $\bigsqcup T$ is called the *nullstate*, or $\square$, which is part of every state. When $T = S$, $\bigsqcup T$ is called the *fullstate*, or $\blacksquare$, which has every state as part. The greatest lower bound of $T = \{t_1, t_2, \ldots\}$ w.r.t. $\sqsubseteq$—the greatest common part of the elements of T—will be denoted $\bigsqcap T = t_1 \sqcap t_2 \sqcap \ldots$[10]

Definition 2 Let $(S, \sqsubseteq)$ be a state-space.

- A *unilateral proposition* on $(S, \sqsubseteq)$ is any non-empty subset P of S
- A *bilateral proposition* on $(S, \sqsubseteq)$ is any pair $\mathbf{P} = (\mathbf{P}^+, \mathbf{P}^-)$ of unilateral propositions

We think of the members of a unilateral proposition P as its verifiers, and of the members of $\mathbf{P}^+$ ($\mathbf{P}^-$) as the verifiers (falsifiers) of a bilateral proposition $\mathbf{P}$.

We turn to the definition of the boolean operations. We take a disjunction, as one might expect, to be verified by exactly those states that verify one of the disjuncts. A conjunction we take to be verified by exactly those states that may be obtained by taking the fusion of a verifier of the one conjunct and a verifier of the other conjunct. Dually, we take a disjunction to be falsified by the fusions of falsifiers of the disjuncts, and a conjunction to be falsified by the falsifiers of the conjuncts. Negations are verified by the falsifiers, and falsified by the verifiers, of their negatum.

Definition 3 Let P, Q be unilateral propositions and $\mathbf{P}, \mathbf{Q}$ bilateral propositions on $(S, \sqsubseteq)$.

- $P \vee Q = P \cup Q$
- $P \wedge Q = \{p \sqcup q : p \in P \text{ and } q \in Q\}$

[9] An upper bound of T is any state that has every member of T as part. A least upper bound of T is any upper bound of T that is part of every upper bound of T. By anti-symmetry, least upper bounds are unique if they exist.

[10] That greatest lower bounds always exist follows from the fact that every subset of S, including the empty set, has a least upper bound. Indeed, we can always form the greatest lower bound of a subset of S by taking the fusion of all the lower bounds, i.e. all the states that are part of every member of the given subset of S.

- $\neg\mathbf{P} = (\mathbf{P}^-, \mathbf{P}^+)$
- $\mathbf{P} \wedge \mathbf{Q} = (\mathbf{P}^+ \wedge \mathbf{Q}^+, \mathbf{P}^- \vee \mathbf{Q}^-)$
- $\mathbf{P} \vee \mathbf{P} = (\mathbf{P}^+ \vee \mathbf{Q}^+, \mathbf{P}^- \wedge \mathbf{Q}^-)$

So far, we have a very liberal conception of propositions.[11] Note, for instance, that we have not excluded the possibility of one and the same state both verifying and falsifying a proposition. Intuitively, verifiers and falsifiers of a proposition should not only be distinct, but incompatible. To capture this fact, we need to incorporate a distinction between possible and impossible, or *consistent* and *inconsistent* states within our state-space.

Definition 4 A *modalized* state-space is a triple $(S, S^\diamond, \sqsubseteq)$, where

1. $(S, \sqsubseteq)$ is a state-space,
2. $S^\diamond$ is a non-empty subset of S such that $s \in S^\diamond$ whenever $s \sqsubseteq t$ for $t \in S^\diamond$

We call the members of $S^\diamond$ the *possible* or *consistent* states, and we say that some states $s_1\, s_2 \ldots$ are *compatible* iff their fusion is a possible state. We call a state a *(possible) world* iff it is a possible state that has every state it is compatible with as a part. A modalized state-space is called a *W-space* iff every possible state is part of a world. Our interest henceforth will be in W-spaces.

Definition 5 Let $(S, S^\diamond, \sqsubseteq)$ be a W-space and $\mathbf{P}$ a bilateral proposition on that W-space.

- $\mathbf{P}$ is *exclusive* iff no member of $\mathbf{P}^+ \wedge \mathbf{P}^-$ is consistent
- $\mathbf{P}$ is *exhaustive* iff every world contains some member of $\mathbf{P}^+ \vee \mathbf{P}^-$

It may be shown that the properties of exhaustivity and exclusivity are preserved under the boolean operations. Moreover, say that $\mathbf{P}$ is true (false) at a world w iff w contains some verifier (falsifier) of $\mathbf{P}$ as a part. Then for exclusive and exhaustive $\mathbf{P}$ and $\mathbf{Q}$, at any world w, $\mathbf{P}$ is true or false but not both, and the boolean operations behave classically: $\neg\mathbf{P}$ is true at w iff $\mathbf{P}$ is not, $\mathbf{P} \wedge \mathbf{Q}$ is true at w iff both $\mathbf{P}$ and $\mathbf{Q}$ are, and $\mathbf{P} \vee \mathbf{Q}$ is true at w iff either $\mathbf{P}$ or $\mathbf{Q}$ is (cf. [4, pp. 665f]).

There are a number of natural relations of entailment that may be defined between bilateral propositions. For present purposes, the most important ones are those of *inexact* and *loose* entailment, respectively.

Definition 6 Let $\mathbf{P}, \mathbf{Q}$ be bilateral propositions on some W-space.

- $\mathbf{P}$ inexactly entails $\mathbf{Q}$ iff every verifier of $\mathbf{P}$ contains a verifier of $\mathbf{Q}$
- $\mathbf{P}$ loosely entails $\mathbf{Q}$ iff every world containing a verifier of $\mathbf{P}$ contains a verifier of $\mathbf{Q}$

[11] In many applications of truthmaker semantics, requirements are imposed to the effect that the set of verifiers of a given proposition be closed under non-empty fusion—so that $\bigsqcup T \in P$ for every non-empty subset T of P—and convex—so that $t \in P$ whenever $s, u \in P$ and $s \sqsubseteq t \sqsubseteq u$. In the present context, for reasons that will become clear later, it seems preferable to me not to impose these requirements.

It is known that the propositional logic of loose entailment over exclusive and exhaustive propositions is classical, while that of inexact entailment is the logic of first-degree entailment (cf. [4, p. 669]).

4 Totality operators

The central constraint on our intended totality operator Δ is that it turn all and only weakly complete truths into strongly complete ones:

($\Delta-$S): $\Delta\mathbf{P}$ is a strongly complete truth iff $\mathbf{P}$ is a weakly complete truth.

It is helpful to contrast this constraint with a slightly different one:

($\Delta-$T): $\Delta\mathbf{P}$ is true iff $\mathbf{P}$ is a weakly complete truth.

This constraint would motivate an interpretation of $\Delta\mathbf{P}$ as simply *saying* that $\mathbf{P}$ is a weakly complete truth. Plausibly, such an interpretation would also satisfy ($\Delta-$S).[12] But as we shall see, there are natural interpretations of Δ which satisfy ($\Delta-$S) but not ($\Delta-$T).

To make the constraint precise, we need to explicate the notions of a strongly and a weakly complete truth within the truthmaker framework. To do this, it may be useful to have an intensional explication available for comparison. Suppose, then, that we adopt the familiar intensional conception of propositions as sets of possible worlds, and let @ stand for the actual world. Then a proposition is true iff @ is a member of it, and it entails every truth iff it is the singleton set {@}. So the only sensible understanding of a strongly complete truth here is as the proposition {@}. To characterize weakly complete truths, a natural approach—first suggested by [1] and studied in detail in [8]—is to appeal to a relation of *outstripping* among possible worlds. Roughly speaking, the idea is that a world outstrips another iff it contains more positive facts. A proposition P may then be regarded as a weakly complete truth iff P is a truth, and P is false at every world that is outstripped by @—if it takes all of @, as it were, to make P true. P is a weakly complete truth, therefore, iff the actual world is a minimal P-world.

Turning back to the hyperintensional framework of truthmaker semantics, we first need to define a notion of truth for propositions in a given W-space. To that end, we designate one of the worlds in a given W-space as the actual world @. For a proposition to be true is then for one of its verifiers to be part of @. As before, we want a strongly complete truth to entail every truth. But now there are several ways we can understand this condition, depending on which entailment relation we take to be at issue. If we appeal to loose entailment, we obtain an intensional criterion, since intensionally equivalent propositions agree on their loose consequences. Since we are aiming for a hyperintensional account, a more appropriate choice is *inexact entailment*. For surely a truth verified only by @ should count as strongly complete.

[12]The crucial assumption is that if $\mathbf{P}$ is a weakly complete truth, then the truth that $\mathbf{P}$ is a weakly complete truth is strongly complete. But this is prima facie plausible: if $\mathbf{P}$ falls short of completeness only by failing to entail some negative truths, then *that* truth seems to entail the negation of any positive truth not entailed by $\mathbf{P}$, and so to be strongly complete.

Any such truth will indeed inexactly entail every truth, and inexact entailment is the only natural hyperintensional entailment relation within truthmaker semantics in which a truth verified only by @ stands to every truth.

Note, though, that a truth may have more verifiers than just @ and still inexactly entail every truth, as long as every one of its additional verifiers has @ as a part. (Any such state will then be inconsistent.) Thus, suppose **P** is verified just by @, and suppose that there are no ghosts in the actual world. Then the proposition **P** ∨ (**P**∧ there are ghosts) also inexactly entails every truth. Still, since it is distinct from **P**, and since there is a non-actual way for it to be true—albeit an inconsistent one—it seems more natural to deny it the title of *the whole truth*. So we shall count a truth as strongly complete iff it is verified by @, and only by @.[13] We can give an equivalent definition in terms of inexact entailment and truth:[14]

(SCT): A proposition **P** is a strongly complete truth iff:
P is true, and
P inexactly entails every truth, and
P is inexactly entailed by every truth that inexactly entails every truth.

We now turn to the notion of a weakly complete truth. Given the way we have introduced the notion, just as we want a strongly complete truth to entail every truth, we want a weakly complete truth to entail every *positive* truth. To make this condition precise, we need not give a definition of a positive proposition or even a positive truth. Let us simply assume as given some part $@^p$ of @ as the complete positive part of @. It seems very plausible that a proposition verified only by $@^p$ should count as a positive truth, and that a proposition is a positive truth only if some part of $@^p$ is one of its verifiers. It is then easy to see that a proposition **P** inexactly entails every positive truth iff every verifier of **P** contains $@^p$ as a part, and some verifier of **P** is part of @. For similar reasons as before, we should still exclude some truths of this sort. For suppose **P** is weakly complete, and that @ contains no ghosts. Then the truth that **P** ∨ (**P**∧ there are ghosts) inexactly entails every positive truth. But it shouldn't count as weakly complete. Roughly speaking, it does not *just* fail to entail the negative truth that there are no ghosts, it explicitly allows for the possibility that there are ghosts. So I propose that we count a truth as weakly complete iff it is verified by some part of @, and all its verifiers contain $@^p$ and are part of @. There is no straightforward way to give an equivalent definition parallel to (SCT), but once we have characterized a suitable totality operator that conforms to ($\Delta-$S), we will be able to give an equivalent definition in terms of Δ and strong completeness.

[13] If our talk of possible *worlds* is taken at face value, then this explication captures thoroughly *unrestricted* notions of complete truths, with no restriction to subject matter. But if our interest is in interpreting a particular language, which merely describes some specific restricted subject matter, then we may perhaps regard the state-space, and hence the worlds in our semantics as likewise restricted to that subject matter, and in this way also handle various restricted notions of completeness. What we cannot do, in this way, is to study the interaction between notions of completeness with *different* subject matter restrictions. To do this, we should need to explicitly relativize our notions of completeness. Krämer [6] contains some discussion of how to go about this in an intensional and in a hyperintensional setting, but for the purposes of this paper we shall set these issues aside.

[14] I owe this observation to Kit Fine's response to my [6], to be published in the same volume.

So let us now turn to developing a suitable truthmaker account of Δ. To give such an account, we need to characterize the truthmakers and the falsitymakers of a proposition $\Delta\mathbf{P}$ given the truthmakers and falsitymakers of $\mathbf{P}$. The issues raised on the side of the truthmakers are largely separate from those on the side of the falsitymakers, so we shall begin by considering just the truthmakers. First of all, it seems very plausible that if the proposition $\mathbf{P}$ has just a single truthmaker s, $\Delta\mathbf{P}$ will also have just a single truthmaker, which we may think of as the state that (s obtains, and that's it). Intuitively, there is nothing disjunctive about it being the case that s obtains, and that's it, so we are justified in taking this condition to correspond to a single state. So we may take there to be a function δ on the states, which maps any given state s to the state δs to the effect that s obtains, and that's it. The task of obtaining an account of the truthmakers of $\Delta\mathbf{P}$ then divides into two subtasks: that of determining the properties of the completion operation δ on the states, and that of identifying the truthmakers of $\Delta\mathbf{P}$ when $\mathbf{P}$ is disjunctive, i.e. has more than one verifier. We address these tasks in turn, before we consider the matter of the falsifiers.

4.1 State-Completions

What can we say about the properties of δs? Part of what it is for it to be the case that s obtains, and that's it, is for it to be the case that s obtains. So it is plausible to hold that always $s \sqsubseteq \delta s$.

δ-Containment: $s \sqsubseteq \delta s$

Now since δs is the state that s obtains, *and that's it*, in some sense, δs must not contain anything in excess of s. Of course, we cannot demand that δs strictly contain nothing beyond what is contained in s, since then δs would always be identical with s. Rather, the idea has to be that δs contains no *positive* state that is not already contained in s. Indeed, any positive state not contained in s would constitute a counter-example to the claim that s obtains, and *that's it*, so δs must be incompatible with any such state.

One strategy for capturing these ideas formally is to assume as given a division of the states into positive and negative, and simply lay down as a further requirement on δ that every positive state that is part of δs also be part of s, and that every other positive state is incompatible with δs. Other things being equal, however, it would seem preferable not to rely on a multitude of primitive notions—the δ function and a positive/negative divide—within our semantics for Δ, and to stick to a single primitive if possible. So for now, I propose to instead appeal to a distinction between positive and negative states only in an informal, heuristic capacity to motivate constraints on δ that can be stated without appeal to that distinction. Below, we will see that there is a natural way to define the notion of a positive state in terms of δ.

We begin by considering some modal properties of δ. First of all, we can argue that

δ-Completeness: δs is a world if consistent.

For recall that $\Delta\mathbf{P}$ is to be strongly complete if $\mathbf{P}$ is weakly complete. Moreover, it is natural to think that this should hold of necessity, so if $\mathbf{P}$ is a weakly complete truth with respect to any world w, then $\Delta\mathbf{P}$ is a strongly complete truth with respect to w.

Now suppose **P** has a single verifier s, which is consistent and which contains the entire positive part of some given world w. Then **P** inexactly entails every positive truth at w, and so $\Delta\mathbf{P}$ must inexactly entail every truth at w. So in this case, δs must be identical to w. Note that we are assuming, as seems plausible, that for any positive state, there is at most one world whose positive part is that state.

Suppose now that s is consistent but does not contain the positive part of any world. In that case, δs must be regarded as inconsistent. For suppose it is consistent. If it is a world, then δs contains a positive state not contained in s, contrary to our informal desiderata. If it is a proper part of a world, it is compatible with the part of that world, which is by assumption not contained in s, which is likewise incompatible with our desiderata. Finally, if s is inconsistent, δs must also be inconsistent by δ-Containment. It follows that δs is a world if consistent.

From the principle that no two worlds share their positive part, we may also draw some conclusions about the conditions under which δs and δt are identical. Firstly, if δs is consistent, and $s \sqsubseteq t \sqsubseteq \delta s$, then $\delta t = \delta s$. For if δs is consistent, and hence a world, s contains the positive part of δs. Since $s \sqsubseteq t$, so does t. So t is consistent and contains the entire positive part of δs, and so $\delta t = \delta s$. Secondly, suppose $\delta s = \delta t$ is a world. Then both s and t contain the positive part of that world, hence so does $s \sqcap t$. So $s \sqcap t$ is part of a world and contains the entire positive part of that world, so $\delta(s \sqcap t)$ is that world. And since there is at most one world with a given positive part, it is the same world as δs and δt.

For both principles, it seems very plausible to take them to apply in the same way if δs is not consistent. For consistent or not, δs goes beyond s in a purely negative way, merely excluding the obtaining of any positive state not already in s. If t is between s and δs, then t, too, goes beyond s at most in excluding the obtaining of some positive states not already in s, and δt goes beyond t only in excluding the obtaining of any positive state not even in t. But the positive states not contained in t are exactly those not contained in s, and so there would seem to be no basis for any distinction between δs and δt.[15]

Moreover, if $\delta s = \delta t$ is inconsistent, we may still infer that s and t have the same positive part, since otherwise δs and δt would differ with respect to what positive states they exclude. In general, we may think of δs as dividing into (i) the positive part of s, (ii) the part excluding any positive state not contained in s, as well as (iii) any other non-positive components of s.[16] The positive part of s and t must be the

[15] In [6], I suggested that $\delta s = \delta t$ if s and t have the same positive part, but this is plausible only under the assumption that s and t are consistent.

[16] This assumes a picture whereby the completion of a state is formed by fusing it with certain negative states that would normally be disjoint from that state. An alternative picture is that completions of a state always sit immediately on top, as it were, of the completed state in the part-whole ordering, so that every proper part of δs is a part of s. While this would violate the mereological principle of weak supplementation, the general constraints on state-spaces do not exclude this possibility, and it does not seem out of the question that the mereology of states should be non-standard in this kind of way. If one adopted this alternative picture, the identity principles would of course still hold, since we would then have that $\delta s = \delta t$ implies $s = t$.—It is worth noting that if we wish to allow for the existence of subject-matter-restricted completions, that is some reason not to adopt this alternative picture. For consider a state s, and some restricted subject-matter m. Then the state that (s obtains, and that's it with respect to m) is plausibly regarded as a part of δs that is typically disjoint from s.

same, given that $\delta s = \delta t$, and so it must also be the same as the positive part of $s \sqcap t$. Consequently, the part of δs excluding any positive state not contained in s must also be the same as the corresponding part of δt and $\delta(s \sqcap t)$. Now consider any other non-positive components of s. Since they are also to be found in δt, they must be among the other non-positive components of t, and hence also parts of $s \sqcap t$. We may thus infer that $\delta s = \delta(s \sqcap t)$.

δ-Identity(1): $\delta s = \delta t$ if $s \sqsubseteq t \sqsubseteq \delta s$
δ-Identity(2): $\delta s = \delta(s \sqcap t)$ if $\delta s = \delta t$

Finally, it is plausible that in certain cases, the application of δ will be redundant. In particular, if s is already a world, δs must be taken to be the same as s. By δ-Containment, the only alternative would be for δs to be inconsistent. Yet if s is a world, then it is surely not inconsistent that s obtains, and that's it. More generally, it is very plausible that $\delta\delta s = \delta s$. For $\delta\delta s$ goes beyond δs at most in saying that no positive state not contained in δs obtains. But since δs already says that no positive state not contained in s obtains, this condition is already imposed by δs.

δ-Redundancy(1): $\delta s = s$ if s is a world
δ-Redundancy(2): $\delta\delta s = \delta s$

There are perhaps other constraints that might plausibly be imposed on δ. As we shall see, though, these constraints suffice to render sound a natural propositional logic for Δ. Inspection of the completeness proof will show, moreover, that a number of further plausible constraints on δ have no effect on the logic obtained (although this may change, of course, when we consider languages with more expressive resources such as other modal operators).

Definition 7 A *completed* state-space (short: C-space) is a quadruple $(S, S^{\diamond}, \sqsubseteq, \delta)$, where

1. $(S, S^{\diamond}, \sqsubseteq)$ is a W-space,
2. $\delta : S \to S$ satisfies the constraints of δ-Containment, δ-Completeness, δ-Identity, and δ-Redundancy.

We are now in a position to argue that $(\Delta{-}S)$ holds for the case of propositions with just a single verifier, using two extremely plausible assumptions about the positive part $@^p$ of $@$. The first is that the completion of $@^p$ is $@$, i.e. $\delta(@^p) = @$. The second is that $@^p$ is the smallest state whose completion is $@$, i.e. $\delta s = @$ implies $@^p \sqsubseteq s$.[17] Let **P** have just a single verifier s. As we have argued, Δ**P** then has just the single verifier δs. Now suppose **P** is to be a weakly complete truth. Then its sole verifier s has to contain the positive part $@^p$ of $@$, and also to be part of $@$. But $\delta(@^p) = @$, and so by δ-Identity(1), $\delta s = @$. So Δ**P** is a strongly complete truth. Conversely, suppose Δ**P** is a strongly complete truth, hence that $\delta s = @ = \delta(@^p)$. By δ-Containment, $s \sqsubseteq @$. By the second assumption, $@^p \sqsubseteq s$. So **P** is a weakly complete truth.

[17] In Section 4.5, I propose a way to define the notion of the positive part of a world in terms of δ under which these assumptions can be derived.

It is worth noting that when $\mathbf{P}$ has just a single verifier s, even (Δ−T) holds. For then $\Delta\mathbf{P}$ has just the single verifier δs, and so if $\Delta\mathbf{P}$ is to be true at all, δs must be @, so $\Delta\mathbf{P}$ is again a strongly complete truth.

4.2 Examples of C-Spaces

It is clear that there are C-spaces. Indeed, there are very small ones. Take any object x, and let $S = \{x\}$, $S^{\Diamond} = S$, $\sqsubseteq = S \times S$, and let $\delta x = x$. Then $(S, S^{\Diamond}, \sqsubseteq, \delta)$ is readily seen to be a C-space. Any unmodalized state-space may be extended into a C-space in a trivial way by letting $S^{\Diamond} = S$ and $\delta s = \blacksquare$ for all $s \in S$. Still, such C-spaces are not very natural. Since we have $\blacksquare = \delta\square$ and δs is supposed to always extend s in a purely negative way, in these C-spaces we are, in effect, regarding all states (except, perhaps, the nullstate) as negative.

But there are also more interesting and natural C-spaces. We specify two methods for constructing such spaces. Let $(S_0, \sqsubseteq_0)$ be any unmodalized state-space. We may construct a C-space by regarding every member of S_0 as wholly positive, and adding for each state in S_0 a distinct δ-state immediately 'on top'. First, with each pair of states s, t with $s \sqsubseteq_0 t$, we associate a distinct item $f(s, t)$ that is not a member of S_0. Intuitively, we may think of $f(s, t)$ as saying that t obtains, and that no positive state not contained in s obtains. We pick an additional item $\blacksquare$ to play the role of the fullstate in the C-space to be constructed. Let N—the new states—be the set comprising $\blacksquare$ and the values of f, and let $S = S_0 \cup N$. We define the partial order on S indirectly, by first specifying the fusion operation $\bigsqcup$ and then letting $s \sqsubseteq t$ iff $\bigsqcup\{s, t\} = t$. For $T \subseteq S$, we distinguish three cases. First, if $T \subseteq S_0$ then $\bigsqcup T = \bigsqcup_0 T$. Second, if T has at least two members that belong to N, $\bigsqcup T = \blacksquare$. Third, if $f(s, t)$ is the sole member of N in T, then $\bigsqcup T = f(s, \bigsqcup(T \cap S_0))$. It is routine to show that $\sqsubseteq$ as defined from $\bigsqcup$ is a complete partial order on S, with $\bigsqcup T$ the least upper bound of $T \subseteq S$. Now let δs be $f(s, s)$ if $s \in S_0$, and s otherwise, and let $S^{\Diamond} = S_0 \cup \{\delta s : s \in S_0\}$.

Definition 8 Let $(S_0, \sqsubseteq_0)$ be any state-space. Let f be a one-one function from $\sqsubseteq_0$ onto a set D disjoint from S_0, and let $\blacksquare$ be an item not in S_0 or D. The *simple C-space based on* $(S, \sqsubseteq_0)$ with f and $\blacksquare$ is $(S, S^{\Diamond}, \sqsubseteq, \delta)$, with

- $S = S_0 \cup D \cup \{\blacksquare\}$
- $\bigsqcup : \mathrm{wp}(S) \to S$ with

$$\bigsqcup T \mapsto \begin{cases} \bigsqcup_0 T & \text{if } T \subseteq S_0 \\ f(s, t \sqcup_0 \bigsqcup_0 (T \cap S_0)) & \text{if } T \cap N = \{f(s, t)\} \\ \blacksquare & \text{otherwise} \end{cases}$$

- $s \sqsubseteq t$ iff $\bigsqcup\{s, t\} = t$
- $\delta s = f(s, s)$ if $s \in S_0$ and s otherwise
- $S^{\Diamond} = S_0 \cup \{\delta s : s \in S_0\}$

Proposition 1 *Let* $(S_0, \sqsubseteq_0)$*,* f *and* $\blacksquare$ *be as in the previous definition. Then the simple C-space based on* $(S, \sqsubseteq_0)$ *with* f *and* $\blacksquare$ *is a C-space.*

Proof Omitted. □

Instead of regarding all the elements of S_0 and their δ-completions as consistent, we can also pick any non-empty subset D of $\{\delta s : s \in S_0\}$ and let $S^\Diamond$ be the closure under part of D. The result will then still be a C-space.

There is also a natural way to construct a C-space from a given set of atomic sentences (short: atoms) and a division of the atoms into positive and negative ones. Let At be some fixed set of sentence letters, and let $At^P \subseteq At$. The set of literals Lit is $At \cup \{\neg p : p \in At\}$. The *converse* φ' of a literal φ is $\neg\varphi$ if φ is an atom, and ψ if $\varphi = \neg\psi$. For $\varphi \in Lit$, a set $s \subseteq Lit$ is said to be φ-neutral iff neither φ nor φ' are members of s.

Let $At^N = At \backslash At^P$. Let P – the set of positive literals – be $At \cup \{\neg\varphi : \varphi \in At^N\}$, and let N – the set of negative literals – be $Lit \setminus P$. Note that $N = At^N \cup \{\neg\varphi : \varphi \in At^P\}$. For φ a literal, we let $p(\varphi)$ and $n(\varphi)$ be the positive and negative members of $\{\varphi, \varphi'\}$, respectively.

Definition 9 The *canonical C-space based on* $At,\ At^P$ is $(S, S^\Diamond, \sqsubseteq, \delta)$, with

- $S = \text{wp}(Lit)$
- $S^\Diamond = \{s \in S : \{p, \neg p\} \not\subseteq s$ for all $p \in At\}$
- $\sqsubseteq = \subseteq \restriction S$
- $\delta s = s \cup \{n(\varphi) : \varphi \in Lit$ and s is φ-neutral$\}$

Proposition 2 *Let At be some set of sentence letters, and let $P \subseteq At$. The canonical C-space based on At, P is a C-space.*

Proof It is easily established that $(S, S^\Diamond, \sqsubseteq)$ is a W-space (cf. [4 p. 647]). It remains to prove that δ satisfies the additional conditions imposed on a C-space.

δ-Containment: Immediate from the definition of $\sqsubseteq$ and δ.

δ-Completeness: Suppose δs is consistent and suppose $t \in S$ is compatible with δs. Let φ be any literal in t. We show that $\varphi \in \delta s$. If $\varphi \in s$ the claim follows immediately, so suppose $\varphi \notin s$. Since t is compatible with δs and hence s, also $\varphi' \notin s$, so s is φ-neutral. By definition of δ, $n(\varphi) \in \delta s$. Since t is compatible with δs, we may infer that $\varphi = n(\varphi)$ and hence that $\varphi \in \delta s$.

δ-Identity(1): Suppose $s \sqsubseteq t \sqsubseteq \delta s$. We show first that $\delta t \sqsubseteq \delta s$. Since $t \sqsubseteq \delta s$ and $\delta t = t \cup \{n(\varphi) : \varphi \in Lit$ and t is φ-neutral$\}$, it suffices to show that $\{n() : \varphi \in Lit$ and t is φ-neutral$\} \subseteq \delta s$. But since $s \sqsubseteq t$, s is φ-neutral whenever t is, so $\{n(\varphi) : \varphi \in Lit$ and t is φ-neutral$\} \subseteq s \cup \{n(\varphi) : \varphi \in Lit$ and s is φ-neutral$\} \subseteq \delta s$.

We now show that also $\delta s \sqsubseteq \delta t$. Since $s \sqsubseteq t \sqsubseteq \delta t$, it suffices to show that $\{n(\varphi) : \varphi \in Lit$ and s is φ-neutral$\} \subseteq \delta t$. Consider any $\varphi \in Lit$ with s φ-neutral. If t is also -neutral, then $n(\varphi) \in \delta t$ by definition of δ. If t is not φ-neutral, then either $n(\varphi)$ or $(n(\varphi))'$ is a member of t. Suppose that $(n(\varphi))' \in t$. Then since $t \sqsubseteq \delta s$, $(n(\varphi))' \in \delta s$, contrary to the assumption that s is φ-neutral. So $n(\varphi) \in t$ and hence $n(\varphi) \in \delta t$, as required.

δ-Identity(2): Suppose $\delta s = \delta t$. We need to show that $\delta s = \delta(s \sqcap t)$. Suppose first that $\varphi \in \delta s$. Then (a) $\in s$ or (b) $\varphi = n(\varphi)$ and s is φ-neutral. Since $\delta s = \delta t$, also $\in \delta t$, and so either (c) $\varphi \in t$ or (d) $\varphi = n(\varphi)$ and t is φ-neutral. If either (b) or (d) hold, then $\varphi = n(\varphi)$ and $s \sqcap t$ is φ-neutral, and so $\varphi \in \delta(s \sqcap t)$. But if both (a) and (c) hold, then $\varphi \in s \sqcap t$ and hence also $\varphi \in \delta(s \sqcap t)$.

Now suppose $\varphi \in \delta(s \sqcap t)$. Then either (a) $\varphi \in s \sqcap t$, or (b) $\varphi = n(\varphi)$ and $s \sqcap t$ is φ-neutral. Clearly if (a) then $\varphi \in s$ and hence $\varphi \in \delta s$. So assume (b). If s is φ-neutral, then clearly $\varphi = n(\varphi) \in \delta s$. If s is not φ-neutral, it follows that either φ or φ' is a member of s. Suppose for reductio that $\varphi' \in s$. Then since $s \sqcap t$ is by assumption φ-neutral, φ' is not a member of $s \sqcap t$, and hence φ' is not a member of t. But since $\delta s = \delta t$ and so $\varphi' \in \delta t$, it then follows that $\varphi' = n(\varphi)$, contrary to our assumption that $\varphi = n(\varphi)$.

δ-Redundancy(1): From the fact that a world is φ-neutral for no φ.

δ-Redundancy(2): From the fact that δs is φ-neutral for no φ. □

4.3 Disjunction

We have given an account of the verifiers of $\Delta\mathbf{P}$ when $\mathbf{P}$ has just one verifier, in the form of various principles about the state-completion function δ. We now need to extend the account to the case in which $\mathbf{P}$ is disjunctive in the sense that it has several verifiers. What should we take the verifiers of $\Delta\mathbf{P}$ to be in this case? Unfortunately, it is much less clear what we should say here. Most paradigm uses of 'that's it' and its ilk seem to be applications to non-disjunctive statements. A typical disjunctive statement explicitly leaves open, to an extent, how matters stand: are they as described in the one disjunct, or are they as described in the other disjunct? As a result, when a disjunctive statement has been made, there is normally something more to be said, namely which disjunct is true. So there is typically a degree of oddity to following up an explicitly disjunctive statement by stating: *and that's it*. At any rate, the application of Δ to disjunctive arguments raises distinctive issues of interpretation that do not come up when Δ is applied to propositions with a single verifier.

I want to consider two natural approaches to the problem. One is to take the verifiers of $\Delta\mathbf{P}$ to be exactly the states of the form δs when s verifies $\mathbf{P}$.[18] On this view, Δ distributes over disjunction in the sense that $\Delta(\mathbf{P} \vee \mathbf{Q})$ and $\Delta\mathbf{P} \vee \Delta\mathbf{Q}$ have exactly the same verifiers. I will thefore call it the *disjunctive* conception of Δ. $\Delta\mathbf{P}$ will then be true just in case *some* verifier of $\mathbf{P}$ has the actual world as its completion.

This view conforms to (Δ−S). To see this, suppose $\Delta\mathbf{P}$ is strongly complete, so it is verified by @, and only by @. Let s be a verifier of $\mathbf{P}$. Then $\delta s = @ = \delta(@^p)$, and thus by the same reasoning as above, $@^p \sqsubseteq s \sqsubseteq @$. Conversely, suppose $\mathbf{P}$ is weakly complete, so all its verifiers are between $@^p$ and @. As before, for any such

[18] This is also the view I proposed in [6].

state s, $\delta s = @$, so $\Delta\mathbf{P}$ is verified only by @. The view does not conform to ($\Delta-$T), on which $\Delta\mathbf{P}$ might be read as saying that **P** is a weakly complete truth. For if **P** is a weakly complete truth with a single verifier, then $\Delta\mathbf{P}$ is a strongly complete truth and therefore true. But then $\Delta\mathbf{P} \vee \Delta\mathbf{Q}$ is also true, and by the distributivity principle, so is $\Delta(\mathbf{P} \vee \mathbf{Q})$. ($\Delta$-T) would allow us to infer that $\mathbf{P} \vee \mathbf{Q}$ is weakly complete. But **Q** was arbitrary, and so might be verified, for example, by some small proper part of $@^p$, rendering $\mathbf{P} \vee \mathbf{Q}$ clearly not weakly complete.

On another approach, we take $\Delta\mathbf{P}$ to be verified by the *fusion* of all states of the form δs when s verifies **P**. On this view, Δ distributes over disjunction in the sense that $\Delta(\mathbf{P} \vee \mathbf{Q})$ and $\Delta\mathbf{P} \wedge \Delta\mathbf{Q}$ have exactly the same verifiers. I will thefore call it the *conjunctive* conception of Δ. $\Delta\mathbf{P}$ will then be true just in case *every* verifier of **P** has the actual world as its completion. This view conforms to ($\Delta-$T) as well as to ($\Delta-$S), and thus fits with a reading of $\Delta\mathbf{P}$ as saying that **P** is a weakly complete truth. For note that $\Delta\mathbf{P}$ will always have exactly one verifier. That verifier will be @ iff $@ = \delta s$ for every verifier s of **P** and thus if and only if **P** is weakly complete. Otherwise it will be the fusion of several δ-states, and hence inconsistent, so $\Delta\mathbf{P}$ will be false.

Both views, I believe, correspond to useful and legitimate conceptions of a totality operator, suited for slightly different purposes. They agree on the paradigm applications of totality operators to non-disjunctive arguments, and differ drastically on the applications to disjunctive arguments. For disjunctive **P**, the conjunctive conception makes it almost impossible for $\Delta\mathbf{P}$ to be true, while the disjunctive conception makes it fairly easy for $\Delta\mathbf{P}$ to be true.

The conjunctive conception, as highlighted, fits a reading of Δ as expressing the notion of weak completeness.[19] Especially in the kind of metaphysical contexts described in the introduction, that notion is of central interest, and so it is useful to have an operator expressing it, and worthwhile studying its logic. The disjunctive conception, on the other hand, seems to fit better than the conjunctive one with many ordinary language applications of 'that's it' to disjunctive arguments. Under this conception, $\Delta(\mathbf{P} \vee \mathbf{Q})$ is true exactly when $\Delta\mathbf{P}$ or $\Delta\mathbf{Q}$ is true. Taken as an account of 'that's it', the view thus predicts that utterances of the form 'A or B, and that's it' will be felicitous just when 'A, and that's it, or B, and that's it' will sound felicitous (modulo the awkward repetitiveness of the second formulation). And that prediction seems to be borne out in many cases. Suppose I describe my breakfasting habits by saying: I always have porridge and tea or porridge and coffee, and that's it. Under the most natural reading of my statement, it is equivalent to, and sounds just as good

[19] I would like here to thank an anonymous referee whose criticial comments on a previous version of the paper prompted me to formulate and develop the conjunctive conception in addition to the disjunctive one. The attractiveness of a conception of Δ in line with ($\Delta-$T) is also emphasized by Kit Fine in his response to my [6]. Fine's proposal for such a notion is slightly different than mine; a comparison will have to wait for another occasion.

as, the claim that I always have porridge and tea, and that's it, or porridge and coffee, and that's it.[20]

Of course, such disjunctive answers do not always sound good. Suppose I am asked what I had for breakfast, and I answer that I had cereal or toast, and that's it. This answer seems less than ideal, and invites the question which it was. But that is just what we would expect also for the disjunction 'I had cereal, and that's it, or I had toast, and that's it'. It is normally safe to assume that the speaker will know which disjunct obtains, and so their answer is less informative than it could be and seems called for.

Relatedly, when the hypothesis that the speaker doesn't know which disjunct obtains seems more reasonable, disjunctive 'that's it' statements sound fine. Suppose Bob asks how much Bill had to drink last night. Bill responds: 'I just had two or three beers, and that's it!' Under the most natural interpretation, what Bill said is true iff he either had two beers, and that's it, or he had three beers, and that's it, and the disjunctive answer is fine assuming Bill isn't quite sure whether it was two or three. So there seems to be a systematic pattern in our uses of 'that's it' in application to disjunctive arguments which seems to conform very well to the disjunctive interpretation proposed for Δ. It may be worth making explicit that the usage pattern here does *not* fit well with any interpretation of Δ that sustains (Δ−T). On such an interpretation, Bill's 'I had two or three beers, and that's it' is equivalent to 'That I had two or three beers is a weakly complete truth'. But that statement is incompatible with Bill having had three beers, since that would be a positive truth not entailed by the truth that he had two or three beers.

Admittedly, there are also some 'that's it' statements that are counted as true under the proposed account, which do sound quite bad. For a proposition may have a state between $@^p$ and @ among its verifiers without being particularly informative. The reason is that while exact truthmaking is subject to a relevance constraint, it is not subject to a minimality constraint, and even very large states can be wholly relevant to very weak propositions. For example, it is plausible that $@^p$ is among the verifiers of the proposition that something exists, or that something is the case. As a result, the claim that something exists, and that's it, comes out true, bizarre though it sounds.

[20]In some contexts, it may be more natural to read 'A or B, and that's it' as also allowing for the possibility that A *and* B, and that's it. One way to accommodate this datum would be to adopt the so-called *inclusive* clause for disjunction, on which where $A \vee B$ is verified by the verifiers of A, the verifiers of B, *and any fusion of such verifiers*. However, there are also contexts in which it is more natural to read 'A or B, and that's it' as *not* allowing for the possibility that A *and* B. For instance, I might describe the very minimal breakfast enjoyed by another hotel guest by saying: he just had some coffee or tea, and that's it. (Assume I didn't see whether it was tea or coffee.) The most natural interpretation excludes the possibility that the person had both tea and coffee. So there are good reasons, in the present context, not to rely on the inclusive conception of disjunction across the board. (Note that some other natural language phenomena, to do especially with counterfactuals, seem also to speak against the inclusive conception; cf. [2] and [5].) A better view, then, seems to be that 'A or B' is sometimes, but not always, used in accordance with the inclusive conception—equivalently, as shorthand for 'A or B or both'—and to regiment such utterances accordinly as '$\Delta(A \vee B \vee (A \wedge B))$'. On this way of accommodating the datum, we can stick to our original, non-inclusive clause for disjunction as well as the proposed clause for Δ. Thanks here to an anonymous referee for discussion and the references.

In defence of the disjunctive account, we may note that it actually *predicts* that an utterance of this claim, though true, should seem bizarre. For on this account, $\Delta\mathbf{P}$ is true iff some verifier of **P** both obtains and contains every positive fact (pertaining to the subject matter). So what one does by applying Δ to an argument **P**—or by appending 'and that's it'—is in effect to *exclude* any positive state that does *not* help verify **P**. But when **P** is a proposition that is partially verified by *any* positive state, no matter how big, then the application of Δ can't serve its purpose, as it excludes nothing. So we should expect an application of 'that's it' to such statements to sound bizarre.[21]

4.4 Falsifiers

We now turn to the question of the exact falsifiers of totality statements. An exact falsifier of a given statement needs to be 'big enough' to bring about that the statement in question is false, and 'small enough' so as not to include any parts irrelevant to the falsity of that statement. In addition, our account of the falsifiers needs to render totality statements exclusive and exhaustive. That is, we need it to be the case that every verifier is incompatible with every falsifier, and that every possible world contains either a verifier or a falsifier.

First of all, since **P** seems to relate to $\Delta\mathbf{P}$ much like a conjunct relates to a conjunction, it is very plausible to take any state falsifying **P** to also falsify $\Delta\mathbf{P}$. But of course falsifying **P** cannot be the only way to falsify $\Delta\mathbf{P}$, another way is to merely exclude what $\Delta\mathbf{P}$ says beyond **P**. Let us again begin by considering the simple case in which **P** has only a single verifier s, so $\Delta\mathbf{P}$ likewise has just a single verifier δs. A natural suggestion is to then take any δ-state distinct from δs to falsify $\Delta\mathbf{P}$. Any such state is of course incompatible with δs, so the condition of exclusivity is met, and it seems wholly relevant to making $\Delta\mathbf{P}$ false. Roughly speaking, δs tells us what the world is like in every single respect, and any distinct δ-state answers the same question, but in a different and therefore incompatible way. Moreover, including all these δ-states as falsifiers also ensures exhaustivity, since every world is either identical to δs or to some other δ-state.

One might also plausibly regard some smaller states as falsifiers of $\Delta\mathbf{P}$. For example, by invoking the notion of a positive state that I define below, one might take any positive state not contained in δs also to falsify $\Delta\mathbf{P}$. It turns out, however, that such an extension of the set of falsifiers does not change the propositional logic of our totality operators, given the very plausible assumption that any falsifier of $\Delta\mathbf{P}$ is either a falsifier of **P** or part of some δ-state distinct from δs. For present purposes, it is therefore more convenient to stick to the simpler proposal.

It remains to extend our account to the general case with multiple verifiers for $\Delta\mathbf{P}$. Recall that in general, the falsifiers of a conjunction are obtained by taking the disjunction of the falsifier-sets of the conjuncts, and the falsifiers of a disjunction are

[21] An anonymous referee pointed out that 'that's it' is infelicitous in application to 'at least' statements: it is at best rather odd say that say Ben had an least an apple, and that's it. That datum seems to allow for a similar explanation.

obtained by taking the conjunction of the falsifier-sets of the disjuncts. So a natural strategy is to apply the same idea to the conjunctive and disjunctive interpretations of Δ as well. Call a state t an *incompletion* of a state s iff t is a δ-state distinct from δs. Let $I(s)$ be the set of s's incompletions, and let us also write $C(s)$ for the singleton set of s's completion δs. Then we define the disjunctive and conjunctive Δ-operations as follows:

Definition 10 Let $\mathbf{P}$ be a bilateral proposition on a C-space $(S, S^{\Diamond}, \sqsubseteq, \delta)$. Then

$$\Delta_C\mathbf{P} = (\bigwedge\{C(t) : t \in \mathbf{P}^+\}, \mathbf{P}^- \vee \bigvee\{I(t) : t \in \mathbf{P}^+\}$$

$$\Delta_D\mathbf{P} = (\bigvee\{C(t) : t \in \mathbf{P}^+\}, \mathbf{P}^- \vee \bigwedge\{I(t) : t \in \mathbf{P}^+\}$$

Lemma 3 *Let* $\mathbf{P}$ *be an exclusive and exhaustive bilateral proposition on a C-space* $(S, S^{\Diamond}, \sqsubseteq, \delta)$. *Then* $\Delta_C\mathbf{P}$ *and* $\Delta_D\mathbf{P}$ *are also exclusive and exhaustive.*

Proof *Exclusivity of* $\Delta_C\mathbf{P}$: Let s verify $\Delta_C\mathbf{P}$ and let t falsify $\Delta_C\mathbf{P}$. If s is inconsistent, then s is trivially incompatible with t. If s is consistent, $s = \delta u$ for every u verifying $\mathbf{P}$. If t falsifies $\mathbf{P}$, by exclusivity of $\mathbf{P}$, t is incompatible with u, and hence with $\delta u = s$. Otherwise, t is an incompletion of u, and hence a δ-state distinct from s, and therefore again incompatible with s.

Exhaustivity of $\Delta_C\mathbf{P}$: Let w be a world. Suppose w does not contain a falsifier of $\Delta_C\mathbf{P}$. So w is not an incompletion of any verifier of $\mathbf{P}$. Since w is a δ-state, it follows that $w = \delta s$ for every verifier s of $\mathbf{P}$, and hence that w verifies $\Delta_C\mathbf{P}$.

Exclusivity of $\Delta_D\mathbf{P}$: Let s verify $\Delta_D\mathbf{P}$ and let t falsify $\Delta_D\mathbf{P}$. Then $s = \delta u$ for some u verifying $\mathbf{P}$. If t falsifies $\mathbf{P}$, by exclusivity of $\mathbf{P}$, t is incompatible with u and hence with s. If t does not falsify $\mathbf{P}$, t is either inconsistent or an incompletion of u, so again, t is incompatible with s.

Exhaustivity of $\Delta_D\mathbf{P}$: Let w be a world and suppose w does not contain a verifier of $\Delta_D\mathbf{P}$. Then w is an incompletion of every verifier of $\mathbf{P}$, and hence a falsifier of $\Delta_D\mathbf{P}$. □

Exclusivity and exhaustivity ensure that negation has the classical modal profile when applied to Δ. The central remaining question for our logics of totality is how negated Δ-statements behave within the hyperintensional context of another occurrence of Δ, i.e. under what conditions $\Delta\neg\Delta A$ might be true. Under the disjunctive interpretation, this will be so exactly when $\neg\Delta A$ is true. For if $\neg\Delta A$ is true, then it is verified by—perhaps among other things—the actual world, and then so is $\Delta\neg\Delta A$. Note that this would still be the case if we were to add further falsifiers to ΔA. Under the conjunctive interpretation on the other hand, it will almost be impossible for $\Delta\neg\Delta A$ to be true. Specifically, $\Delta\neg\Delta A$ is false whenever there is at least one world w in our state-space which is non-empty in the sense that $w \neq \delta\square$. For in that case, there are at least three δ-states—$\delta\square$, δw, and $\delta\blacksquare$—of which at most one will verify ΔA. So at least two δ-states will verify $\neg\Delta A$, and hence $\Delta\neg\Delta A$ will be verified by a fusion of distinct δ-states, which is bound to be inconsistent. Again, the point is stable under the addition of more falsifiers.

In models with just two δ-states, and therefore just one, empty world @, $\Delta\neg\Delta A$ holds iff $\Delta\neg A$ does. For suppose $\Delta\neg\Delta A$ holds. Then $\neg\Delta A$ is verified only by parts of @, and since every verifier of $\neg A$ verifies $\neg\Delta A$, it follows that $\neg A$ is verified only by parts of @. Since @ is empty, any such state has @ as its completion, so $\Delta\neg A$ holds. Conversely, if $\Delta\neg A$ holds, $\neg A$ is verified only by parts of @, so A is verified only by states not contained in @, which will be inconsistent. So ΔA will be verified by the δ-state distinct from @, which must be ■. Then any falsifier of A, and any part of any incompletion of a verifier of A, and therefore any verifier of $\neg\Delta A$, is a part of @. So $\Delta\neg\Delta A$ is verified solely by @.

4.5 Positivity

Our informal reasoning about the behaviour of the completion operation δ has been strongly informed by the idea that δs goes beyond s in excluding the obtaining of any *positive* state not contained in s. Let us consider the relation between the notions in a bit more detail. It turns out that there is a plausible way to define a notion of a state's being positive in terms of δ and the mereology of the state-space.

First of all, it is plausible that δ-Identity(2) may be strengthened in the following way. Say that states s and t are δ-equivalent iff $\delta s = \delta t$. Now let T be a set of δ-equivalent states. Then δ-Identity(2) implies that if T is *finite*, then $\delta \bigsqcap T = \delta s$ for all $s \in T$. But surely the principle is just as plausible in the case of an infinite set T. So suppose it holds in general. Then for any state s, the state $\bigsqcap\{t \in S : \delta s = \delta t\}$ will be the smallest state δ-equivalent to s. We may call that state the δ-core of s and denote it as s^δ; note that always $s^\delta \sqsubseteq s$.

Now in terms of the notion of a δ-core, we can plausibly define the notion of the positive part w^p of a world w. For it seems clear that this will always simply the δ-core of the world in question, so that $w^p = w^\delta$. From this definition, we can then straightforwardly *derive* the two assumptions about $@^p$ we used in order to establish (Δ−S): that $\delta(@^p) = @$ and that $@^p \sqsubseteq s$ whenever $\delta s = @$.

Defining the monadic notion of a wholly positive state is a little more tricky. Call a state s δ-minimal iff it is its own δ-core, i.e. $s = s^\delta$. Roughly, a δ-minimal state contains no parts that are rendundant in the sense that if one were to remove them from s, they would be added back in when we form the completion of the result. At first glance, it is tempting to suppose that all and only the wholly positive states will be δ-minimal. It is clear, firstly, that a wholly positive state must be δ-minimal. This is because if a state s is not δ-minimal, it must have a proper part s' with $\delta s' = \delta s$, and so some part of s must be added to s' in forming its completion $\delta s'$. But since nothing positive may be added to a state in forming its completion, it follows that s is not wholly positive.

But the converse is not plausible. For given some wholly positive state s, we may plausibly extend it by a negative state that excludes some parts of s. Call the resulting state s^+. Then s^+ should be δ-minimal. For suppose t is a proper part of s^+. If t lacks some positive part of s^+, then clearly t and s^+ are δ-inequivalent. And if t only lacks some negative part of s^+, then that part will exclude part of t, and therefore not be added back in upon forming the completion δt. So again t and s^+ must be δ-inequivalent.

So δ-minimality is not sufficient for positivity. However, it is plausible that any state that is not wholly positive, although possibly δ-minimal, will have a *part* that is not δ-minimal. For instance, in the case of s^+, the result of removing from s^+ any positive part that is excluded by s^+ will not be δ-minimal. A wholly positive state, on the other hand, will not only itself be δ-minimal, but will also have exclusively δ-minimal parts. So we may plausibly maintain that:

(Df. δ-Positivity) A state s is wholly positive iff $t^\delta = t$ for all $t \sqsubseteq s$

It follows from this definition that any part of a wholly positive state is itself wholly positive, as one would expect.[22]

5 The Logics of Totality

In our choice of language, we largely follow the lead of [8] and extend a standard propositional language by a one-place sentential operator Δ[23] to express the target concept of totality. More precisely, we take as given a set At of (non-logical) sentence letters or *atoms*. The formulas of our *language of totality* $\mathcal{L}_\Delta$ are: the members of At, a logical constant for triviality $\top$, as well as $\neg A$, $(A \wedge B)$, $(A \vee B)$, and ΔA whenever A and B are formulas. We abbreviate $\neg\top$ by $\bot$.

The point of including a special purpose triviality constant $\top$ among the formulas is to enable us to express within $\mathcal{L}_\Delta$ that the world is empty. For $\top$ will be interpreted as verified only by the nullstate (and falsified only by the fullstate), so $\Delta\top$ will be verified only by $\delta\square$, the one wholly negative world (if consistent).[24]

Definition 11 A model is a tuple $\mathcal{M} = (S, S^\Diamond, \sqsubseteq, \delta, @, [\cdot])$ where $(S, S^\Diamond, \sqsubseteq, \delta)$ is a C-space, @ is a world, and $[\cdot]$ maps each member of At to an exclusive and exhaustive bilateral proposition.

Given a model $\mathcal{M}$, we define extensions of $[\cdot]$ to all formulas of $\mathcal{L}_\Delta$, one for the conjunctive and one for the disjunctive conception of Δ, written with a C or a D as subscript. For φ an atom and A and B any formula of $\mathcal{L}_\Delta$,

- $[\varphi]_C = [\varphi]_D = [\varphi]$
- $[\top]_C = [\top]_D = (\{\square\}, \{\blacksquare\})$
- $[\neg A]_C = \neg[A]_C$ and $[\neg A]_D = \neg[A]_D$
- $[A \wedge B]_C = [A]_C \wedge [B]_C$ and $[A \wedge B]_D = [A]_D \wedge [B]_D$

[22] It is also very plausible to hold that the fusion of some wholly positive states must be wholly positive as well. As it stands, this is not guaranteed by our definition and our constraints on C-spaces. As we shall see, requiring C-spaces to satisfy this additional condition makes no difference to the logic, so I have opted against it.

[23] Leuenberger, naturally enough, uses T. I prefer Δ because of the relationship to the state-completion function δ (for which t would have been unfortunate, since $s, t, \ldots$ etc. are standardly used as variables ranging over states).

[24] It would not serve our purposes to introduce $\top$ as an abbreviation of an arbitrary tautology, since these, although necessarily true, are not in general verified by the nullstate.

- $[A \vee B]_C = [A]_C \vee [B]_C$ and $[A \vee B]_D = [A]_D \vee [B]_D$
- $[\Delta A]_C = \Delta_C[A]_C$ and $[\Delta A]_D = \Delta_D[A]_D$

Definition 12 Let $A \in \mathcal{L}_\Delta$ and $\Phi \subseteq \mathcal{L}_\Delta$.

- For $\mathcal{M}$ any model, $\mathcal{M} \models_C A$ iff $s \in [A]^+_C$ for some $s \sqsubseteq @$
- $\Phi \models_C A$ iff for all models $\mathcal{M}$, $\mathcal{M} \models_C A$ if $\mathcal{M} \models_C B$ for all $B \in \Phi$
- For $\mathcal{M}$ any model, $\mathcal{M} \models_D A$ iff $s \in [A]^+_D$ for some $s \sqsubseteq @$
- $\Phi \models_D A$ iff for all models $\mathcal{M}$, $\mathcal{M} \models_D A$ if $\mathcal{M} \models_D B$ for all $B \in \Phi$

To specify adequate deductive systems we proceed in two steps. First, we specify systems for establishing when two formulas A and B will be interchangeable in the scope of Δ due to their logical form. For disjunctive Δ, we can use a version of the system shown in [7] to prove $A \approx B$ exactly when A and B have the same truthmakers in every model, given that neither closure nor convexity is assumed.[25]

Collapse($\vee$)	$A \vee A \approx A$
Commutativity($\vee$)	$A \vee B \approx B \vee A$
Associativity($\vee$)	$A \vee (B \vee C) \approx (A \vee B) \vee C$
ECollapse($\wedge$)	$A \vee (A \wedge A) \approx A \wedge A$
$\top$Collapse($\wedge$)	$A \wedge \top \approx A$
$\bot$Collapse($\wedge$)	$A \wedge \bot \approx \bot$
Commutativity($\wedge$)	$A \wedge B \approx B \wedge A$
Associativity($\wedge$)	$A \wedge (B \wedge C) \approx (A \wedge B) \wedge C$
Collapse($\neg\neg$)	$\neg\neg A \approx A$
DeMorgan($\neg\vee$)	$\neg(A \vee B) \approx \neg A \wedge \neg B$
DeMorgan($\neg\wedge$)	$\neg(A \wedge B) \approx \neg A \vee \neg B$
Distributivity($\wedge/\vee$)	$A \wedge (B \vee C) \approx (A \wedge B) \vee (A \wedge C)$
Preservation($\vee$)	$A \approx B \ / \ A \vee C \approx B \vee C$
Preservation($\wedge$)	$A \approx B \ / \ A \wedge C \approx B \wedge C$
Symmetry	$A \approx B \ / \ B \approx A$
Transitivity	$A \approx B, \ B \approx C \ / \ A \approx C$

We write $\vdash_{eD} A \approx B$ iff A and B are formulas of $\mathcal{L}_\Delta$ for which $A \approx B$ can be derived using these axioms and rules.

For conjunctive Δ, we have to slightly strengthen this system. Recall that under the conjunctive conception, ΔA is true iff every verifier of A is between $@^p$ and $@$. This condition is satisfied by the set of verifiers of A iff it is satisfied by that set's closure under non-empty fusions. As a result, A and $A \wedge A$ are interchangeable in

[25] Krämer [7] uses a language that does not include the logical constant $\top$, so the logic does not have the axioms $\top$Collapse($\wedge$) and $\bot$Collapse($\wedge$). The soundness of these axioms is readily verified. The completeness proof in [7] is by disjunctive normal forms, and may be adapted to the present setting by adjusting the relevant notion of normal form, and simplifying conjunctions of literals to $\bot$ if $\bot$ is among them, and removing all occurrences of $\top$ as a literal otherwise.

the context of conjunctive Δ. So for that context, we replace ECollapse($\wedge$) by the stronger

Collapse($\wedge$) $A \wedge A \approx A$

We write $\vdash_{eC} A \approx B$ iff A and B are formulas of $\mathcal{L}_\Delta$ for which $A \approx B$ can be derived using the resulting set of axioms and rules.

Next, we turn to derivations within $\mathcal{L}_\Delta$. For the logic of conjunctive Δ, we extend any standard axiomatization of classical propositional logic by the instances of[26]

Δ-EquivC	$\Delta A \leftrightarrow \Delta B$ whenever $\vdash_{eC} A \approx B$
Δ-Fact	$\Delta A \rightarrow A$
Δ-PFix	$\Delta A \rightarrow \Delta\Delta A$
Δ-DisjC	$\Delta(A \vee B) \leftrightarrow (\Delta A \wedge \Delta B)$
Δ-AbsP	$(\Delta A \wedge \Delta(B \wedge C)) \rightarrow \Delta(A \wedge B)$
Δ-NE1	$\Delta(\neg\Delta A \wedge B) \rightarrow \Delta\neg\Delta\bot$
Δ-NE2	$\Delta\neg\Delta\bot \rightarrow (\Delta\neg A \leftrightarrow \Delta\neg\Delta A)$

and write $\Phi \vdash_C A$ iff $A \in \mathcal{L}_\Delta$ is derivable from $\Phi \subseteq \mathcal{L}_\Delta$ in the resulting system. We prove in the appendix that this system is sound and complete with respect to $\models_C$, so that $\Phi \vdash_C A$ iff $\Phi \models_C A$.

The final two axioms require some explanation. As we noted, $\Delta\top$ says that the actual world of the model is empty, so that $@ = \delta\Box$. $\Delta\neg\Delta\bot$ functions as a strengthening of this; it holds iff the actual world is empty, and the only δ-state aside from @ is $\blacksquare$. So Δ-NE1 says that the negation of a Δ-claim can only form part of a weakly complete truth in this sort of model, and Δ-NE2 says that in such a model, $\neg A$ is a weakly complete truth iff $\neg\Delta A$ is. Note that by instantiating Δ-NE2 with $\bot$ and using the fact that $\vdash_{eC} \top \approx \neg\bot$, we can derive $\Delta\neg\Delta\bot \rightarrow \Delta\top$.

For the logic of disjunctive Δ, we extend a classical propositional logic by the instances of

Δ-EquivD	$\Delta A \leftrightarrow \Delta B$ whenever $\vdash_{eD} A \approx B$
Δ-Fact	$\Delta A \rightarrow A$
Δ-PFix	$\Delta A \rightarrow \Delta\Delta A$
Δ-DisjD	$\Delta(A \vee B) \leftrightarrow (\Delta A \vee \Delta B)$
Δ-AbsT	$(\Delta A \wedge B) \rightarrow \Delta(A \wedge B)$
Δ-NFix	$\neg\Delta A \rightarrow \Delta\neg\Delta A$

and write $\Phi \vdash_D A$ iff $A \in \mathcal{L}_\Delta$ is derivable from $\Phi \subseteq \mathcal{L}_\Delta$ in the resulting system. We prove in the appendix that this system is sound and complete with respect to $\models_D$, so that $\Phi \vdash_D A$ iff $\Phi \models_D A$.[27]

[26] Strictly speaking, our language does not include the material conditional or bi-conditional, but we can regard them as meta-linguistic abbreviations in the usual way. Note that while the various logically equivalent candidates for formulas to abbreviate using $\rightarrow$ and $\leftrightarrow$ may differ with regard to their truth-makers, since the conditionals do not appear within the scope of Δ in our axioms, these hyperintensional differences do not lead to a difference in the logic.

[27] If one wished to endorse Collapse($\wedge$) even in the context of the logic for disjunctive Δ, one would simply need to replace Δ-EquivD by Δ-EquivC and modify Δ-DisjD to read $\Delta(A\vee B) \leftrightarrow (\Delta A \vee \Delta B \vee \Delta(A\wedge B))$. Disjunction would then need to be given the inclusive interpretation described above, and we would impose a general constraint to the effect that every unilateral proposition by closed under non-empty fusions.

6 Comparison

We now turn to the comparison between our hyperintensional logics of totality and Leuenberger's intensional one. As one might expect, there is significant overlap, but there are also significant differences, resulting from our hyperintensional orientation and, relatedly, our conjunctively or disjunctively distributive interpretation of Δ with respect to disjunctive arguments.

Following the suggestion in [1], Leuenberger takes ΔA to be true at a world w iff A is true at w, but at no world outstripped by w. In other words, w satisfies ΔA iff w is a minimal A-world. The conditions under which ΔA is true at a world w under our two semantics can also be formulated partially in terms of outstripping and minimality, where a world w is taken to outstrip another v iff $v^p \sqsubset w^p$. We can bring out the relationship between the accounts as follows, ordering them by increasing strength of the satisfaction condition for ΔA.

1. Under the disjunctive interpretation, w satisfies ΔA iff for some state s verifying A, w is a minimal s-world.
2. Under the intensional interpretation, w satisfies ΔA iff w is minimal among the worlds containing some verifier s of A.
3. Under the conjunctive interpretation, w satisfies ΔA iff for every state s verifying A, w is a minimal s-world.

This difference is reflected in the logical behaviour of Δ with respect to disjunction. As we observed, under the disjunctive interpretation, we have $\Delta(A \vee B) \leftrightarrow (\Delta A \vee \Delta B)$, and under the conjunctive interpretation we have $\Delta(A \vee B) \leftrightarrow (\Delta A \wedge \Delta B)$. The intensional interpretation is strictly in between, we have $\Delta(A \vee B) \rightarrow (\Delta A \vee \Delta B)$, but not the converse, and we have $(\Delta A \wedge \Delta B) \rightarrow \Delta(A \vee B)$, but not the converse. Indeed, it is clear that the converses are unacceptable under any intensional approach. For suppose ΔA. Then firstly, $\Delta A \vee \Delta\top$, so if we can infer from this, as under our disjunctive interpretation, that $\Delta(A \vee \top)$, then by intensionality it would follow that $\Delta\top$. And secondly, by intensionality we can infer from ΔA that $\Delta(A \vee (A \wedge \bot))$, from which we had better not be able to infer $\Delta(A \wedge \bot)$, as we could under the conjunctive interpretation. What we can infer, under Leuenberger's account, from $\Delta(A \vee B)$, is $A \rightarrow \Delta A$ and $B \rightarrow \Delta B$. So under this account, every *true* disjunct of a total truth is itself total (where A is called a total truth iff ΔA holds).

It is interesting also to consider under what conditions $\Delta(A \vee B)$ may be inferred from ΔA. Under the conjunctive interpretation, we may do so only if ΔB. Under the disjunctive interpretation, we may always do so. Under the intensional account, we may do so if, and only if, for every A-minimal world w, no world outstripped by w makes true B, i.e. if B cannot be made true with strictly less material, as it were, than is available in w. This will hold when B requires *more* than is available in w, but also when B just requires *something else* than is available in w. For example, suppose that w is a physicalist world, and that A states all the physical facts in w, so ΔA is true at w. Now let B say that there are angels. Then $\Delta(A \vee B)$ will also be true at w, since any B-world will contain angels, and no such world is outstripped by w.

I do not think that these observations constitute a compelling objection to Leuenberger's account. All three accounts constitute, from their respective semantic points of view, natural ways to handle disjunction in the scope of Δ. Something of an advantage may perhaps be claimed for the hyperintensional approaches in that the disjunctive reading does seem to align with a natural use of 'that's it' in some ordinary contexts, while the conjunctive one fits a reading of 'that's it' as expressing weak completeness. As the angels-example just presented suggests, it is not clear that there is any similarly intuitive reading that is tracked by the specific conditions under which, on the intensional account, we may infer $\Delta(A \vee B)$ from ΔA. But since it is not Leuenberger's aim to track the ordinary usage of 'that's it', or to characterize an operator expressing weak completeness, that is not by itself an objection to his account.

Let us consider the overall logic obtained under the intensional account. Leuenberger first presents a base system C, which he proves sound and complete relative to the class of all *totality frames*, which consist of a set of worlds W and some binary relation of outstripping $< \quad \subseteq W \times W$, with ΔA being true at a world w iff A is true at w and at no world outstripped by w. The system C consists of the propositional tautologies, the rule of modus ponens, as well as the following axioms and rules distinctively concerned with Δ:

(A1) $\Delta A \rightarrow A$
(A2) $\Delta A \wedge \Delta B \rightarrow \Delta(A \vee B)$
(RIM) If $\vdash A \rightarrow B$, then $\vdash A \rightarrow (\Delta B \rightarrow \Delta A)$

The intended relation of outstripping, much like the relation of proper part, is clearly a strict partial order, i.e. transitive and asymmetric. Leuenberger shows that C is not complete with respect to the class of partial order frames, but that a sound and complete axiomatization of the class of the partial order frames, and indeed the wider class of transitive frames, is obtained by adding to C the somewhat hard to interpret axiom schema

(A3) $\Delta A \rightarrow \Delta(A \leftrightarrow (\Delta B \rightarrow \Delta(A \vee B)))$

He calls the resulting system $C3$.

(A1) and (A2) are of course valid in both our logics. I shall set aside the 'transitivity axiom' (A3). As Leuenberger highlights, its connection to the transitivity of outstripping is not exactly transparent, as a result of the fact that Δ is related to outstripping in a tight but somewhat complicated way. Moreover, to make the connection to transitivity, (A3) exploits the distinctive behaviour of Leuenberger's Δ with respect to disjunctive arguments, and for this reason it does not say in our contexts quite what it says in the intensional context. In particular, it does not appear to be tied in any clear way to the transitivity of outstripping, which is guaranteed under the proposed definition of w outstripping v as $w^p \sqsupset v^p$. (RIM)—the Rule of Inverse Monotonicity—is valid in neither. Nor should it be; it is unacceptable given our general hyperintensional orientation regarding Δ. For (RIM) says that any truth that logically entails a total truth is itself a total truth, which means that in constructing a total truth, logical consequences come for free, contrary to our hyperintensionalism.

In addition to the stated axioms and rules, Leuenberger considers, but does not endorse, some additional axioms corresponding to the existence of unrestrictedly minimal worlds and related constraints. They are:

(A4) $\neg\Delta\top$
(A5) $\neg\Delta\neg\Delta A$
(A6) $\Delta\neg\Delta A \to \Delta\top$

In effect, within the context of Leuenberger's semantics, (A4) says that no world is minimal, (A5) says that every world outstrips a world which itself outstrips a world, and (A6) says that every non-minimal world outstrips a world that outstrips a world.

Since we are giving $\top$ a non-disjunctive interpretation, (A4)'s status is the same under both our accounts, and it says something close to what it says in the intensional setting: that the actual world is not null. Since any world in a C-space is eligible to be the actual world in a model, the validity of (A4) comes down to the condition on C-spaces that $\delta\Box \notin S^{\Diamond}$. We might impose such a constraint, but, like Leuenberger, I see no clear reason to exclude the empty world. (Such an exclusion seems especially questionable when we consider applications in which worlds merely represent complete ways for things to stand with respect to the possibly restricted subject matter of the language under consideration.)

It is worth noting that we can use (A4) to exclude the empty world, without thereby prohibiting *minimal* worlds. Let At be a set of atoms and consider the canonical C-space based on At, At in which exactly the atoms are regarded as positive literals. We may exclude the empty world by regarding the set of all negative literals as inconsistent, but then have many minimal worlds, namely the worlds with exactly one atom as member. So this is one way in which our framework allows us to express somewhat more fine-grained distinctions than the intensional one regarding the structure of the space of worlds in terms of outstripping and positivity.

(A5) and (A6) concern the potential of negated Δ-claims to be total truths, so both their status and what they say about minimal worlds differs strongly between our two approaches. Under the conjunctive interpretation, (A5) is invalid, but valid in the class of models with more than two δ-states, and hence in the class of models in which some world is non-empty. (A6) is valid, and registers the fact that the negation of a Δ-claim can only be a weakly complete truth if the actual world is empty.

Under the disjunctive interpretation, (A5) and (A6) are invalid, and indeed $\Delta\neg\Delta A$ holds whenever $\neg\Delta A$. This is another instance in which the different orientation towards disjunction makes comparison between the accounts difficult. For us, the fact that $@^{p}$ is an exact falsifier of any given false statement ΔA means that $\Delta\neg\Delta A$ must then be true. Now, as we have observed, $\Delta\neg\Delta A$ will then not normally be the complete truth, but we can even give an example of a C-space without minimal worlds and a selection of a world as actual so that a proposition of the form $\Delta\neg\Delta A$ is a complete truth.[28]

[28] Let S be the set of natural numbers plus ω, and let $\sqsubseteq$ be $\geq$, so ω is the nullstate, and there is no minimal non-null state. Now modify the simple C-space as defined above simply by letting $\delta\omega$ be inconsistent, so that while there is a minimal δ-state, there is no minimal world. Let $@ = \delta 0$ and let **P** be verified by all and only the proper parts of @, and falsified by @ and only @. Then @ is the only δ-state not verifying Δ**P**, and therefore verifying $\neg\Delta$**P**, and so @ will be the sole verifier of $\Delta\neg\Delta$**P**.

Having discussed how the principle Leuenberger considers fare within our system, let us briefly turn to the converse question of how our axioms fare within Leuenberger's logics. We have already seen that the disjunction-rules Δ-DisjC and Δ-DisjD and the negated-Δ-rules Δ-N⊥ and Δ-NFix are not valid and why. The remaining rules are readily seen to hold within Leuenberger's base system C.

And that's it.

Appendix A: Soundness

Theorem 4 *Let* $\Phi \subseteq \mathcal{L}_\Delta$ *and* $\varphi \in \mathcal{L}_\Delta$*. Then* $\Phi \vdash_C \varphi$ *implies* $\Phi \models_C \varphi$*.*

Proof The soundness of classical propositional logic is easily established using the fact that @ is a world, and that any formula of $\mathcal{L}_\Delta$ is assigned an exclusive and exhaustive proposition. It remains to show that our additional axioms are true in every model.

The soundness of Δ-EquivC follows from the soundness for $\vdash_{eC}$ for the interpretation of $A \approx B$ as true iff the sets of verifiers of A and B have the same closure under (non-empty) fusions and the fact that the truth-value of ΔA and ΔB depends only on the closure under fusion of the verifier-sets of A and B.

For Δ-Fact and Δ-PFix, assume ΔA is true in a model, so some part s of @ is the sole verifier of ΔA. By the clause for Δ, some δ-state is part of s. But the only δ-state contained in @ is @, so @ is the sole verifier of ΔA. By the clause for Δ again, A has a verifier and every verifier of A is part of @, so A is true in the model, and @ is also the sole verifier of $\Delta\Delta A$, so $\Delta\Delta A$ is also true in the model.

For Δ-DisjC, note that $\Delta(A \vee B)$ is true in a model iff @ is the sole verifier of $\Delta(A \vee B)$, which holds iff @ is the fusion of all δs where s verifies $A \vee B$. This is so just in case $@ = \delta s$ for all s verifying $A \vee B$ and hence iff $@ = \delta s$ for all s verifying A and for all s verifying B, so iff @ is the sole verifier of both ΔA and ΔB, and hence of their conjunction, so iff $\Delta A \wedge \Delta B$ is true in the model.

For Δ-AbsP, suppose $\Delta A \wedge \Delta(B \wedge C)$ is true in a model and hence verified solely by @. Then the same holds for both ΔA and $\Delta(B \wedge C)$. It follows that every verifier of A, and every verifier of $B \wedge C$ has @ as its completion, and so that every verifier of B is part of @. But then using δ-Identity(1), we may infer that every verifier of $A \wedge B$ has @ as its completion, and thus that $\Delta(A \wedge B)$ is verified by @ and hence true in our model.

For Δ-NE1, suppose $\Delta(\neg\Delta A \wedge B)$ is true in a model. Then $\neg\Delta A$ is verified only by parts of @. Since it is verified by any incompletion of any verifier of A, @ must be the only incompletion of a verifier of A. So every δ-state other than @ must be the completion of every verifier of A, so there can only be two δ-states. Since ■ is a δ-state distinct from @, it follows that @ and ■ are the only δ-states. By δ-Identity(1), $\delta\square$ cannot be ■, and so must be @. Then $\Delta\bot$ is falsified by the one incompletion of

$\blacksquare$, which will be @, and by $\square$, the falsifier of $\bot$. Both have @ as their completion, so $\Delta\neg\Delta\bot$ is verified by @, and hence true in the model.[29]

For Δ-NE2, suppose $\Delta\neg\Delta\bot$ is true in a model. By the same reasoning as before, the model can only have two δ-states, with $@ = \delta\square$. Now suppose $\Delta\neg A$ is true in the model, so every falsifier of A is part of @. Then every verifier of A has $\blacksquare$ as its completion, so @ is the only incompletion of any verifier of A. It follows that $\neg\Delta A$ is verified only by parts of @, and thus by states that have @ as their completion, so $\Delta\neg\Delta A$ is true in the model. Conversely, suppose $\Delta\neg\Delta A$ is true in the model. Then all verifiers of $\neg\Delta A$, and hence all falsifiers of A, must be parts of @. But then they all have @ as their completion, so $\Delta\neg A$ is true in the model. □

Theorem 5 *Let* $\Phi \subseteq \mathcal{L}_\Delta$ *and* $\varphi \in \mathcal{L}_\Delta$. *Then* $\Phi \vdash_D \varphi$ *implies* $\Phi \models_D \varphi$.

Proof The argument for the soundness of classical propositional logic remains the same. The soundness of Δ-EquivD is immediate from soundness for $\vdash_{eD}$ for the interpretation of $A \approx B$ as true iff the sets of verifiers of A and B are the same and the fact that the truth-value of ΔA and ΔB depends only on the verifiers of A and B. Δ-DisjD is immediate by the clause for Δ. Δ-Fact and Δ-PFix are immediate given the fact that always $\delta s \sqsupseteq s$ and $\delta\delta s = \delta s$.

For Δ-NFix, suppose $\neg\Delta A$ is true, and so verified by some part s of @. Since falsifiers of ΔA are all δ-states, $s = @$. Since @ is a world, $\delta @ = @$, and so @ verifies $\Delta\neg\Delta A$.[30]

For Δ-AbsT, suppose part s of @ verifies $\Delta A \wedge B$. Then $s = \delta t \sqcup u$ for some t and u verifying A and B, respectively. Since @ is consistent, s is consistent, and so δt is consistent, and therefore a world. And since $\delta t \sqsubseteq @$, it follows that $\delta t = @$, und $u \sqsubseteq @$. Then also $t \sqsubseteq @$ and hence $t \sqcup u \sqsubseteq @$. So $t \sqsubseteq t \sqcup u \sqsubseteq @ = \delta t$, and hence $\delta(t \sqcup u) = @$. □

Appendix B: Normal Forms

We establish completeness via normal forms. For each logic, we prove a reduction theorem to the effect that any formula of $\mathcal{L}_\Delta$ is provably equivalent to a formula in a kind of disjunctive normal form.

For these purposes, we count as *atoms* the members of At, $\top$, as well as—somewhat unusually—the formulas $\Delta\top$ and $\Delta\bot$. Any atom and any negation of an atom is a *literal*; the members of At and their negations are the non-logical literals and their set is denoted by Lit. We use the variables φ, ψ, χ as well as variously adorned versions of them to range over the literals of $\mathcal{L}_\Delta$. The *converse* φ' of a

[29]Note that $\Delta\neg\Delta\bot$ will still be true if we countenance more falsifiers of Δ-statements, as long as these are all parts of an incompletion of a verifier of the argument. For in the present case, this will mean that these added falsifiers are still parts of @, and hence have @ as their completion.

[30]Note that we do not require the assumption that only δ-states falsify ΔA to obtain the result; it is enough that *some* state t with $\delta t = @$ falsifies ΔA. So as long as @, or at least its positive part, is included among the falsifiers of any given false ΔA, Δ-NFix will be sound.

literal φ is $\neg\varphi$ if is an atom, and ψ if $\varphi = \neg\psi$. Formulas of the form $\Delta(_1 \wedge \ldots \wedge \varphi_n)$ and their negations will be called Δ-*literals*. A $\Delta - DNF$ is any disjunction of conjunctions in which each conjunct is either a literal or a Δ-literal.

We will show that under either logic, every formula of $\mathcal{L}_\Delta$ is provably equivalent to a Δ-DNF. To do this, we classify formulas of $\mathcal{L}_\Delta$ in accordance with their maximum nesting depth of Δ. For each n, we define a set of n-atoms, n-literals, and n-formulas. The 0-atoms are the atoms, the 0-literals are the 0-atoms and their negations, and the 0-formulas are all the formulas built up from 0-atoms without the use of Δ. The $(n+1)$-atoms are the n-atoms as well as all expressions of the form ΔA, where A is an n-formula. The $(n+1)$-literals are the $n+1$-atoms and their negations, and the $(n+1)$-formulas are all the formulas built up from $(n+1)$-atoms without the use of Δ. A conjunction of n-literals will be called an n-CF, and a disjunction of n-CFs will be called an n-DNF. The *degree* of a formula A is the smallest n for which A is an n-formula.

B.1 Conjunctive Δ

Throughout this subsection, reference to provability is to be understood in terms of $\vdash_C$, and reference to provable exact equivalence in terms of $\vdash_{eC}$. We first establish some useful facts in preparation of the reduction theorem.

Lemma 4

1. *Any n-formula is provably exactly equivalent to an n-DNF.*
2. *For A an n-formula, ΔA is provably equivalent to a formula $\Delta A'$, where A' is an n-DNF.*
3. $\Delta(\Delta A \wedge B)$ *is provably equivalent to* $\Delta A \wedge \Delta(A \wedge B)$
4. $\Delta(\neg\Delta A \wedge B)$ *is provably equivalent to* $\Delta\neg\Delta\bot \wedge \Delta\neg A \wedge \Delta B$.
5. *If A is an n-CF of degree $n > 0$, then ΔA is provably equivalent to some n-formula.*
6. *For $n > 0$, if A is a formula of degree n, then ΔA is provably equivalent to some n-formula.*

Proof (1): This is a straightforward consequence of Theorem 15 in [3, p. 214].

(2): By (1) and Δ-EquivC.

(3): By Δ-Fact, Δ-PFix, and Δ-AbsP.

(4): Left-to-right: By Δ-NE1, from $\Delta(\neg\Delta A \wedge B)$ we infer $\Delta\neg\Delta\bot$, from which we obtain, using Δ-NE1, $\Delta\top$. Using $\vdash_{eC}\ A \approx A \wedge \top$ and Δ-AbsP, we can derive every instance of $\Delta\top \to (\Delta(A \wedge B) \to \Delta A))$ which allows us to infer both $\Delta\neg\Delta A$ and ΔB from $\Delta(\neg\Delta A \wedge B)$. Using Δ-NE2 again, we may infer $\Delta\neg A$. Right-to-left: Using Δ-NE2, we may infer $\Delta\neg\Delta A$ from $\Delta\neg\Delta\bot$ and $\Delta\neg A$. Using Δ-AbsP and ΔB, we obtain $\Delta(\neg\Delta A \wedge B)$.

(5): Suppose A is an n-CF of degree $n > 0$. So A is a conjunction of n-literals $A_1, \ldots, A_m$, at least one of which is of degree n. Let k be the number of degree n literals among $A_1, \ldots, A_m$. We may then use an 'extraction' procedure to obtain a formula provably equivalent to ΔA, which is a conjunction of an n-formula and a sentence $\Delta A'$ with A' an n-CF with $k-1$ degree n literals among its conjuncts. So by repeated application of that procedure, we obtain a sentence provably equivalent to ΔA which is a conjunction of n-formulas, and hence itself an n-formula.

The extraction procedure is as follows. First, we pick one of the degree n literals among $A_1, \ldots, A_m$, call it B_1, and turn A into a provably exactly equivalent n-CF of degree n $B = B_1 \wedge \ldots \wedge B_m$. Then ΔA is provably equivalent to ΔB. B_1 is either of the form ΔC or of the form $\neg\Delta C$. In the former case, ΔB is $\Delta(\Delta C \wedge B_2 \wedge \ldots \wedge B_m)$, and hence by (3) provably equivalent to $\Delta C \wedge \Delta(C \wedge B_2 \wedge \ldots \wedge B_m)$. Since C is an $(n-1)$-formula, ΔC is an n-formula and $\Delta(C \wedge B_2 \wedge \ldots \wedge B_m)$ has only $k-1$ degree n literals, as desired. In the latter case, ΔB is $\Delta(\neg\Delta C \wedge B_2 \wedge \ldots \wedge B_m)$, and hence by (4) provably equivalent to $\Delta\neg\Delta\bot \wedge \Delta\neg C \wedge \Delta(B_2 \wedge \ldots \wedge B_m)$. Since C is an $(n-1)$-formula, $\Delta\neg\Delta\bot \wedge \Delta\Box C$ is an n-formula, and $\Delta(B_2 \wedge \ldots \wedge B_m)$ has only $k-1$ degree n literals, as desired.

(5): By (2), ΔA is provably equivalent to ΔB for some n-DNF $B = B_1 \vee \ldots \vee B_k$. Then by Δ-DisjC, ΔA is provably equivalent to $\Delta B_1 \wedge \ldots \wedge \Delta B_k$. Now consider each conjunct ΔB_i. If it is of degree $< n$, we do nothing. If it is of degree n, then by (3) it may be replaced by a provably equivalent n-formula. In this way we obtain an n-formula provably equivalent to ΔA, as desired. □

Theorem 6 *Any formula is provably equivalent within* $\vdash_C$ *to a* Δ*-DNF.*

Proof The case in which A is of degree 0 is trivial. Suppose that A is degree 1, and consider any degree 1 atom in A. This will be of the form ΔB, with B a 0-formula. By Lemma 6 (2), there is a 0-DNF C with ΔB provably equivalent to ΔC. By Δ-DisjC, ΔC in turn is provably equivalent to a conjunction of Δ-literals. So we may replace any degree 1 atom in A by a provably equivalent conjunction of Δ-literals. The result is provably equivalent to A, and, by propositional logic, provably equivalent to a Δ-DNF.

Now suppose A is degree $n+1$ with $n > 0$, and suppose the claim holds for all formulas up to degree n (IH). Consider any degree $n+1$ atom in A, which will be of the form ΔB, with B a degree n-formula. Then by Lemma 6 (6), ΔB is provably equivalent to some n-formula. We may thus replace each degree $n+1$ atom in A by a provably equivalent n-formula. The result will itself be an n-formula provably equivalent to A. By IH, that formula, and therefore A, is provably equivalent to a Δ-DNF. □

B.2 Disjunctive Δ

Throughout this subsection, reference to provability is to be understood in terms of $\vdash_D$, and reference to provable exact equivalence in terms of $\vdash_{eD}$.

Lemma 5

1. *Any* n*-formula is provably exactly equivalent to an* n*-DNF.*
2. *For* A *an* n*-formula,* ΔA *is provably equivalent to a formula* $\Delta A'$*, where* A' *is an* n*-DNF.*

3. $\Delta(\Delta A \wedge B))$ *is provably equivalent to* $\Delta A \wedge B$.
4. $\Delta(\neg\Delta A \wedge B))$ *is provably equivalent to* $\neg\Delta A \wedge B$.
5. *If* A *is an* n*-CF of degree* $n > 0$*, then* ΔA *is provably equivalent to* A.
6. *For* $n > 0$*, if* A *is a formula of degree* n*, then* ΔA *is provably equivalent to some* n*-formula.*

Proof (1) and (2) are as before. (3) and (4) may be established using Δ-Fact, Δ-AbsT, and Δ-PFix or Δ-NFix, respectively.

(5): A has at least one conjunct of degree $n > 0$, i.e. an n-literal of the form ΔB or $\neg\Delta B$. In the former case, A is provably exactly equivalent to some conjunction of the form $\Delta B \wedge B'$. So by Δ-EquivD, ΔA is provably equivalent to $\Delta(\Delta B \wedge B')$, which by (3) is provably equivalent to $\Delta B \wedge B'$, which is provably equivalent to A. In the latter case, A is provably exactly equivalent to some conjunction of the form $\neg\Delta B \wedge B'$. So by Δ-EquivD, ΔA is provably equivalent to $\Delta(\neg\Delta B \wedge B')$, which by (4) is provably equivalent to $\neg\Delta B \wedge B'$, which is provably equivalent to A.

(6): By (2), ΔA is provably equivalent to ΔB for some n-DNF $B = B_1 \vee \ldots \vee B_k$. Then by Δ-DisjD, ΔA is provably equivalent to $\Delta B_1 \vee \ldots \vee \Delta B_k$. Now consider each disjunct ΔB_i. If it is of degree $< n$, we do nothing. If it is of degree n, then we replace it by B_i. Since B_i is an n-CF, by (5), ΔB_i is provably equivalent to B_i. So through these replacements we obtain an n-formula provably equivalent to ΔA, as desired. □

Theorem 7 *Any formula is provably equivalent within* $\vdash_D$ *to a* Δ*-DNF.*
Proof The reasoning is as for Theorem 7, using Δ-DisjD in place of Δ-DisjC and Lemma 8 (6) in place of Lemma 6 (6). □

Appendix C: Completeness

Since our logics extend classical propositional logic, strong completeness—that $\Phi \vdash \varphi$ whenever $\Phi \models \varphi$—follows from the claim that every consistent subset of $\mathcal{L}_\Delta$ has a model. As usual, any consistent subset of $\mathcal{L}_\Delta$ can be extended to a maximal consistent subset of $\mathcal{L}_\Delta$, so it suffices to show for each logic that every maximal consistent (within that logic) subset of $\mathcal{L}_\Delta$ has a model. So in this section we show how to construct appropriate canonical models. We start with the strong, conjunctive interpretation of Δ, since the construction is somewhat simpler in this case.

C.1 Conjunctive Δ

Let Φ be a subset of $\mathcal{L}_\Delta$ that is maximal consistent with respect to $\vdash_C$. We begin by defining some important notions in terms of Φ.

Definition 13 Let $\varphi \in Lit$ and A a 0-CF in $\mathcal{L}_\Delta$. Then

- A is *complete* iff $\Delta A \in \Phi$
- A is *minimally complete* iff A is complete, and the result of replacing any conjunct of A by $\top$ is not complete
- $m(A)$ is the set of non-logical literals occurring as a conjunct in A
- φ is *actual* iff $\varphi \in \Phi$
- φ is *pure* iff φ is a conjunct in some complete 0-CF
- if φ is actual: φ is *positive* iff either there is no complete 0-CF, or φ is a conjunct in some minimally complete 0-CF; φ is *negative* otherwise
- if φ is non-actual: φ is *positive* iff its actual converse φ' is negative; φ is *negative* otherwise
- $n\varphi$ is the negative member of $\{\varphi, \varphi'\}$

We write @ for the set of actual non-logical literals and $@^p$ for the set of positive members of @. Note that $@^p$ is empty just in case $\Delta\top \in \Phi$, in which case there is a complete 0-CF, but no minimally complete 0-CF. We denote the set of all (non-logical)[31] literals by P (since we now regard them as *proto*-states). For $s \subseteq P$, we say that s is minimally complete iff some conjunction of all its members is, and we let $f(s)$ be s if no subset of s is minimally complete, and $s \cup @^p$ otherwise. Let S—the set of *states*—be $\{f(s) : s \subseteq P\}$. (The point of excluding from the set of states those sets of literals which have a minimally complete subset but do not include $@^p$ is that this will make it easy to ensure that all minimally complete 0-CFs are verified by $@^p$.)

Definition 14 The *canonical state-space* for Φ is $(S, \sqsubseteq)$ with $\sqsubseteq = \subseteq \upharpoonright S$

In preparation of the proof that $(S, \sqsubseteq)$ is indeed a state-space, we note some useful facts about f and S.

Lemma 6 *Let $s, t \subseteq P$.*

1. $@ \in S$
2. $@^p \in S$
3. $f(s) \subseteq @^p$ *if* $s \subseteq @^p$
4. $f(s) \subseteq @$ *if* $s \subseteq @$
5. $S = \{s \subseteq P$*: $@^p \subseteq s$ or s has no minimally complete subset*$\}$
6. *S is closed under intersection*
7. *If $s \neq t$ are minimally complete and s has at least two members, then S is not closed under union*
8. $f(s) \subseteq f(t)$ *if* $s \subseteq f(t)$
9. *For any $T \subseteq S$, $f \bigcup T$ is $\bigsqcup T$, the least upper bound of T in S*
10. *If $T \subseteq wp(P)$, then* $\bigsqcup\{f(t) : t \in T\} = f \bigcup T$

[31] This qualification will often remain tacit in what follows.

Proof (1) and (2) are obvious from the definition of f.

(3) and (4) are immediate from the fact that $f(s)$ extends s at most by members of $@^p$, and $@^p \subseteq @$.

(5): It is immediate from the definition of f that $f(s)$ always includes $@^p$ or has no minimally complete subset, and that $s \subseteq P$ is equal to s and hence a member of S whenever s includes $@^p$ or has no minimally complete subset.

(6): Let $T \subseteq S$. Either some member of T has no minimally complete subsets, in which case the same is true of $\bigcap T$, or all members of t have $@^p$ as a subset, in which case again the same is true of $\bigcap T$.

(7): Let $\varphi_i \in s$. Then $\{\varphi_i\}$ and $s \backslash \{\varphi_i\}$ have no minimally complete subset, but their union s does. Since t is also minimally complete and distinct from s, $s \not\supseteq @^p$. So $\{\varphi_i\}$ and $s \backslash \{\varphi_i\}$ are members of S but their union s is not.

(8): Suppose $s \subseteq f(t)$. If s has no minimally complete subset, then $f(s) = s$, so $f(s) \subseteq f(t)$. So suppose that s does have a minimally complete subset. Then so does $f(t)$, and hence $@^p \subseteq f(t)$. It follows that $f(s) = s \cup @^p \subseteq f(t)$.

(9): Let $T \subseteq S$. $f \bigcup T$ is clearly a member of S and an upper bound of T. So let $f(x) \in S$ be any upper bound of T in S. Then $\bigcup T \subseteq f(x)$, and so by (8), $f \bigcup T \subseteq f(x)$.

(10): By (9), $\bigsqcup\{f(t) : t \in T\}$ is $f \bigcup\{f(t) : t \in T\}$. Suppose first that no $t \in T$ has a minimally complete subset. Then $f(t) = t$ for all $t \in T$, and hence $\bigcup\{f(t) : t \in T\} = \bigcup T$, so $f \bigcup\{f(t) : t \in T\} = f \bigcup T$. Suppose now that at least one $t \in T$ has a minimally complete subset. Then $f \bigcup\{f(t) : t \in T\} = f(\bigcup T \cup @^p) = \bigcup T \cup @^p$. But then also $\bigcup T$ has a minimally complete subset, so also $f \bigcup T = \bigcup T \cup @^p$. □

Lemma 7 *The canonical state-space for* Φ, $(S, \sqsubseteq)$, *is a state-space.*

Proof From the fact that the subset-order on any set is a partial order together with Lemma 10 (9). □

Next, say that a member s of S is *consistent* iff there is no literal φ such that $\{\varphi, \varphi'\} \subseteq s$ and let $S^\Diamond = \{s \in S : s \text{ is consistent}\}$. Say that $s \in S$ *decides* a given literal φ iff either φ or φ' is a member of s, and let $\delta s = s \cup \{n\varphi : \varphi \in P \text{ and } s \text{ does not decide } \varphi\}$.

Lemma 8 $(S, S^\Diamond, \sqsubseteq, \delta)$ *is a C-space*

Proof It is clear that $S^\Diamond$ is non-empty and closed under part. Next, we show that δs is a world if s is consistent. So suppose s is consistent. Then clearly δs is also consistent. Now let t be any state not contained in δs. Let φ be a literal which is in t but not in δs. Suppose for reductio that $\varphi' \notin \delta s$. Then s does not decide φ, and so $n\varphi \in \delta s$. But $n\varphi$ is either φ or φ'. Contradiction. So $\varphi' \in \delta s$ and hence δs is incompatible with t. It follows that δs is a world. Moreover, δ-Containment is immediate from the definition of δs. So we may conclude that every consistent state is part of a world—and hence the conditions on a W-space are satisfied—, and that δ-Completeness holds.

δ-Redundancy(1): Suppose s is a world. Suppose for reductio that s does not decide φ. Then at least one of $s \cup \{\varphi\}$ and $s \cup \{\varphi'\}$ is a state, and hence a consistent state, contrary to s being a world. So s decides every literal, and hence by definition, $\delta s = s$.

δ-Redundancy(2): Immediate from the fact that by construction, δs decides every literal.

δ-Identity(1): Suppose $s \sqsubseteq t \sqsubseteq \delta s$. Suppose that $\varphi \in \delta t$. If $\varphi \in t$ then $\varphi \in \delta s$ is immediate. If $\varphi \notin t$, then $\varphi = n\varphi$ with t not deciding φ. But then s does not decide φ, and hence $\varphi \in \delta s$. Suppose $\varphi \in \delta s$. If $\varphi \in t$, then $\varphi \in \delta t$ by δ-Containment. If $\varphi \notin t$, then $\varphi = n\varphi$ with s not deciding φ. We need to show that t does not decide φ. Suppose otherwise. Then since $\varphi \notin t$, it follows that $\varphi' \in t$. Then $\varphi' \in \delta s$. Since $\varphi' \neq n\varphi$, it follows that $\varphi' \in s$. But then s decides φ. Contradiction. So t does not decide φ, and hence $\varphi = n\varphi \in \delta t$. So $\delta s = \delta t$, as desired.

δ-Identity(2): Suppose $\delta s = \delta t$. Suppose $\varphi \in \delta s$. Suppose first that $\varphi \in s$. If also $\varphi \in t$, then $\varphi \in s \sqcap t$ and so $\varphi \in \delta(s \sqcap t)$. If $\varphi \notin t$, then $\varphi = n\varphi$ and t does not decide φ. But then $s \sqcap t$ also does not decide φ, and so again $\varphi \in \delta(s \sqcap t)$. Suppose $\varphi \in \delta(s \sqcap t)$. Clearly if $\varphi \in s \sqcap t$, then $\varphi \in \delta s$. If $\varphi \notin s \sqcap t$, then $\varphi = n\varphi$ and $s \sqcap t$ does not decide φ. If s also does not decide φ, then $\varphi \in \delta s$. If s decides φ by having φ as member, then again $\varphi \in \delta s$. So suppose s decides φ by having φ' as a member. Then $\varphi' \in \delta s$ and hence $\varphi' \in \delta t$. Since $\varphi = n\varphi$, it follows that $\varphi' \in t$, and hence that $\varphi' \in s \sqcap t$. But then $s \sqcap t$ decides φ. Contradiction. So s does not have φ' as a member, and hence $\varphi \in \delta s$. □

It is straightforward to verify that @ is a world in this C-space, that $\delta s = @$ iff $@^p \sqsubseteq s \sqsubseteq @$ and that $@^p$ is the positive part of @ under the definition proposed in Section 4.5. Moreover, a state is wholly positive under the definition from that section just in case it does not contain $n\varphi$ for any φ, so the set of wholly positive states is closed under fusion, as one would expect.

We now define the interpretation function for the atoms; it is straightforward that it satisfies the conditions of exclusivity and exhaustivity.

Definition 15 For $\varphi \in At$, let

- $[\varphi]^+ = \{f\{\varphi\}\}$ if φ is pure and $\{f\{\varphi\}, \blacksquare\}$ otherwise
- $[\varphi]^- = \{f\{\neg\varphi\}\}$ if $\neg\varphi$ is pure and $\{f\{\neg\varphi\}, \blacksquare\}$ otherwise

Lemma 9 *For all $\varphi \in At$,*

1. *s and t are incompatible whenever $s \in [\varphi]^+$ and $t \in [\varphi]^-$*
2. *for every world $w \in S$, there is some $s \in [\varphi]^+ \cup [\varphi]^-$ with $w \sqsupseteq s$*

Definition 16 The *canonical model* $\mathcal{M}$ of Φ is $(S, S^\Diamond, \sqsubseteq, \delta, @, [\cdot])$

Lemma 10 *$\mathcal{M}$ is a model.*

Proof By Lemmas 12 and 13. □

We now show that in the canonical model, the literals and Δ-literals in Φ are verified by parts of @. Note that every formula of $\mathcal{L}_\Delta$ has a verifier in the canonical model.

Lemma 11 *Let $\mathcal{M}$ be the canonical model, φ any literal and A any 0-CF. If $\Delta\neg\Delta\bot \notin \Phi$, then*

1. $\mathcal{M} \models \varphi$ *if* $\varphi \in \Phi$.
2. $\mathcal{M} \models \Delta A$ *if* $\Delta A \in \Phi$.
3. $\mathcal{M} \models \neg\Delta A$ *if* $\neg\Delta A \in \Phi$.

Proof (1): Let $\varphi \in \Phi$. Suppose first that φ is a non-logical literal. Then $\varphi \in @$. By definition of $[\cdot]$, φ is verified by $f\{\varphi\}$. By definition of f, either $f\{\varphi\} = \{\varphi\}$ or $f\{\varphi\} = @^p$. Either way, $f\{\varphi\} \sqsubseteq @$, and so $\mathcal{M} \models \varphi$. Suppose now that φ is a logical literal. $\bot$ and $\Delta\bot$ are inconsistent, so there are four logical literals that may occur in Φ. (i) $\top$: $\Box$ verifies $\top$ by definition of $[\cdot]$. (ii) $\neg\Delta\bot$: Since $\Box$ falsifies $\bot$, and every falsifier of $\bot$ falsifies $\Delta\bot$, $\Box$ verifies $\neg\Delta\bot$. (iii) $\Delta\top$: By definition, $@^p$ is empty if $\Delta\top \in \Phi$, so then $@^p = \Box$ and hence $\Delta\top$ is verified by $@ = \delta\Box$. (iv): $\neg\Delta\top$: By definition, $@^p$ is non-empty if $\neg\Delta\top \in \Phi$, so @ is an incompletion of $\Box$ and hence verifies $\neg\Delta\top$.

(2): Suppose $\Delta A \in \Phi$. Then A is complete, so every non-logical literal φ occurring as a conjunct in A is pure, and therefore is verified only by $f\{\varphi\}$. By Lemma 10 (10) and the clause for conjunction, the sole verifier of the conjunction of non-logical conjuncts of A is $fm(A)$. Note that since $\Delta\neg\Delta\bot \notin \Phi$, the only logical literals that can occur in A are $\top$ and $\Delta\top$. Suppose first that $\Delta\top \in \Phi$, so $@^p = \Box$, and $\Delta\top$ is verified only by @. Then $fm(A) = m(A)$. By Δ-Fact and conjunction elimination, every literal in $m(A)$ is in Φ, so $m(A) \sqsubseteq @$. Since $\top$ is verified only by $\Box$ and $\Delta\top$ only by @, it follows that the sole verifier of A is part of @, and hence ΔA is verified by @. Suppose now that $\Delta\top \notin \Phi$. Then $\Delta\top$ also does not occur as a conjunct in A, and so the sole verifier of A is $fm(A)$. Since A is complete but $\top$ is not complete, some subset of $m(A)$ is minimally complete, so $fm(A) = m(A) \cup @^p$. By Δ-Fact and conjunction elimination, every literal in $m(A)$ is in Φ, so $m(A) \subseteq @$. It follows that $@^p \sqsubseteq fm(A) \sqsubseteq @$, and since $fm(A)$ is the sole verifier of A, that @ verifies ΔA. So $\mathcal{M} \models \Delta A$.

(3): Suppose $\neg\Delta A \in \Phi$. By the falsifier clause for Δ, @ will be a verifier of $\neg\Delta A$ provided that @ is not the sole verifier of ΔA. So suppose for contradiction that @ is the sole verifier of ΔA, and hence that every verifier of A is between $@^p$ and @. It follows that no literal that occurs as a conjunct in A has a non-actual verifier. So every non-logical literal $\varphi \in m(A)$ is pure, and hence occurs in some complete 0-CF B_φ, so that $\Delta B_\varphi \in \Phi$. Any such φ is then verified only by $f\{\varphi\}$, so by Lemma 10 (10), the sole verifier of their conjunction is $fm(A)$. Moreover, the only logical literals possibly occurring in A are $\top$ and $\Delta\top$, and the latter can occur in A only if it is in Φ, since it has non-actual verifiers otherwise.

So suppose first that $\Delta\top \in \Phi$. Then $\Delta B \in \Phi$ whenever B is a conjunction solely of $\top$ and $\Delta\top$. So A must include some non-logical literals, which must then be pure. But then using the ΔB_φ for all non-logical literals φ in A, using Δ-AbsP and the assumption $\Delta\top$, we can derive ΔA, contrary to the fact that Φ is consistent and $\neg\Delta A \in \Phi$.

Suppose now that $\Delta\top \notin \Phi$. Then $fm(A)$ is the sole verifier of A, and hence $@^p \subseteq fm(A)$, so some subset M of $m(A)$ is minimally complete. Then for some conjunction B with $m(B) = M$, $\Delta B \in \Phi$. By Δ-Equiv and Δ-AbsP, from ΔB and ΔB_φ for each non-logical $\varphi \in m(A)$ we can again derive ΔA, contrary to the consistency of Φ. So @ is not the sole verifier of ΔA and hence @ verifies $\neg\Delta A$, and so $\mathcal{M} \models \neg\Delta A$. □

We now construct a very simple, 'empty' model $\mathcal{E}$ for the case in which Φ does include $\Delta\neg\Delta\bot$. Let $S_e = \{0, 1\}$ and $S_e^\Diamond = \{0\}$, $@ = 0$, $\delta_e(0) = 0$ and $\delta_e(1) = 1$, and let $\sqsubseteq_e$ be the partial order on S_e in which $0 \sqsubseteq_e 1$. It is readily verified that $(S_e, S_e^\Diamond, \sqsubseteq_e, \delta_e)$ is a C-space. For $\varphi \in At$, we let $[\varphi]_e^+ = \{0\}$ if φ is pure, $\{1\}$ if $\varphi \notin \Phi$, and $\{0, 1\}$ otherwise, and $[\varphi]_e^- = \{0\}$ if $\neg\varphi$ is pure, $\{1\}$ if $\neg\varphi \notin \Phi$, and $\{0, 1\}$ otherwise. Again, exclusivity and exhaustivity are straightforward.

Lemma 12 *Let $\mathcal{E}$ be the canonical empty model, φ any literal, and A any 0-CF. If $\Delta\neg\Delta\bot \in \Phi$, then*

1. $\mathcal{E} \models \varphi$ *if* $\varphi \in \Phi$.
2. $\mathcal{E} \models \Delta A$ *if* $\Delta A \in \Phi$.
3. $\mathcal{E} \models \neg\Delta A$ *if* $\neg\Delta A \in \Phi$.

Proof (1): Let $\varphi \in \Phi$. If φ is a non-logical literal, then by definition of $[\cdot]_e$, φ is verified by $0 = @$, so $\mathcal{E} \models \varphi$. If φ is a logical literal then φ is one of $\top$, $\Delta\top$, and $\neg\Delta\bot$. It is readily verified that in either case, φ is again verified by 0, and indeed only by 0. So again, $\mathcal{E} \models \varphi$.

(2): Suppose $\Delta A \in \Phi$. Then every non-logical literal in A is pure, and hence by definition of $[\cdot]_e$ verified only by 0. Every logical literal in A is in Φ, and hence by (1) also verified only by 0. So ΔA, too, is verified by 0, and so $\mathcal{E} \models \Delta A$.

(3): Suppose $\neg\Delta A \in \Phi$. Then since $\Delta\top \in \Phi$, it follows that some literal in A is not in Φ or not pure. Either way, it is verified by 1, and hence so is A. Since 0 is an incompletion of 1, it follows that $\neg\Delta A$ is verified by 0, and hence $\mathcal{E} \models \neg\Delta A$. □

Lemma 13

1. *If* $\Delta\neg\Delta\bot \notin \Phi$, *then* $\mathcal{M} \models \varphi$ *for all* $\varphi \in \Phi$.
2. *If* $\Delta\neg\Delta\bot \in \Phi$, *then* $\mathcal{E} \models \varphi$ *for all* $\varphi \in \Phi$.

Proof Suppose $\varphi \in \Phi$. By Theorem 7, φ is provably equivalent to some Δ-DNF $\varphi' \in \Phi$. Since Φ is maximal consistent, some disjunct φ^* of φ' is also in Φ. Since φ' is a Δ-DNF, φ^* is a conjunction each conjunct of which is a literal or a Δ-literal and also a member of Φ. By Lemmas 15 and 16, every such conjunct, and hence their conjunction φ^*, and thus also φ', is true in the relevant model. By soundness, it follows that so is φ. □

From this lemma, completeness follows straightforwardly.

Theorem 8 *Let $\Phi \subseteq \mathcal{L}_\Delta$ and $\varphi \in \mathcal{L}_\Delta$. Then $\Phi \models_C \varphi$ implies $\Phi \vdash_C \varphi$.*

C.2 Disjunctive Δ

Let Φ be a subset of $\mathcal{L}_\Delta$ that is maximal consistent with respect to $\vdash_D$. Again, we begin by defining some important notions in terms of Φ.

Definition 17 Let $\varphi \in Lit$ and A a 0-CF in $\mathcal{L}_\Delta$. Then

- A is *complete* iff $\Delta A \in \Phi$
- A is *minimally complete* iff A is complete, and the result of replacing any conjunct of A by $\top$ is not complete
- $o(\varphi, A)$ is the number of times φ occurs as a conjunct in A
- $p(\varphi)$ is $\max\{o(\varphi, A) : A \text{ is minimally complete}\}$

With each (non-logical) literal $\varphi \in \mathcal{L}_\Delta$, we associate countably many *indexed literals* (short: i-literals) φ_i ($i \in \mathbb{N}$), and with each indexed literal φ_i we associate a unique *shadow* $\overline{\varphi_i}$. Roughly, the various φ_i will correspond to different ways for φ to be true, and $\overline{\varphi_i}$ acts as a negation of φ_i. We countenance an infinite number of ways for each literal to be true in order to prevent the sets of verifiers and falsifiers from being closed under fusion. The need for negations of (some of) these will become clear later on.

Definition 18 Let $\varphi \in Lit$ and $i \in \mathbb{N}$. Then

- $b\varphi = b(\varphi_i)$ is the sentence letter on which φ is based
- φ' is $\neg\varphi$ if $\varphi = b$ and $b\varphi$ otherwise
- $t\varphi = t(_i)$ is the unique member of $\{\varphi, \varphi'\} \cap \Phi$
- φ_i is *actual* iff $\varphi \in \Phi$ and either $i = 0$ or $i < p(\varphi)$
- We let $n\varphi$ and $n(\varphi_i)$ be: $(t\varphi)_0$ if $\Delta\top \in \Phi$; $(t\varphi)_0$ if $p(t\varphi) = 0$ and $p(\psi) > 0$ for some $\psi \in Lit$; $(t\varphi)'_0$ otherwise
- φ_i is negative iff $\varphi_i = n\varphi$ and positive otherwise

Intuitively, we may think of $t\varphi$ as the true, or obtaining state, and of $n\varphi$ as the negative state w.r.t. φ. We are taking $(t\varphi)_0$ to be negative iff the actual world is empty according to Φ, or some non-logical literals occur in minimally complete 0-CFs, but $t\varphi$ is not one of them. (Note that the actual world may be non-empty while no non-logical literals occur in minimally complete 0-CFs, because it would take an infinite 0-CF to exhaust the positive part of the world.) Then $n\varphi$ is non-actual if the truth w.r.t. φ is positive, and actual otherwise. We denote the set of actual indexed literals by @, and the set of actual and positive (short: *a-positive*) ones by $@^p$. Note that $@^p$ is empty iff $\Delta\top \in \Phi$.

Next, with any 0-CF $A \in \mathcal{L}_\Delta$ including only non-logical conjuncts, we associate a unique *matching* set of indexed literals $m(A) = \{\varphi_i : i < o(\varphi, A)\}$. Note that $m(A) = m(B)$ iff A and B correspond to the same multi-set of literals, i.e. A and B contain the same conjuncts, and they contain them the same number of times, though possibly in a different order. Moreover, say that a 0-CF B is *included* in a 0-CF A iff each literal occurring in B as a conjunct occurs at least as many times as a conjunct in A. Then $m(B) \subseteq m(A)$ implies that B is included in A.

We call a set s of i-literals *minimally complete* iff $s = m(A)$ for some minimally complete conjunction of non-logical literals A. Now suppose A is such a conjunction. Suppose $\varphi_i \in m(A)$, so φ occurs at least $i+1$ times in A. Then $i < p(\varphi)$, so $\varphi_i \in @^p$. So $m(A) \subseteq @^p$ whenever A is a minimally complete conjunction of non-logical conjuncts.

We construct the set of states similarly as before, but using as the set of proto-states P not the set of literals, but the set of the i-literals together with the shadows of the *a-positive* i-literals. Then for $s \subseteq P$, let $f(s) = s$ if no subset of is minimally complete, and $s \cup @^p$ otherwise, and let S be $\{f(s) : s \subseteq P\}$.

Definition 19 The *canonical state-space* for Φ is $(S, \sqsubseteq)$ with $\sqsubseteq = \subseteq \restriction S$

Since f is defined from P just as before, the Lemma 10 and its proof carry over unchanged.

Lemma 14 *Let* $s, t \subseteq P$.

1. $@ \in S$
2. $@^p \in S$
3. $f(s) \subseteq @^p$ *if* $s \subseteq @^p$
4. $f(s) \subseteq @$ *if* $s \subseteq @$
5. $S = \{s \subseteq P$*:* $@^p \subseteq s$ *or* s *has no minimally complete subset*$\}$
6. S *is closed under intersection*
7. *If* $s \neq t$ *are minimally complete and* s *has at least two members, then* S *is not closed under union*
8. $f(s) \subseteq f(t)$ *if* $s \subseteq f(t)$
9. *For any* $T \subseteq S$, $f \bigcup T$ *is* $\bigsqcup T$*, the least upper bound of* T *in* S
10. *If* $T \subseteq wp(P)$*, then* $\bigsqcup\{f(t) : t \in T\} = f \bigcup T$

Lemma 15 *The canonical state-space for* Φ*,* $(S, \sqsubseteq)$*, is a state-space.*

Proof From the fact that the subset-order on any set is a partial order together with Lemma 19 (9). □

To extend $(S, \sqsubseteq)$ to a canonical C-space, we need to define a set of *consistent* states and the δ-function, both of which is slightly more complicated than before.

Definition 20 Let $\varphi_i, \psi_j \in ILit$, and set $s \in S$. Then

- φ_i and ψ_j are *co-literals* iff $\varphi = \psi$, or $\varphi = \neg\psi$, or $\psi = \neg\varphi$
- φ_i and ψ_j are *incompatible* iff they are (a) co-literals and (b) not both actual
- s is *consistent* iff (a) no two members of s are incompatible i-literals, and (b) no member of s is the shadow of a member of s

- $S^{\Diamond} = \{s \in S : S \text{ is consistent}\}$

Next, if φ is a literal, say that a state s decides φ iff some co-literal of is a member of s. Then to form the completion of any given state, we extend it by the negative state w.r.t. any literal it does not decide, as well as the shadows of any a-positive i-literal that are not members of the state.

Definition 21 For $s \in S$, let

- $c(s) = \{n\varphi : s \text{ does not decide } \varphi\} \cup \{\overline{\varphi_i} : \varphi_i \in @^p \backslash s\}$
- $\delta s = s \cup c(s)$

We take note of a useful fact about the c-function used in defining δ:

Lemma 16 *For $s,\ t \in S$, if $s \sqsubseteq t$ then $c(t) \subseteq c(s)$*

Proof Since s decides no literal not decided by t, and contains no member of $@^p$ not contained in t, if $s \sqsubseteq t$. ☐

Lemma 17 *Let $s \in S^{\Diamond}$. Then δs is a world in $(S, S^{\Diamond}, \sqsubseteq, \delta)$.*

Proof δs is a state: Since $n\varphi$ is always a non-positive i-literal, and no shadow is a positive i-literal, δs has a minimally complete subset iff s does, and consequently δs is a state since s is.

δs is consistent: Firstly, no shadow of any member of δs is a member of s. For only a-positive i-literals have shadows, and the only a-positive i-literals in δs are already in s. Since s is consistent, it contains no shadows of its a-positive members, and beyond the states in s, by definition δs contains only shadows of i-literals not in s. Secondly, no pair of i-literals in δs is incompatible, for only co-literals are incompatible, and by construction, any pair of co-literals in δs is already in s, which was assumed to be consistent.

δs contains every state it is compatible with: Suppose that δs is compatible with t. First, suppose $\overline{\varphi_i}$ is a shadow in t. Since t is compatible with δs, $\varphi_i \notin \delta s$, and hence $\varphi_i \notin s$, so $\overline{\varphi_i} \in \delta s$.

Now suppose φ_i is an an i-literal in t. Suppose first that s does not decide φ. Then $n\varphi \in \delta s$. If $n\varphi \neq \varphi_i$, then $n\varphi$ and φ_i are distinct co-literals, and since $n\varphi$ is non-positive, they are incompatible, contrary to our assumption that δs is compatible with t. So $n\varphi = \varphi_i$ and hence $_i \in \delta s$. So suppose instead s does decide φ, so some member ψ_j of s is a co-literal of φ_i. If $\psi_j = \varphi_i$ then $\varphi_i \in \delta s$. If $\psi_j \neq \varphi_i$, then since δs and t are compatible, φ_i and ψ_k are both a-positive. But then since t is compatible with δs, $\overline{\varphi_i} \notin \delta s$, so by construction of δs we may infer that $\varphi_i \in s$ and hence $\varphi_i \in \delta s$. It follows that $t \sqsubseteq \delta s$, as desired. ☐

Lemma 23 *For all $s \in S$, $\delta s = @$ iff $@^p \sqsubseteq s \sqsubseteq @$*

Proof For the left-to-right direction, suppose $\delta s = @$. Then clearly $s \sqsubseteq @$, so it suffices to show that $@^p \sqsubseteq s$. So suppose otherwise. Then since s is a state, it follows that s has no minimally complete subset. So let φ_i be an a-positive literal which is not a member of s. Then by construction, $\overline{\varphi_i} \in \delta s$, contrary to the assumption that $\delta s = @$.

For the right-to-left direction, assume $@^p \sqsubseteq s \sqsubseteq @$. We show first that $\delta s \sqsubseteq @$. Since $@^p \sqsubseteq s$, $\delta s = s \cup \{n\varphi : s \text{ does not decide } \varphi\}$. So it suffices to show that $n\varphi \in @$ whenever s does not decide φ. But if s does not decide φ, then $(t\varphi)_0$ is not in s, and hence not in $@^p$, so is not a-positive, so $n\varphi = (t\varphi)_0$. But $t\varphi \in \Phi$, so $(t\varphi)_0 = n\varphi \in @$, as desired.

Finally, we show that $@ \sqsubseteq \delta s$. Suppose $\varphi_i \in @$. If $\varphi_i \in s$ then clearly $\varphi_i \in \delta s$. So suppose $\varphi_i \notin s$. It follows that φ_i is not a-positive. Since φ_i is actual, it follows that $i = 0$, and that $t\varphi = \varphi$, so $_i = (t\varphi)_0$. Since φ_i is not a-positive, $n\varphi = (t\varphi)_0$. So to show that $\varphi_i \in \delta s$, it suffices to show that s does not decide φ. Suppose otherwise, so s contains some co-literal ψ_j of φ_i. Then $\psi_j \neq \varphi_i$, since by assumption $_i \notin s$. Since $s \sqsubseteq @$, ψ_j is actual, so $\psi \in \Phi$. By consistency of Φ, $\psi = \varphi$, and so $j \neq 0$. But by the definition of actuality it then follows that $j < p(\varphi)$, and hence that $i < p(\varphi)$, contrary φ_i not being a-positive. So s does not decide φ, and hence $\varphi_i = n\varphi \in \delta s$. □

Lemma 24 $(S, S^\Diamond, \sqsubseteq, \delta)$ *is a C-space*

Proof From Lemmas 20 (9) and 22 it follows that $(S, S^\Diamond, \sqsubseteq)$ is a W-space,

δ-Containment: Immediate from the definitions of δs.

δ-Completeness: Immediate from the stronger Lemma 22.

δ-Redundancy(1): Suppose s is a world. $\delta s = s \cup \{n\varphi : s \text{ is } \varphi\text{-neutral}\} \cup \{\overline{\varphi_i} : \varphi_i \in @^p \backslash s\}$. Since s is consistent, δs is a world, and since s is a world and $s \sqsubseteq \delta s$, we have $s = \delta s$.

δ-Redundancy(2): It is easily verified that δs decides every literal and contains every positive i-literal contained in s. From that observation, it is immediate that $\delta\delta s = \delta s$.

δ-Identity(1): Suppose $s \sqsubseteq t \sqsubseteq \delta s$. We show first that $\delta t \sqsubseteq \delta s$. By assumption, $t \sqsubseteq \delta s$. By Lemma 21, $c(t) \subseteq c(s)$, so $c(t) \subseteq \delta s$, and hence $\delta t \sqsubseteq \delta s$. We now show that $\delta s \sqsubseteq \delta t$. By assumption, $s \sqsubseteq t \sqsubseteq \delta t$. So let $x \in c(s)$. Then either (a) $x = n\varphi$ with φ not decided by s, or (b) $x = \overline{\varphi_i}$ with φ_i an a-positive i-literal not in s. Suppose (a). Note that $n\varphi$ is the only co-literal of φ in δs. If φ is not decided by t, then $x = n\varphi$ is also in δt. If φ is decided by t, t contains a co-literal of φ. But $t \sqsubseteq \delta s$ and $n\varphi$ is the only co-literal of φ in δs, so $n\varphi \in t$ and hence $x = n\varphi \in \delta t$. Suppose (b). Then $\varphi_i \notin \delta s$, and by $t \sqsubseteq \delta s$ also $\varphi_i \notin t$, so again $x = \overline{\varphi_i} \in \delta t$.

δ-Identity(2): Suppose $\delta s = \delta t$. We show first that $\delta(s \sqcap t) \sqsubseteq \delta s$. Any member x of $\delta(s \sqcap t)$ is either (a) a member of $s \sqcap t$, or (b) $n\varphi$ with φ not decided by $s \sqcap t$, or (c) $\overline{\varphi_i}$ with φ_i an a-positive i-literal not in $s \sqcap t$. Suppose (a). Then $x \in s$ and hence $x \in \delta s$. Suppose (b). If either s or t also do not decide φ, then clearly $n\varphi \in \delta s = \delta t$. So suppose both s and t decide φ. Since $s \sqcap t$ does not, no co-literal of φ is a member of both s and t. Let ψ_j be a co-literal of φ that is a member of s but not t. Since $\delta s = \delta t$, $\psi_j \in \delta t$. Since ψ_j is neither in t nor a shadow, it follows that $\psi_j = n\varphi$, and hence that $n\varphi \in \delta s$. And if (c), then at least one of s and t also does not have φ_i as a member, so again $x \in \delta s = \delta t$.

We show finally that $\delta s \sqsubseteq \delta(s \sqcap t)$. By Lemma 21, since $s \sqcap t \sqsubseteq s$, $c(s) \subseteq c(s \sqcap t)$. It remains to show that $s \subseteq \delta(s \sqcap t)$. Since $\delta s = \delta t$, any member of s is a member of t or a member of $c(t)$. In the first case, it is a member of $s \sqcap t$ and hence of $\delta(s \sqcap t)$. In the second case, since $s \sqcap t \sqsubseteq t$, it is a member of $c(s \sqcap t)$ and so a member of $\delta(s \sqcap t)$. □

Again $@^p$ is the positive part of @ and the set of wholly positive states is closed under fusion under the definitions from Section 4.5; here, positivity is equivalent to not having any members which are either shadows or equal to $n\varphi$ for some φ.

We now define the interpretation function for the atoms and prove that it satisfies the conditions of exclusivity and exhaustivity.

Definition 22 For $\varphi \in At$, let $[\varphi]^+ = \{f\{\varphi_i\} : i \in \mathbb{N}\}$, and $[\varphi]^- = \{f\{\neg\varphi_i\} : i \in \mathbb{N}\}$

Lemma 25 *For all* $\varphi \in At$,

1. *s and t are incompatible whenever $s \in [\varphi]^+$ and $t \in [\varphi]^-$*
2. *for every world $w \in S$, there is some $s \in [\varphi]^+ \cup [\varphi]^-$ with $w \sqsupseteq s$*

Proof (1): Suppose $s \in [\varphi]^+$ and $t \in [\varphi]^-$, so for some i, j, $s = f\{\varphi_i\}$ and $t = f\{\neg\varphi_j\}$. By consistency of Φ, φ_i and $\neg\varphi_j$ are not both actual. Since they are co-literals, they are incompatible. Since $_i \in s$ and $\neg\varphi_j \in t$, s and t are incompatible.

(2): Let w be a world. By consistency of Φ, either φ_0 or $\neg\varphi_0$ is non-actual. Suppose φ_0 is non-actual. Then $f\{\varphi_0\} = \{\varphi_0\}$, so $\varphi_0 \notin w$. Since w is a world, w is incompatible with $\{\varphi_0\}$. Since φ_0 is non-actual, it does not have a shadow, so it follows that w has some co-literal of φ_0 as a member. Any co-literal of φ_0 is identical with either some φ_n or some $\neg\varphi_n$, so for some n, w contains either $f\{\varphi_n\}$ or $f\{\neg\varphi_n\}$. The case of $\neg\varphi_0$ non-actual is analogous. □

Definition 23 The *canonical model* $\mathcal{M}$ of Φ is $(S, S^\Diamond, \sqsubseteq, \delta, @, [\cdot])$

Lemma 26 *$\mathcal{M}$ is a model.*

Proof From Lemmas 24 and 25. □

In the canonical model, the literals and Δ-literals in Φ are verified by parts of @.

Lemma 27 *In the canonical model $\mathcal{M}$, for φ any literal and A any 0-CF:*

1. *$\mathcal{M} \models \varphi$ if $\varphi \in \Phi$*
2. *If A contains only non-logical conjuncts, then $fm(A)$ verifies A*

3. $\mathcal{M} \models \Delta A$ *if* $\Delta A \in \Phi$
4. $\mathcal{M} \models \Delta A$ *if* $\neg\Delta A \in \Phi$

Proof (1): Let $\varphi \in \Phi$ and suppose first that φ is a logical literal. If φ is $\top$ or $\neg\Delta\bot$, then it is verified by $\Box$ and hence by part of @. Since Φ is consistent, φ cannot be $\bot$ or $\Delta\bot$. If φ is $\Delta\top$, then by definition, $@^p$ is empty, and hence $@ = \delta\Box$ verifies φ. If φ is $\neg\Delta\top$, then by definition, $@^p$ is not empty, so @ is an incompletion of $\Box$ and hence verifies $\neg\Delta\top$. Suppose next that φ is a non-logical literal. Then φ_0 is an actual i-literal, so $\{\varphi_0\} \subseteq @$. By Lemma 19(4), $f\{\varphi_0\} \sqsubseteq @$. By definition of $[\cdot]$, $f\{\varphi_0\}$ verifies φ. So $\mathcal{M} \models \varphi$.

(2): We may set up a one-one correspondence between occurrences of literals in A and members of $m(A)$ by pairing each $\varphi_i \in m(A)$ with the $(i+1)^{\text{th}}$ occurrence of φ in A. Since always $f\{\varphi_i\} \in [\varphi]^+$, by the clause for conjunction it follows that $\bigsqcup\{f(\varphi_i) : \varphi_i \in m(A)\} \in [A]^+$. By Lemma 19(10), $\bigsqcup\{f(\varphi_i) : \varphi_i \in m(A)\} = f\bigcup\{\varphi_i : \varphi_i \in m(A)\} = fm(A)$.

(3): Suppose $\Delta A \in \Phi$. By Δ-Fact, every conjunct of A is in Φ, and so by (1) is verified by some part of @, so some verifier s of A is part of @. Suppose first that $\Delta\top \in \Phi$, so $@^p = \Box$. Then s is between $@^p$ and @, and so $\delta s = @$, so @ verifies ΔA. Suppose now that $\Delta\top \notin \Phi$. There are three non-logical literals that can occur as conjuncts in A, namely $\top$, $\neg\Delta\top$, and $\neg\Delta\bot$. The latter two are verified by @, so if either is a conjunct of A, then @ verifies A and hence ΔA. If at most $\top$ is a logical conjunct of A, then the result of removing $\top$ from A as a conjunct is still complete and extends some minimally complete conjunction A' of only non-logical literals. By (2), $fm(A')$ verifies A', and since A' is minimally complete, $fm(A') = @^p$. It follows that some verifier s of A is between $@^p$ and @ and hence that $\delta s = @$ verifies ΔA, so $\mathcal{M} \models \Delta A$.

(4): Suppose $\neg\Delta A \in \Phi$. Given the clause for the falsifiers of ΔA, it suffices to show that @ is not a verifier of ΔA, and hence that no state between $@^p$ and @ verifies A. So suppose for contradiction that s verifies A and $@^p \sqsubseteq s \sqsubseteq @$. Then $\mathcal{M} \models \varphi$ for every conjunct φ of A, and so by (1) and Φ being maximal consistent, $A \in \Phi$. We show that then $\Delta A \in \Phi$, contrary to the consistency of Φ. Note that by Δ-AbsT, if suffices to show that $\Delta B \in \Phi$ for some 0-CF B included in A. If $\Delta\top \in \Phi$, then $\Delta A \wedge \top \in \Phi$ and hence $\Delta A \in \Phi$. So suppose $\Delta\top \notin \Phi$. If $\neg\Delta\bot$ is a conjunct in A, then since $\vdash_C \Delta\neg\Delta\bot$, we obtain $\Delta A \in \Phi$. If $\neg\Delta\top$ is a conjunct in A, then by Δ-NFix, $\Delta\neg\Delta\top \in \Phi$ and hence again $\Delta A \in \Phi$. In every other case, there can be no logical conjuncts in A except possibly $\top$. It follows that the conjunction A' of the non-logical conjuncts in A is also verified by s. So suppose let $A' = \varphi^0 \wedge \ldots \varphi^n$ and let $s = s_1 \sqcup \ldots \sqcup s_n$ with every s_i verifying φ^i. Note that for all i, there is a $j(i)$ such that $s_i = f\{\varphi^i_{j(i)}\}$, and by Lemma 19(10), $s = f\{\varphi^i_{j(i))} : 0 \leq i \leq n\}$. Since by assumption $s \sqsupseteq @^p$, by definition of f it follows that some subset x of $\{\varphi^i_{j(i))} : 0 \leq i \leq n\}$ is minimally complete, so $x = m(B)$ for some minimally complete 0-CF B. It follows that $\Delta B \in \Phi$, and that B is included in A, so again $\Delta A \in \Phi$. $\Box$

Lemma 28 $\mathcal{M} \models \varphi$ *for all* $\varphi \in \Phi$.

Proof As before, now using Theorem 9 and Lemma 27. □

From this lemma, completeness follows straightforwardly.

Theorem 29 *Let* $\Phi \subseteq \mathcal{L}_\Delta$ *and* $\varphi \in \mathcal{L}_\Delta$*. Then* $\Phi \models_D \varphi$ *implies* $\Phi \vdash_D \varphi$*.*

Acknowledgements For helpful comments on earlier drafts of this paper and/or discussion of the ideas here developed, I would like to thank the Relevance crew: Stefan Roski, Martin Glazier, and Singa Behrens, as well as Stephan Leuenberger, Bruno Whittle, Josh Schechter, Johannes Korbmacher, and especially Kit Fine. I am also grateful to two anonymous reviewers for this journal whose comments led to highly significant extensions of and improvements to the paper. My work on this paper was done as part of the Emmy Noether project *Relevance*, funded by the *Deutsche Forschungsgemeinschaft* (Grant number KR 4516/2-1); I am very grateful for that support.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

1. Chalmers, D. J., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *The Philosophical Review*, *110*, 315–361. https://doi.org/10.2307/2693648.
2. Ciardelli, I., Zhang, L., & Champollion, L. (2018). Two switches in the theory of counterfactuals. *Linguistics and Philosophy*, *41*, 577–621. https://doi.org/10.1007/s10988-018-9232-4.
3. Fine, K. (2016). Angellic Content. *Journal of Philosophical Logic*, *45*(2), 199–226. https://doi.org/10.1007/s10992-015-9371-9.
4. Fine, K. (2017). A theory of truthmaker content I: conjunction, disjunction and negation. *Journal of Philosophical Logic*, *46*, 625–674. https://doi.org/10.1007/s10992-016-9413-y.
5. Güngör, H. (2022). Counterfactuals, hyperintensionality, and Hurford disjunctions. Linguistics and Philosophy, https://doi.org/10.1007/s10988-022-09360-3.
6. Krämer, S. (202Xa). The whole truth. Forthcoming. In F. Leuenberger, & F. van de Putte (Eds.) *Kit Fine on Truthmakers, Relevance, and Non-Classical Logic. Springer*.
7. Krämer, S. (202Xb). Truthmaker equivalence. Forthcoming. In S. Leuenberger, & A. Rieger (Eds.) *Themes from Alan Weir*. Springer Library.
8. Leuenberger, S. (2014). Total logic. *The Review of Symbolic Logic*, *7*(3), 529–547. https://doi.org/10.1017/S1755020314000124.
9. Spector, B. (2016). Comparing exhaustivity operators. *Semantics & Pragmatics*, *9*, 1–33. https://doi.org/10.3765/sp.911.

Synthese (2017) 194:2917–2930
DOI 10.1007/s11229-016-1078-0

A hyperintensional criterion of irrelevance

Stephan Krämer[1]

Received: 8 February 2016 / Accepted: 22 March 2016 / Published online: 29 March 2016

Abstract On one important notion of irrelevance, evidence that is irrelevant in an inquiry may rationally be discarded, and attempts to obtain evidence amount to a waste of resources if they are directed at irrelevant evidence. The familiar Bayesian criterion of irrelevance, whatever its merits, is not adequate with respect to this notion. I show that a modification of the criterion due to Ken Gemes, though a significant improvement, still has highly implausible consequences. To make progress, I argue, we need to adopt a hyperintensional conception of content. I go on to formulate a better, hyperintensional criterion of irrelevance, drawing heavily on the framework of the truthmaker conception of propositions as recently developed by Kit Fine.

Keywords Relevance · Hyperintensionality · Partial content · Truthmaker semantics

1 Introduction

In the context of any inquiry, we need to distinguish between evidence that is relevant to the problem at hand, and evidence that is irrelevant. The distinction is of some importance. Discarding relevant evidence as irrelevant increases the likelihood of error and bad decision-making, and treating irrelevant evidence as relevant results at best in a waste of resources. In view of its centrality to our cognitive and practical

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

1 Department of Philosophy, University of Hamburg, Von-Melle-Park 6, 20146 Hamburg, Germany

lives, the notion of relevance that is in play here carries great philosophical interest, and an adequate explication of the notion would be highly desirable.[1]

At first glance, it may appear that such an explication is afforded by the understanding of relevance in terms of a change in probabilities. Roughly speaking, a piece of evidence is relevant to a given hypothesis just in case it makes the hypothesis either more or less probable than it would otherwise be. This suggestion is given a precise formulation in the shape of the usual Bayesian standard for irrelevance[2]:

(BI) A piece of evidence E is irrelevant to a hypothesis H iff $Pr(H|E) = Pr(H)$.

Unfortunately, as Gemes (2007) has convincingly argued, this explication of our notion is not satisfactory.[3] In a rough approximation, the problem is that evidence may bear on a hypothesis both in a positive and a negative way, so that its probabilistic effects cancel each other out.[4] In such a case, (BI) issues a verdict of irrelevance. But in real situations of inquiry and decision-making, Gemes points out, we would not, and we should not, discard such pieces of evidence as irrelevant (cf. Gemes 2007, p. 162f).

Gemes then goes on to propose a new account of irrelevance which avoids the problematic feature of (BI). The account has two main components. The first is a modified criterion of Gemesian irrelevance that replaces (BI)[5]:

(GI) A piece of evidence E is irrelevant to a hypothesis H iff for every part E' of E and every part H' of H, $Pr(H'|E') = Pr(H')$.

In contrast to (BI), (GI) invokes a notion of a *part* of a content or proposition, such as a piece of evidence or a hypothesis.[6] Although we have some intuitive grip on that notion, it clearly stands in need of explication. The second component of Gemes'

[1] As Cohen has emphasized (cf. 1994, p. 171f), there are also important applications of the notion of relevance which cannot readily be represented as concerning a relation of evidence to hypothesis. In this paper, I restrict attention to applications which can naturally be so understood. As I mention below, however, I suspect that the tools I shall employ in accounting for these instances of relevance can also fruitfully be applied in a much greater range of cases.

[2] $Pr(H)$ denotes the (prior) probability of the hypothesis H, and $Pr(H|E)$ denotes the (posterior) probability of H given the evidence E. The standard provided by (BI), and any of its refinements to be considered below, is thus relative to a suitable prior probability distribution. Following common practice in the debate, we won't worry here about the exact nature of the probabilities in question. The notion of irrelevance is often further relativized to a body of background information K and is then taken to be characterized by the condition that $Pr(H|E \wedge K) = Pr(H|K)$. I have left out reference to K throughout. Doing so facilitates comparison with Gemes' (2007) account to be considered below, and the difference is of no import for our purposes.

[3] This is not to say that (BI) may not amount to an adequate explication of *a* useful notion of irrelevance. The claim is that the notion sketched above is not adequately explicated by (BI).

[4] Cf. (Gemes 2007, p. 162). As Gemes highlights, this kind of complaint against (BI)'s identification of relevance with probabilistic relevance is not new; it was already made in 1929 by John Maynard Keynes (1929, p. 79).

[5] See (Gemes 2007, p. 165). Gemes' formulation has 'content part' in place of 'part', which is Gemes' term for his explication of the notion of a part of a content. For our purposes, it is better to view as separate the specific account of that notion and the proposed revision of (BI). The term 'part', here and throughout, is to be understood as 'proper or improper part'.

[6] I use the terms 'content' and 'proposition' interchangeably. I have found it convenient to speak of evidence and hypotheses as themselves propositions rather than sentences expressing these propositions.

proposal is accordingly a precise account of this notion, which Gemes has developed and defended in more detail in earlier papers (Gemes 1994, 1997).

I have two primary aims in this paper, one destructive, the other constructive. The destructive aim is to reveal some problems both for the criterion (GI) and for Gemes' specific account of content parts. The constructive aim is to formulate a better account of relevance. This may be done, I argue, by firstly, replacing the Gemesian account of content parts with the rival account recently developed by Kit Fine (2015), and secondly, by tweaking (GI) somewhat. Both Fine's account of partial content[7] and the tweaks to (GI) that I propose draw heavily on the hyperintensional distinctions made available in Fine's *truthmaker* account of content. This leads to a secondary aim for the paper. For, the great utility of the truthmaker framework in accounting for the present notion of evidential relevance may plausibly be taken as an indication of a deeper and more general connection between notions of relevance on the one hand, and the central concepts of the truthmaker account on the other. The results of the paper are therefore suggestive of the possibility of a more encompassing and general theory of relevance within the framework of truthmaker semantics. I hope to further pursue this matter in future research.

The plan for the paper is as follows. In Sect. 2, I present an example in which Gemes' account yields a counter-intuitive verdict of irrelevance. In this case, I argue, the problem lies with Gemes' account of content parts. On an intuitive understanding of that notion, the cases are consistent with (GI). Section 3 argues that the difficulties arising for the Gemesian view may be traced to his insistence that logically equivalent sentences must have the same content. I show that we obtain a much more plausible version of (GI) once we replace Gemes' intensional account of content parts with its hyperintensional Finean rival, which allows for logically equivalent sentences to have different contents, and hence content parts. Section 4 argues that intuitively, even under the Finean interpretation of 'part', (GI) is subject to counter-examples. To avoid them, two changes are proposed. In a first step, we broaden our attention to consider not just parts of the hypothesis, but parts of any disjunct of the hypothesis. In a second step, we restrict attention to what I call *helpful* parts. Roughly, these are parts whose truth brings us closer to the truth of the hypothesis than we are independently of them. Section 5 concludes.

2 A counter-example

I shall begin by briefly reviewing the case that Gemes uses to argue against (BI) and to motivate his own account (cf. Gemes 2007, p. 162f). Suppose two dice A and B have been tossed, and consider the following pair of evidence and hypothesis:

(E_1) A came up 1, and B came up either 1, 3, 5, or 6.
(H_1) A came up odd, and B came up even.

[7] I use the phrases 'content part' and 'partial content' interchangeably. The former is favoured by Gemes, the latter by Fine.

In this case (given natural background assumptions), $\Pr(H_1|E_1) = \Pr(H_1)$.[8] According to (BI), therefore, the evidence E_1 is irrelevant to H_1. But this seems to be the wrong result. It would appear irrational, in the context of an inquiry into H_1, to discard the information that E_1 holds as irrelevant. Gemes' example therefore constitutes a strong case against (BI) as an explication of the target notion of irrelevance.

The case also lends some support to Gemes' own account, which implies that E_1 is relevant to H_1. Clearly, the proposition

(E'_1) A came up 1.

is probabilistically relevant to H_1, for $\Pr(H_1|E'_1) = 1/2 \neq 1/4 = \Pr(H_1)$. Intuitively, E'_1 would certainly seem to qualify as a part of E_1, in which case (GI) implies that E_1 is relevant to H_1. Similarly, E_1 itself is probabilistically relevant to the hypothesis

(H'_1) A came up odd.

For $\Pr(H'_1|E_1) = 1 \neq 1/2 = \Pr(H'_1)$. Intuitively, H'_1 certainly qualifies as a part of H_1, which again would mean that (GI) counts E_1 relevant to H_1. As Gemes shows, the relevant claims of parthood are vindicated on his account of content parts. So it seems that Gemes can handle the example quite convincingly.[9]

Unfortunately, there are many other cases which intuitively, as far as relevance is concerned, are of exactly the same sort as the previous example, but nevertheless receive the opposite verdict on Gemes' account. Here is one such case. Suppose two fair coins A and B have been tossed, and consider the following pair of evidence and hypothesis:

(E_2) A came up heads.
(H_2) (A came up heads or B came up tails) and (A came up tails or B came up heads).

It is easy to see that $\Pr(H_2) = \Pr(H_2|E_2) = 1/2$. The outcomes that render H_2 true are exactly the both-heads and the both-tails outcomes, and the probability that one of these two obtains is 1/2 independently of the evidence that A came up heads as well as given that evidence. So according to (BI), E_2 is irrelevant to H_2. But just as in Gemes' case, this result is implausible. It would appear irrational, in the context of an inquiry into H_2, to discard the information that E_2 holds as irrelevant. Moreover, on an intuitive understanding of content parts, the case also seems to fit with (GI). Clearly, E_2 is probabilistically relevant to the hypothesis

(H'_2) A came up heads or B came up tails.

[8] The obvious prior probability of H_1 is 1/4, which is also the probability of H_1 given E_1, since E_1 is compatible with four equally probable outcomes of the tosses, exactly one of which makes H_1 true.

[9] Cf. (Gemes 2007, p. 165f). Note that to deal with this particular example, a more modest deviation from (BI) than is embodied in (GI) would have been sufficient. For as we saw, (GI) overdetermines, as it were, the result that E_1 is relevant to H_1 in that we have both that part of the evidence is probabilistically relevant to the hypothesis as a whole, and that the evidence as a whole is probabilistically relevant to part of the hypothesis. But this feature is specific to the current example. The next example in the main text plausibly can only be captured by allowing as sufficient for relevance the probabilistic relevance of the evidence for part of the hypothesis. For an example of the 'converse' sort, consider any hypothesis H and the corresponding evidence that Bill said H was false and Bob said H was true, where Bill and Bob are generally reliable sources equally likely to be wrong (or lying) with respect to H or any part of H.

For $\Pr(H_2'|E_2) = 1 \neq 3/4 = \Pr(H_2')$. Intuitively, H_2' certainly seems to qualify as part of H_2, in which case (GI) yields the desirable result that E_2 is relevant to H_2.

However, on Gemes' account, neither H_2' nor H_2's other conjunct that A came up tails or B came up heads is a content part of H_2. On this account, for the content of a sentence α to be part of the content of a sentence β is for α to be a special kind of logical consequence of β. Gemes offers two characterizations of the additional condition that has to be satisfied, one relatively informal, the other more formal. Suppose that α is a logical consequence of β. Then on the informal version, in order for α's content to be a Gemesian part of the content of β, it must be the case that: there is no logical consequence γ of β such that γ logically entails, but is not logically entailed by, α, and 'all the vocabulary of [γ] occurs (essentially) in α' (Gemes 2007, p. 164).[10]

More formally characterized, the condition a logical consequence α of β must satisfy in order for α's content to qualify as as Gemesian part of the content of β is that every relevant model of α can be extended to a relevant model of β:

(GP) The content of α is part of the content of β iff $\beta \vdash \alpha$, and every relevant model of α can be extended to a relevant model of β.

Here, a relevant model is one that assigns truth-values only to the relevant sentence letters in a formula, and a sentence letter is relevant in a formula iff changing its truth-value within a given model can change the truth-value of the formula (cf. ibid). Thus, 'P' is relevant in '$P \vee Q$' and in '$\neg P \wedge Q$', but irrelevant in both '$P \vee \neg P$' and '$Q \vee (P \wedge Q)$'.

To apply the account to our example, we first formalize evidence and hypothesis within a propositional language. Let 'P' and 'Q' stand for the propositions that coin A came up heads and that coin B came up heads, respectively. Our hypothesis H_2 may then be written as '$(P \vee \neg Q) \wedge (\neg P \vee Q)$'. We now see that '$P \vee \neg Q$' has a relevant model—the model assigning True to 'P' and False to 'Q'—which cannot be extended to a relevant model of '$(P \vee \neg Q) \wedge (\neg P \vee Q)$'. And likewise '$\neg P \vee Q$' has a relevant model—the model assigning False to 'P' and True to 'Q'—which also cannot be extended to a relevant model of '$(P \vee \neg Q) \wedge (\neg P \vee Q)$'. Nor are there other parts of the content of '$(P \vee \neg Q) \wedge (\neg P \vee Q)$' to which the evidence P might be probabilistically relevant. For in order for the content of α to be part of the content of '$(P \vee \neg Q) \wedge (\neg P \vee Q)$', α must not contain any sentence letters except for 'P' and 'Q'. If α contains only 'P' or only 'Q', then in order for it to be a logical consequence of '$(P \vee \neg Q) \wedge (\neg P \vee Q)$', it must be logically true, so the evidence cannot be probabilistically relevant to it. If α contains both 'P' and 'Q', then α must be

[10] The quoted phrase, and in particular the parenthetical qualification, is the bit which requires further clarification. Firstly, the vocabulary of γ in question is supposed to be non-logical vocabulary, otherwise the fact that $P \wedge Q$ is a stronger consequence of $P \wedge Q$ than $P \vee Q$ would not prevent the latter from being a part of $P \wedge Q$, as Gemes clearly intends it to do (cf. Gemes 1994, p. 603). Secondly, from the official definition offered in (Gemes 1994, p. 605), we may extract that a piece of non-logical vocabulary is said to occur essentially in a sentence just in case there is no logically equivalent sentence in which it does not occur. The motivation for the restriction to α's essential vocabulary is to ensure that logically equivalent sentences stand in the same parthood relations (cf. Gemes 1994, p. 604f). To see the point, note that without the restriction, P is part of $P \wedge Q$, but the logically equivalent $P \wedge (Q \vee \neg Q)$ is not, since $P \wedge Q$ itself is a stronger logical consequence of $P \wedge Q$ which contains all the vocabulary in $P \wedge (Q \vee \neg Q)$.

logically equivalent to '$(P \vee \neg Q) \wedge (\neg P \vee Q)$', otherwise it will have a relevant model that is not, and cannot be extended to, a relevant model for '$(P \vee \neg Q) \wedge (\neg P \vee Q)$'.[11]

3 Partial content

At first glance, it seems an odd feature of (GP) that the conjuncts of a conjunction may fail to be parts of the conjunction. The relation between a conjunct P and a corresponding conjunction $P \wedge Q$ seems to be the very paradigm of the relation of content part; it is no accident that the content parts of evidence and hypothesis figuring in the example by which Gemes seeks to motivate (GI) are conjuncts of evidence and hypothesis, respectively. So one may wonder why Gemes chooses to restrict the relation of content part in a way that rules out some instances of the conjunct-conjunction relation. The reason is that he is forced to do so given two other principles he wishes to uphold. I explain them in turn.

As Gemes highlights, the most important feature, for his purposes, of the notion of content part is that not every logical consequence of P, and in particular not every disjunction $P \vee Q$, counts as a part of P. For if it did, (GI) would yield an almost empty notion of irrelevance. The reason is that in almost all cases of a piece of evidence E and a hypothesis H, the disjunction E $\vee$ H is probabilistically relevant to H. Hence if E $\vee$ H is a part of E, by (GI), E automatically comes out relevant to H. It is intuitively quite plausible to deny that $P \vee Q$ is always a part of P. To vary a point made by Gemes, we should otherwise have to say for arbitrary P and Q that the evidence Q conclusively confirms part of the hypothesis P, which seems counter-intuitive. Similarly, we are in no way tempted to count P as partially true, purely on the strength of the truth of Q, and hence $P \vee Q$. (Cf. Gemes 2007, p. 164; Gemes 1994, p. 597ff; see also Fine 2013.)

The second principle is a principle of *intensionality*. It says that pairs of logically equivalent sentences always have the same content, and thus the same content parts.[12] Given this principle, the claim that every conjunction $P \wedge Q$ contains its conjuncts as parts implies that an arbitrary logical consequence Q of a given proposition P is always also a part of P, in contradiction of the first principle. For suppose Q is a logical consequence of P. Then P is logically equivalent to $P \wedge Q$. But by the conjunction principle, Q is part of $P \wedge Q$, and hence of P.

Given these two principles, then, Gemes has no choice but to restrict the principle that conjuncts of a conjunction are parts of the conjunction. As we have seen, the restriction he chooses yields some counter-intuitive denials of parthood, and in

[11] We can construct a similar case pertaining to parts of the evidence. Consider some hypothesis H with prior probability 1/2 and two generally reliable sources Bill and Bob who are equally likely to be wrong (or lying) about H, and let the evidence be that (Bob or Bill said H is true) and (Bob or Bill said H is false). Let it be given that both Bob and Bill said either that H is true or that H is false, so that the evidence may be represented as $(P \vee Q) \wedge (\neg P \vee \neg Q)$. The evidence as a whole is intuitively, but not probabilistically, relevant for the hypothesis, whereas the conjuncts are probabilistically relevant. For the same reasons as before, however, the conjuncts do not qualify as parts of the evidence on Gemes' account.

[12] This principle is not discussed in Gemes (2007), but it plays a central role in Gemes' development of his view in (Gemes 1994, cf. esp. pp. 601–605).

conjunction with (GI), some counter-intuitive denials of relevance. We should note, moreover, that the denials of parthood cannot be motivated by the sorts of considerations that Gemes uses to rule out arbitrary disjunctions as parts. For intuitively, there is no problem with saying that $(P \vee \neg Q) \wedge (\neg P \vee Q)$ is partially true, or that a part of it has been conclusively confirmed, given that $P \vee \neg Q$ is true.

Indeed, these kinds of intuitive considerations seem to speak heavily in favour of accepting the conjunction principle and instead giving up on the intensionality principle. For intuitively, there is also no problem with saying that $P \wedge (P \vee Q)$ is partially true, or that a part of it has been conclusively confirmed, given that Q is true.[13] At the same time, there is no temptation at all to infer from this that the logically equivalent statement P is partially true, or that part of P has been conclusively confirmed.

All this suggests that a hyperintensional conception of content may be better suited for the explication of the notion of content part, and thus for giving a plausible interpretation to the irrelevance criterion (GI). To confirm this conjecture, I will now briefly sketch the *truthmaker* conception of content recently developed by Kit Fine and explain how Fine proposes to explicate the notion of parthood within his framework.[14] I will then apply Fine's view to our case and show that it gives the desired results.

The best way to introduce Fine's truthmaker conception of content is to contrast it with the familiar possible worlds conception of content. On the latter view, a proposition is identified with the set of possible worlds at which the proposition is true. On the truthmaker view, a proposition is instead described in terms of the set of *states* which make the proposition true.[15] Roughly speaking, states are like possible worlds except in that they need not be possible, and they need not be (complete) worlds. That is, whereas a world is complete in the sense that it settles the truth-value of every proposition, a state may be incomplete and leave open the truth-values of many propositions.

States are assumed to have mereological structure. In particular, given any states $s, t, u, \ldots$ we may form their *fusion* $s \sqcup t \sqcup u \sqcup \ldots$ which is the smallest state containing

[13] An anonymous referee has suggested to me that it might after all appear counter-intuitive that one can turn anything into a partial truth just by tacking on a logically idle conjunct. I concede that our unreflective judgement with respect to this principle—especially in this particular, somewhat leading wording—may be negative or at least sceptical. However, part of the intuitive resistance to the principle appears to vanish already when it is reformulated in a more explicit and neutral fashion. The point is that for any proposition P, there is some logical consequence R of P such that $P \wedge R$ is (at least) partially true. More importantly, it seems to me that any remaining intuitive uneasiness with respect to the principle disappears once one reflects that: (i) it seems intuitively very plausible to say that anything can be turned into a partial truth by tacking on a *true* conjunct; (ii) anything has some true logical consequences, and (iii) even given this principle, if one wants to turn a proposition one does not know to be true into a proposition one knows to be partially true, then one needs to 'invest' some known truth. That is, one needs to add a conjunct which one independently knows to be true.

[14] The fullest published exposition by Fine of the truthmaker conception of content and the notion of partial content is given in Fine (2015) in the context of a discussion of Angell's logic of analytic entailment. A more general presentation and discussion of the framework is contained in the as yet unpublished manuscripts Fine (msa) and Fine (msb).

[15] While this may make it sound as though the views make incompatible claims about the same kind of thing, viz. propositions, it is not necessary for my purposes that we think of the views in this way. We may instead take them to concern different concepts of propositions, suited to different theoretical purposes.

all of $s, t, u, \ldots$ as parts. We shall take for granted two important principles about propositions and the states that make them true, or *verify* them:

(Closure) If a proposition P is verified by each of some states $s, t, u, \ldots$ then it is also verified by their fusion $s \sqcup t \sqcup u \sqcup \ldots$

(Convexity) If a proposition P is verified by each of the states s and u, then P is also verified by any state t which is both part of u and has s as a part.

It is important to note, however, that verification is *not* assumed to be *monotonic* in the sense that if a proposition is verified by a state s, then it is also verified by any bigger state s' of which s is a part. Roughly speaking, on Fine's construal of verification, in order for a state to verify a proposition, every part of the state must play a part, must be actively involved, as it were, in the verification of the proposition.[16]

On the possible worlds picture, by fixing which worlds make a proposition P true, we ipso facto also fix which worlds make P false, namely all worlds which do not make P true. We are thereby in a position to say which worlds make the negation $\neg P$ true, namely those worlds which make P false. On the truthmaker picture, this is not so. A state which does not verify a proposition P need not therefore falsify it, and correspondingly, it need not verify $\neg P$. As a result, we have to separately specify the verifiers and the falsifiers of a given proposition.[17] A proposition P is therefore identified with an ordered pair $\langle P^+, P^- \rangle$ of a non-empty set P^+ of states verifying P and a non-empty set P^- of states falsifying P. Operations of conjunction, disjunction and negation on the propositions are then defined as follows (for X a set of states, we write X° for the smallest closed and convex set containing X):

$$
\begin{aligned}
(\neg P)^+ &= P^- \\
(\neg P)^- &= P^+ \\
(P \wedge Q)^+ &= \{s \sqcup t : s \in P^+ \text{ and } t \in Q^+\}^\circ \\
(P \wedge Q)^- &= (P^- \cup Q^-)^\circ \\
(P \vee Q)^+ &= (P^+ \cup Q^+)^\circ \\
(P \vee Q)^- &= \{s \sqcup t : s \in P^- \text{ and } t \in Q^-\}^\circ
\end{aligned}
$$

Fine now proposes the following account of what it is for a proposition P to be part of a proposition Q (cf. Fine 2015, p. 8ff, 19):

(FP) P is a part of Q iff (i) every verifier of P is part of a verifier of Q

[16] Strictly speaking, Fine distinguishes a number of different conceptions of verification. I am here concerned with what Fine calls exact verification, which is the basic notion of verification in terms of which he defines other, looser conceptions. Cf. (Fine 2015, pp. 7f, 20f), and (Fine, msa, p. 35f). Fine sometimes describes his notion of exact verification as embodying a constraint of holistic relevance in the sense that for a state to exactly verify a proposition, it must be wholly relevant to the proposition, and so must not have any part that is irrelevant to the proposition (cf. e.g. (Fine, msa, p. 1)). This may invite the worry that some sort of untoward circularity is involved in using Fine's framework to describe and study relations of relevance. However, the appeal to a notion of relevance is confined solely to Fine's informal commentary on his theory, and not part of the theory itself. It might perhaps still be claimed that to the extent that exact verification imposes relevance constraints, an analysis of relevance within the truthmaker framework is in that sense not fully reductive; I would be content to concede that much.

[17] This is a slight exaggeration, since other treatments of negation are possible within truthmaker semantics that do not require a separate specification of falsifiers, appealing instead to modal connections on the states. These approaches to negation will not be considered here.

(ii) every verifier of Q has a part that verifies P
(iii) every falsifier of P is a falsifier of Q

Like (GP), (FP) has the desirable consequence that $P \vee Q$ is not in general a part of P, since a verifier of Q will not in general be part of a verifier of P. (Note that for this result it is important that verification is non-monotonic.) Unlike (GP), (FP) also has the consequence that for any propositions P, Q, both P and Q are parts of $P \wedge Q$. Indeed, we might say that on (FP), being a part of a proposition is the same as being a conjunct of the proposition, for the condition that P is a part of Q turns out to be equivalent to the condition that $P \wedge Q = Q$ (cf. Fine 2015, p. 13f, 19). In particular, then, on (FP), the conjuncts $P \vee \neg Q$ and $\neg P \vee Q$ of the hypothesis $(P \vee \neg Q) \wedge (\neg P \vee Q)$ in our example are classified as parts of that hypothesis. As a result, by interpreting the notion of part invoked in Gemes' criterion (GI) in terms of Fine's hyperintensional explication, we avoid the counter-example of the previous section.

4 Fine-tuning

I turn now to some difficulties for (GI) that do not arise from an inadequate conception of partial content and show how the criterion may be refined to avoid them.

4.1 Parts of disjuncts

In this section I argue that even under the Finean conception of parts of contents, (GI) overgenerates irrelevance, and propose a fix that employs a disjunctive counterpart of the notion of content parts. We begin by considering a slight variation on the case discussed in Sect. 2, namely the following pair of evidence and hypothesis:

(E_3) A came up heads.
(H_3) (A came up heads and B came up heads) or (A came up tails and B came up tails).

Note that the hypothesis H_3, under the obvious formalization $(P \wedge Q) \vee (\neg P \wedge \neg Q)$, is logically equivalent to the previous hypothesis H_2, formalized as $(P \vee \neg Q) \wedge (\neg P \vee Q)$. As before, $E_3 = P$ is intuitively relevant to H_3, even though it does not lower or raise its probability. If we ask why E_3 appears relevant to H_3, the most natural answer goes roughly along the following lines: Firstly, E_3 guarantees the truth of one conjunct of the first disjunct of H_3, and thereby makes it more probable that this disjunct obtains. Secondly, E_3 rules out the truth of the first conjunct of the second disjunct of E_3, and thereby ensures that this disjunct does not obtain.

So again, the intuitive verdict of relevance may be seen to arise from a probabilistic effect of the evidence on propositions that are intimately related, though not identical, to the hypothesis. However, in this case, the propositions in question are *not parts* of the hypothesis on either Gemes' or Fine's account of partial content.[18] The disjuncts

[18] To see this, it suffices to note that none of these propositions—the disjuncts of H_3 and their conjuncts—are even logical consequences of H_3. On Gemes' view, content parts are by definition a special kind of

of H_3 may instead be described as different *ways* for H_3 to hold, and their conjuncts accordingly as parts of ways for H_3 to hold.

Under this diagnosis, the present counter-example to (BI) is suggestive of a different revision of (BI) than that proposed by Gemes. Specifically, the case seems to suggest that if the evidence, or a part of it, is probabilistically relevant to a way for the hypothesis to hold, or perhaps even just to a part of a way for the hypothesis to hold, then this is sufficient for relevance.[19] Put in terms of irrelevance, the envisaged conditions read as follows:

(W) A piece of evidence E is irrelevant to a hypothesis H only if for every way H^* for H to hold, $Pr(H^*|E) = Pr(H^*)$.

(WP) A piece of evidence E is irrelevant to a hypothesis H only if for every part H' of some way H^* for H to hold, $Pr(H'|E) = Pr(H')$.

The notion of a way for a proposition to be true may then be defined as the disjunctive counterpart of the notion of partial content, so that H^* is a way for H to hold just in case $H^* \vee H = H$. In the framework of truthmaker semantics, this is in turn equivalent to the following definition, paralleling in an obvious way the definition of partial content:[20]

(FW) P is a way for Q to hold iff (i) every verifier of P is a verifier of Q
(ii) every falsifier of P is part of a falsifier of Q
(iii) every falsifier of Q has a part that falsifies P

Admittedly, given the counterpart of (W) for parts

(P) A piece of evidence E is irrelevant to a hypothesis H only if for every part H^* for H to hold, $Pr(H^*|E) = Pr(H^*)$.

which is implied by (GI), we are not *forced* by the above example to accept either of (W) or (WP). The reason is that under the Finean interpretation of 'part', (P) already implies that E_3 is relevant to H_3, for the propositions $P \vee \neg Q$ and $\neg P \vee Q$ turn out to be parts of $H_3 = (P \wedge Q) \vee (\neg P \wedge \neg Q)$. But there is no reason to suppose that

Footnote 18 continued
logical consequence. On Fine's view, this is clear from the fact that content parts are conjuncts of the propositions they are part of.

19 There is no corresponding motivation to also take into consideration mere ways for the evidence to hold, or mere parts of such ways. Indeed, for any hypothesis H with $0 < Pr(H) < 1$ and arbitrary P, $(H \vee \neg H) \wedge P = (H \wedge P) \vee (\neg H \wedge P)$ would otherwise turn out relevant to H on the strength of the probabilistic relevance of H to H. This would seem a bad result. Surely, amassing evidence of this sort by procuring arbitrary information P would amount to an objectionable waste of resources in an inquiry into H. Note that the 'converse' claim of relevance, that H is relevant to $(H \vee \neg H) \wedge P$, which I endorse, is not subject to same objection, since it does not yield a recipe for producing lots of irrelevant-seeming evidence. It does, of course, yield a recipe for producing lots of somewhat strange hypotheses to which the evidence at hand is classified as relevant. But although we may at times start with a piece of evidence, and then ask what hypotheses the evidence might sensibly lead us to inquire into, we would not expect a criterion of evidential irrelevance on its own to provide the answer to this question. This point also serves, firstly, to highlight that on my approach—in contrast to the Bayesian and to Gemes' account—irrelevance is *not symmetric*, and secondly, to indicate how this may be justified in terms of the constraints by which I have introduced my target notion of irrelevance. Thanks to an anonymous referee for raising the issue of symmetry here.

20 Cf. (Fine, msa, p. 16); Fine says that P *exactly entails* Q when I say that P is a way for Q to hold.

in general, evidence rendered relevant by (W) is also rendered relevant by (P) or (GI). The pertinent cases concern hypotheses that may be written as disjunctions $H_a \vee H_b$, where the evidence E is probabilistically relevant to each disjunct in such a way that its effects on the disjuncts cancel each other out. There is no reason to infer from this that either E or $H_a \vee H_b$ even *have* proper parts, let alone ones that are probabilistically relevant to each other. I conclude that we have strong reasons for accepting (W).[21]

On the basis of (W) and (P), we can now give an argument for the stronger claim that (WP). For given (W), it is plausible also to accept the strengthening on which the mere relevance (probabilistic or otherwise) of the evidence E to some way H^* for the hypothesis H to hold is sufficient for E's relevance to H. But then suppose that E is probabilistically relevant to part H' of the way H^* for H to hold. Then by (P), E is relevant to H^*. By the strengthening of (W), it follows that E is relevant to H, just as required for (WP).[22]

4.2 Helpful parts

I shall now give an argument that (WP), and even (P), overgenerates relevance, and propose a modification that avoids the problem. Crucial to the argument and the modification is a distinction between what I will call *helpful* and *unhelpful* parts of a way for a proposition to hold. It will help to have a short term for parts of ways for a proposition to hold, so alluding to their status as *c*onjuncts of *d*isjuncts, I will call them cd-parts.

Consider some proposition P and assume that state s is not a verifier of P, but that s is a proper part of some verifier t of P. Then s, we might say, goes some way towards making P true, though not the whole way.[23] However, note that it may still be the case that, as it were, the truth of P is as far away given the state s as it is without s. For we may ask what states can be fused with s so as to yield a verifier of P. And it may be

[21] Note that (W) implies that whenever $0 < \Pr(P) < 1$, P is relevant to $P \vee \neg P$, since P is then probabilistically relevant to P. This is in marked contrast to (BI), and to (GI) on Gemes' account of content parts, on which nothing can be relevant to a logical truth. Since $\Pr(P \vee \neg P)$ is always 1, $P \vee \neg P$ makes for a somewhat peculiar choice of a hypothesis to investigate, so it is not obvious what significance to attach to our result. However, if we wish to allow for rational inquiry into a hypothesis that is a logical truth like $P \vee \neg P$, then the result seems very plausible to me. For P then *is* evidence that bears on $P \vee \neg P$ in a way in which it does not bear on arbitrary $Q \vee \neg Q$, and it seems to me a feature, not a bug, of the present proposal, that it enables us to capture this fact.

[22] It is plausible that (P) may be strengthened in the analogous way, so that consequently E's probabilistic relevance to a way for H' to hold, where H' is part of a way for H to hold, is also sufficient for E's relevance to H. Fortunately, this is already implied by (WP). For in this case, H may be written $((P \vee Q) \wedge R) \vee S$, where E is probabilistically relevant to P. We can then show that $P \wedge R$ is a way for H to hold, and thus P part of a way for H to hold, using that $(P \vee Q) \wedge R = (P \wedge R) \vee (Q \wedge R)$, and thus $H = ((P \wedge R) \vee (Q \wedge R)) \vee S = (P \wedge R) \vee ((Q \wedge R) \vee S)$. The identities used here are implicit in the soundness results of (Fine 2015, Sects. 6, 9).

[23] Note, though, that the fact that s is part of a verifier of P does not rule out that s is also part of a *falsifier* of P. Indeed, there are various possible scenarios in which it would be natural to say that s goes more of the way towards making P false than it goes towards making P true. It bears emphasis, then, that on my use of the phrase, that s goes some way towards making P true does not imply that *on balance*, s goes further along the way to P's truth than to its falsity. (Thanks here to an anonymous referee).

that the only states satisfying this description are themselves already verifiers of P. If so, then in terms of what is still needed to make P true, the state s does not bring us any closer to the truth of P, even though it goes some of the way towards making P true. Now for evidence to be relevant to a hypothesis, I want to suggest, it (or one of its parts) has to be probabilistically relevant to not just any cd-part of the hypothesis, but to one that brings us closer to the hypothesis.

Consider the following example, concerning again a throw of two dice A and B.

(E_4) B came up 1.
(H_4) A came up even or (A came up even and B came up odd).

Let P be the proposition that A came up even, and Q the proposition that B came up odd, so $H_4 = P \vee (P \wedge Q)$. Note that the verifiers and falsifiers of Q stand to truth and falsity of H_4 in the way just described. For consider any state, such as B having come up 1, or 3, or 5, that verifies Q. For any such state s, we may now ask what we can add to that state so that we obtain a verifier of H_4. We then see that the only states of this sort are themselves already verifiers of H_4, namely the states of A having come up 2, or 4, or 6. Likewise consider any state, such as B having come up 2, or 4, or 6, that falsifies Q. For any such state s, we ask what we can add to it so as to obtain a falsifier of H_4. We then see that the only states of this sort are themselves already falsifiers of H_4, namely the states of A having come up 1, or 3, or 5.

Should we consider E_4 relevant to H_4? E_4 is not probabilistically relevant to H_4, which can be seen from the fact that H_4 is logically equivalent to P. But since Q is a cd-part of H_4 and E_4 is probabilistically relevant to Q, E_4 is classified as relevant to H_4 by (WP). Moreover, $H_4 = P \vee (P \wedge Q) = P \wedge (P \vee Q)$, hence $P \vee Q$ is part of H_4. Since E_4 is probabilistically relevant to $P \vee Q$, E_4 is already classified as relevant to H_4 by (P). It seems to me that this is the wrong result. If you were called upon to investigate the hypothesis H_4, it would appear rational for you to discard the information E_4 as irrelevant, and you might rightly be blamed for wasting time if you were to spend it on procuring the information E_4. It is natural to take the reason for this to be that the information E_4 does not get you any closer to an answer to H_4. The above observations concerning the verifiers and falsifiers of E_4 and H_4 give a precise sense in which this is true.

Since the hypothesis H_4 has a somewhat contrived and unnatural logical structure, intuitions about the example are perhaps less firm than we should like them to be to motivate replacing (WP) in the way I have suggested. So let me try to marshal some additional support for this move. Suppose you are interested in a hypothesis H, but it is difficult to obtain any evidence that bears probabilistically on the hypothesis taken as a whole. So you move to considering parts of ways for the hypothesis to hold, to see if data bearing on the probability of these might be more easily obtained. Then the modification I propose can be seen as amounting to the following, very reasonable injunction: Make sure, in selecting a cd-part P of H to collect data on, that any way for P to be true brings us closer to the truth of H, and that any way for P to be false brings us closer to the falsity of H.

Let us state the proposed modification of (WP) more explicitly. First, we define what it is for a proposition P to bring us closer to another Q, or as I shall say, for P to *help* Q.

(H) P helps Q iff for every $s \in P^+ \setminus Q^+$, there is a $t \notin Q^+$ with $s \sqcup t \in Q^+$, and for every $s \in P^- \setminus Q^-$, there is a $t \notin Q^-$ such that $s \sqcup t \in Q^-$

It can be shown that if P helps Q according to this definition, this ensures that P is part of way for Q to hold.[24] We may therefore replace (WP) above with this weaker alternative:

(HELP) A piece of evidence E is irrelevant to a hypothesis H only if for every helper H′ of H, Pr(H′|E) = Pr(H′).

Here, then, is my proposal for a criterion of irrelevance that fits the role intended for the notion of irrelevance in separating what may be discarded and what should be valued in a given context of inquiry:

(IRRE) A piece of evidence E is irrelevant to a hypothesis H iff for every part E′ of E and every helper H′ of H, Pr(H′|E′) = Pr(H′).

5 Conclusion

Gemes has argued convincingly that there is an important distinction between relevant and irrelevant evidence that is not adequately captured by the usual Bayesian criterion. However, his own proposal, while a significant improvement, still has unacceptable consequences. To do better, I have argued, we need to accept that the notion of relevance is *hyperintensional*; it is sensitive to differences in content that may obtain even between logically equivalent propositions.[25] I have then utilized Fine's framework of truthmaker semantics to formulate a hyperintensional criterion of irrelevance that avoids the difficulties that befell the intensional account of Gemes. In view of the advantages the resulting notion of irrelevance enjoys over its Bayesian and Gemesian rivals, it would be very interesting to develop the theory of this notion in detail, determining its formal properties and its relation to other relevant notions definable within the truthmaker framework. This task, however, I have to leave to future work.

[24] We can construct a proposition R of which P is part and which is a way for Q to hold as follows. Let the set of verifiers of R be the set of verifiers of Q that have a part which verifies P. Let the set of falsifiers of R be $(P^- \cup Q^-)^\circ$. It is then straightforward to show that R is a proposition, and that it relates to P and Q in the desired way.

[25] In this respect, the present paper would seem to follow something of a trend. Hyperintensional accounts have in recent years been proposed for many philosophically central concepts, such as essence, ground, conditionals, subject matter, and, closest to our present concerns, confirmation—cf. here esp. Yablo (2015). It is striking that in many cases, considerations of relevance play an important role in motivating the claim to hyperintensionality. It would be very interesting to explore the connections between these debates and the arguments I have here advanced in detail. A particularly tight connection may obtain to *ground*, for which Fine has offered a semantics within the same truthmaker framework we have employed here (cf. Fine (2012a, b)). Indeed, our notion of a cd-part coincides with (a non-factive version of) Fine's notion of a weak partial ground, and our notion of a way for a proposition to hold coincides with (a non-factive version of) Fine's notion of a weak full ground.

Acknowledgements The research for this paper was funded by the Deutsche Forschungsgemeinschaft (Grant KR 4516/1-1), and I gratefully acknowledge the support. An earlier version of this paper was presented at a research colloquium at the University of Hamburg. I thank the members of the audience for their comments and criticisms. Special thanks are due to Sebastian Krug and to my fellow *phlox* members Michael Clark, Yannic Kappes, Martin Lipman, Giovanni Merlo, Stefan Roski, Benjamin Schnieder, and Nathan Wildman. I have also greatly benefitted from a number of conversations with Kit Fine on the general topic of relevance and truthmaker semantics. Finally, I would like to thank an anonymous referee for this journal for some very helpful comments and suggestions which have resulted in a number of improvements.

References

Cohen, L. J. (1994). Some steps towards a general theory of relevance. *Synthese*, *101*, 171–185.
Fine, K. (2012a). Guide to ground. In B. Schnieder & F. Correia (Eds.), *Metaphysical grounding* (pp. 37–80). Cambridge: Cambridge University Press.
Fine, K. (2012b). The pure logic of ground. *The Review of Symbolic Logic*, *5*(1), 1–25.
Fine, K. (2013). A note on partial content. *Analysis*, *73*(3), 413–419.
Fine, K. (2015). Angellic content. *Journal of Philosophical Logic Online First*, 1–28.
Fine, K. (msa) A theory of truth-conditional content I, unpublished manuscript
Fine, K. (msb). A theory of truth-conditional content II, unpublished manuscript
Gemes, K. (1994). A new theory of content I: Basic content. *Journal of Philosophical Logic*, *23*, 596–620.
Gemes, K. (1997). A new theory of content II: Model theory and some alternatives. *Journal of Philosophical Logic*, *26*, 449–476.
Gemes, K. (2007). Irrelevance: Strengthening the Bayesian criteria. *Synthese*, *157*(2), 161–166.
Keynes, J. M. (1929). *A treatise on probability*. London: Macmillan.
Yablo, S. (2015). *Aboutness*. Princeton, NJ: Princeton University Press.

Journal of Philosophical Logic (2022) 51:1175–1213
https://doi.org/10.1007/s10992-022-09663-7

Mighty Belief Revision

Stephan Krämer[1]

Received: 31 August 2021 / Accepted: 24 March 2022 / Published online: 27 June 2022

Abstract
Belief revision theories standardly endorse a principle of intensionality to the effect that ideal doxastic agents do not discriminate between pieces of information that are equivalent within classical logic. I argue that this principle should be rejected. Its failure, on my view, does not require failures of logical omniscience on the part of the agent, but results from a view of the update as *mighty*: as encoding what the agent learns might be the case, as well as what must be. The view is motivated by consideration of a puzzle case, obtained by transposing into the context of belief revision a kind of scenario that Kit Fine has used to argue against intensionalism about counterfactuals. Employing the framework of truthmaker semantics, I go on to develop a novel account of belief revision, based on a conception of the update as mighty, which validates natural hyperintensional counterparts of the usual AGM postulates.

Keywords Belief revision · Truthmaker semantics · Hyperintensionality · AGM · Counterfactuals

1 Introduction

Belief revision theories standardly endorse a principle of *intensionality*, according to which it is a requirement of rationality on ideal doxastic agents that they do not discriminate between pieces of information that are equivalent within classical logic: whatever they are disposed to (come to or continue to) believe upon receiving the one, they are disposed to believe upon receiving the other, and vice versa. In this paper, I argue that, subject to certain qualifications, that principle should be rejected.

✉ Stephan Krämer
stephan.kraemer@uni-hamburg.de

[1] Department of Philosophy, University of Hamburg, Überseering 35, Hamburg, 22297, Germany

Its argued failure does not require failures of logical omniscience on the part of the agent. It results instead from a view of the update as *mighty*—as encoding what the agent learns *might* be the case, as well as what must be.[1] Central to my argument is a puzzle case, obtained by transposing into the context of belief revision a kind of scenario that Kit Fine, in his 'Counterfactuals without Possible Worlds' ([7], see also his [8]), has used to argue against the principle of intensionality for (the antecedents of) counterfactuals.

The structure of the paper is as follows. Section 2 introduces some background assumptions and terminology and gives a more precise statement of the principle of intensionality. Section 3 describes an intensional account of rational belief revision—a form of the popular AGM approach—which is closely related to the standard possible worlds analysis of counterfactuals. Section 4 presents the puzzle cases. Section 5 applies the AGM approach to these cases and argues that it gives the wrong results. Section 6 examines and rejects some prima facie promising ways to respond to the difficulty while retaining intensionality. Section 7 introduces the basic ideas guiding my subsequent development of a *truthmaker-based*, hyperintensional approach. Section 8 introduces the conception of the update as *mighty* underlying the approach and explains how it leads to violations of the principle of intensionality. Section 9 formally articulates some constraints on the rational ways of revising by mighty updates. It is shown that the account delivers the intuitively correct verdicts in the problem cases while retaining those components of the AGM account that are not undermined by those examples. Section 10, finally, describes in more general terms the advantages I take the truthmaker-based approach to offer while identifying some open questions for future research to pursue.

2 Belief Revision and Intensionality

At any given time, doxastic agents like ourselves have a set of beliefs, and they have dispositions to revise their beliefs in certain ways under certain circumstances. For brevity, we shall refer to such dispositions simply as *dispositions to revise*, and we shall refer to relevant circumstances as *occasions for revision*. The combination of a total system of beliefs and a total set of dispositions to revise we may call a *(complete) doxastic state*. Call a complete doxastic state *(ideally rationally) permissible* iff it could be the doxastic state of an ideally rational doxastic agent (short: ideal agent). We may also call a partial doxastic state permissible iff it has a permissible complete extension. The aim of a theory of belief revision, as I here conceive of it, is to capture the general, logico-structural properties that are held by any permissible complete doxastic state.

[1]This sets the view defended here apart from previous approaches rejecting intensionality, which have generally been motivated by the aim of modelling less idealized doxastic agents. For various approaches of this sort, see e.g. [2, 5, 21, 30, 39].

It is standard to assume that for any ideal agent and for any possible occasion for revision, the agent's dispositions to revise determine a *unique* result, i.e. a unique set of beliefs comprising all and only those beliefs the agent would hold after exercising their dispositions. The dispositions to revise of an ideal agent may then be represented by a function mapping every possible occasion for revision to a revised belief system.[2]

Let us call occasions for revision *dynamically equivalent* iff no permissible doxastic state discriminates between them. That is to say, occasions for revision o_1 and o_2 are dynamically equivalent just in case for any function f representing the dispositions to revise in some permissible doxastic state, $f(o_1) = f(o_2)$. It is a standard (if often tacit) assumption that one way to characterize a sufficient condition for dynamic equivalence is in terms of a *proposition* suitably related to the occasion of revision—call this proposition the *update*. This seems quite plausible. Presumably, the rationality or otherwise of a possible response to an occasion for revision can depend only on what the doxastic agent *learns*, or what *information* they *receive*, on that occasion. If what the agent learns on occasion o_1 is the same as what they learn on occasion o_2, then rationality seems to require that the agent make the same adjustments to their beliefs in both situations. Assuming that the totality of what the agent learns can always be represented by a proposition, we may take that proposition to be the update and conclude that occasions for revision with the same update are dynamically equivalent. We shall later say more about how to make these ideas precise. For now, note that given the dynamic equivalence of situations with the same update, for the purposes of a theory of belief revision, we may identify occasions for revisions with their associated updates, and we may represent an agent's dispositions to revise as a function mapping *each possible update*[3] to a revised belief system. We shall also describe possible updates as dynamically equivalent when the associated occasions for revision are. A principle of intensionality for updates may now be stated as follows:

Intensionality For any possible updates P and Q, if P is logically equivalent to Q then P is dynamically equivalent to Q.

In this formulation, the principle presupposes a notion of logical equivalence for updates. The most common approach in the literature is to identify updates with sen-

[2]Of course, there may be occasion to revise the new belief system again. To study the constraints on *iterated* belief revision, we should have to assume either that the initial doxastic state includes dispositions to revise arbitrary belief states (or at least arbitrary ones reachable from the present belief state by some sequence of revisions) by new information, or that the dispositions to revise the initial doxastic state determine not only the new belief state but also new dispositions to revise that belief state. For the purposes of this paper, we restrict attention to singular, i.e. non-iterated belief revision.

[3]By a possible update I mean a proposition which is the update in some possible situation for an ideal doxastic agent. It is a further question whether every update possible in this sense is also possibly true. For present purposes, though, we may assume that this is so.

tences of some formal, propositional language. An alternative is to assume a notion of a logically possible world, and to identify updates with the sets of logically possible worlds in which they are true. Logical equivalence for updates is then simply the identity relation, and so adopting this kind of conception of the update will automatically ensure that *Intensionality* holds. The two approaches may be connected, relative to a chosen formal language, by identifying logically possible worlds with the corresponding maximal consistent sets of sentences of the language.[4]

3 A Possible Worlds Approach

Within a possible worlds framework, we can formulate a prima facie attractive theory of rational belief revision that is closely related to the standard possible worlds account of *counterfactuals*.[5] On this account of counterfactuals, recall, we assume that for any given possible world w, there is an *ordering* of all worlds according to their comparative similarity, in some suitable sense, to w. A counterfactual $A \mathrel{\Box\!\!\rightarrow} C$ is then taken to be true at w iff all those worlds at which A is true which are *closest*—i.e. most similar—to w are worlds at which C is also true.

Under the analogous approach to belief revision, both belief systems and updates are identified with a set of logically possible worlds. A doxastic state accordingly consists of a set B of possible worlds representing the belief system, and a function mapping any set of possible worlds P—the update—to a set of possible worlds $B * P$—the revised belief system. It is assumed that in any ideally rational doxastic state, B is non-empty. The logical constraints on the revision function are stated by appeal to an ordering on the worlds, formally similar to the similarity orderings by which counterfactuals are interpreted.[6] Informally, we may think of the ordering as representing the comparative plausibility of the worlds by the lights of the agent, or perhaps the strength with which the worlds are *excluded* or *disbelieved* by the agent. The worlds at which the agent's beliefs are true are the most plausible ones, which are not excluded or disbelieved at all. All other worlds are excluded, but some more firmly than others, in which case they are treated as less plausible.

[4] Analogous questions of granularity may also be raised with respect to the other component of a doxastic state, i.e. the total system of beliefs. Our focus in this paper, though, will be on the intensionality or otherwise of the update.

[5] The classical sources are Stalnaker's [33] and Lewis's [25].

[6] The idea of basing belief revision theory semantically on an ordering of worlds is familiar in the literature. My presentation here largely follows Huber [20]. The approach based on plausibility orderings can equivalently be stated in terms of plausibility spheres, just like the Lewis/Stalnaker semantics can be stated in terms of similarity spheres instead of similarity orderings. Modulo the subtleties surrounding the condition ($\leq$4) mentioned below, the present approach is thus equivalent to the sphere-based approach first described by Grove [17]. The same kind of ordering of worlds, under the label of *faithful assignments*, is used by Katsuno and Mendelzon [22] to prove a representation theorem for AGM revision operations (the counterpart of ($\leq$4) is not needed there, since the authors assume the underlying language to be based on a finite set of propositional letters). For a useful overview of equivalent characterizations of the AGM model, see chapter 4 of Fermé and Hansson's [6] and especially section 4.1, which discusses the various possible worlds based models.

More formally, given a belief system B, we call a *plausibility ordering centered on B* any two-place relation $\leq$ on the worlds such that for all worlds w, v, u:

(≤ 1)	$w \leq v$ or $v \leq w$
(≤ 2)	if $w \leq v$ and $v \leq u$ then $w \leq u$
(≤ 3)	$w \in B$ iff $w \leq z$ for all $z \in W$
(≤ 4)	if $\emptyset \subset A \subseteq W$, then $\{z \in A: z \leq y$ whenever $y \in A\} \neq \emptyset$

Informally, $w \leq v$ means that w is at least as plausible as v. (≤ 1)–(≤ 3) ensure that the plausibility ordering is transitive, that any two worlds are comparable in terms of their plausibility, and that all and only the members of B are maximally plausible. The final condition (≤ 4), as we shall see, is of special importance for our purposes: it ensures that any non-empty set of worlds has a maximally plausible member. The crucial claim is now that for any ideally rational doxastic state with belief system B and revision function $*$, there exists a plausibility ordering $\leq$ of the worlds centered on B such that for every possible update P, $B * P = \{z \in P: z \leq y$ whenever $y \in P\}$: the revision by update P is always the set of the most plausible P-worlds.

This account of belief revision is near-equivalent to the popular AGM theory of belief revision ([1]).[7] Within AGM, a belief system is modelled by a set K of sentences of a propositional language $\mathcal{L}$, an update is modelled by a single sentence α from $\mathcal{L}$, and the dispositions to revise are modelled by a function mapping K and any such α to a new belief system $K * \alpha$. The theory then includes the following eight postulates to be satisfied by any ideally rational belief set and revision function (where $K + \alpha$ is the closure under logical consequence of $K \cup \{\alpha\}$):

Closure	$K * \alpha$ is closed under logical consequence
Success	$\alpha \in K * \alpha$
Inclusion	$K + \alpha \supseteq K * \alpha$
Vacuity	$K * \alpha \supseteq K + \alpha$ if $K \cup \{\alpha\}$ is consistent
Consistency	$K * \alpha$ is consistent if α is
Intensionality	$K * \alpha = K * \beta$ if α and β are logically equivalent
Superexpansion	$(K * \alpha) + \beta \supseteq K * (\alpha \wedge \beta)$
Subexpansion	$K * (\alpha \wedge \beta) \supseteq (K * \alpha) + \beta$ if $(K * \alpha) \cup \beta$ is consistent

We shall sometimes refer to the last two postulates as the *supplementary* AGM postulates, and to the other six as the *basic* AGM postulates.[8]

It is known that from any AGM belief set K and revision function $*$, one can construct a possible worlds interpretation of $\mathcal{L}$ and an ordering $\leq$ on the worlds, centered on the set of worlds at which K is true, which satisfies conditions (≤ 1)–(≤ 3)

[7] For an accessible introduction, see again [20] or [19].

[8] In the literature they are also called the *basic* and *supplementary Gärdenfors postulates* for revision, respectively; cf. [19, sec. 3]. There is some variation also in the labels for the individual postulates. *Intensionality* is sometimes called *Extensionality*, and *Superexpansion* and *Subexpansion* are sometimes just called *Conjunction 1* and *2*.

as well as a weakened version of ($\leq$4).[9] Say that a formula $\alpha \in \mathcal{L}$ *expresses* a set of worlds (under the given interpretation) iff it is true at exactly those worlds. Then the relevant weakening of ($\leq$4) says that any non-empty set of worlds *expressed by some formula in* $\mathcal{L}$ has a maximally plausible member. Conversely, given a plausibility ordering centered on a set of worlds $B \subseteq W$ and an interpretation of $\mathcal{L}$ relative to W, one can define a corresponding AGM-style revision operator for the belief set true exactly at the members of B which satisfies the AGM postulates.[10] For most of the discussion to follow, we may treat AGM and the plausibility based possible worlds approach as equivalent, and refer to them indiscriminately as the AGM approach or the possible worlds approach.

4 Of Dominos and Matches

In this section, I shall describe some partial doxastic states and argue that they are rationally permissible, i.e. that they have complete extensions that could be the doxastic state of an ideally rational agent. In the next section I will then show that the permissibility of these doxastic states is in conflict with the AGM approach.

For definiteness, imagine a particular doxastic agent, Dom. His relevant beliefs concern an infinite sequence of domino stones, arranged like this[11]

[] [] [] [] . . .

We assume that each stone can only fall to the right, not to the left. We refer to the stones as $s_1, s_2, \ldots$, respectively, with s_1 being the leftmost stone, and s_{n+1} the stone immediately to the right of s_n. Let F_n be the proposition that stone n fell. We suppose that as a matter of fact, no stone fell.

For each n, Dom believes that $\neg F_n$. Furthermore, he has the following dispositions to revise: If Dom were to learn that F_n, then he would come to believe that F_m for all m with $m \geq n$. At the same time, he would retain the belief that $\neg F_m$ for all m with $m < n$.

It will be helpful to introduce some notation to describe the doxastic state more succinctly. Let us write $P \Rightarrow Q$ for the claim that Dom is disposed to (come to or continue to) believe that Q upon learning that P—i.e. on any occasion for revision whose update is the proposition that P.[12] Slightly artificially, we write $\Rightarrow Q$ to say

[9] That only the weakened version of ($\leq$4) is guaranteed is why I said the above account is *near*-equivalent to AGM. We will see below that this detail is somewhat relevant to our purposes.

[10] These results are due to Adam Grove ([17]).

[11] The scenario is essentially identical to the first example described by Fine [7], except that Fine's scenario features rocks instead of domino stones. Note that our case strictly requires only that our agent *has* the relevant beliefs about domino stones, not that these beliefs are accurate. But for presentational purposes it seemed helpful to me to suppose the situation to be as the agent believes it to be.

[12] The reason for using this notation is that it helps bring out more clearly the connection to counterfactual logic. This will help relating the present discussion to Fine's, and in particular means that his central proofs carry over to our setting without any changes.—The idea of interpreting a conditional in terms of belief revision in this way is again familiar from previous work, most notably in connection with the Ramsey Test; see e.g. [14–16] and [26]; see also [6, p. 85f].

that Dom believes that Q (since this is like saying that he is disposed to believe that Q upon learning nothing). We can now summarize the partial doxastic state $\mathcal{D}$ we have ascribed to Dom as follows:

$$(B) \qquad \Rightarrow \neg F_n \text{ for all } n$$
$$(D.+)\ F_n \Rightarrow F_m \text{ for any } m \geq n$$
$$(D.-)\ F_n \Rightarrow \neg F_m \text{ for any } m < n$$

Dom's dispositions to revise may be seen simply as reflecting an awareness of the nature of the setup as described above: since each stone can only fall to the right, knocking over every subsequent stone, if Dom learns F_n he also comes to believe $F_{n+1},\ F_{n+2},\ \ldots$ and accordingly gives up $\neg F_{n+1},\ \neg F_{n+2},\ \ldots$ but since each stone can only fall to the right, he has no reason to give up $\neg F_{n-1}$, or $F_{n-2}, \ldots$ At first glance, it would therefore appear that the doxastic state is permissible.

At second glance, one might worry that perhaps Dom does have some reason to give up $\neg F_{n-1}$ upon learning F_n. For given that, say, the second stone fell, it is natural to ask what caused it to fall. And since one of the things that may have caused this is the first stone falling, perhaps Dom does have some reason to allow for the possibility that the first stone fell as well. This objection may be avoided, however, by modifying the example, at the cost of some additional complexity.

The difficulty arises because in the case of the dominos, the truth of F_n would be *responsible* for the truth of F_{n+1}, and ultimately F_m whenever $m > n$. But this is an inessential feature of the example. Indeed, for roughly similar reasons, Fine has already described a version of the example which lacks this feature ([7, p. 224f]). Transposed to the belief revision setting, the case runs as follows. We imagine another doxastic agent, Matt. His relevant beliefs are that there is an infinity of matches $m_1, m_2, \ldots$, placed in causal isolation from one another, each of them in an environment maximally conducive to the match lighting upon being struck, but none of them actually struck. Now let S_n be the proposition that match m_n is struck, let L_n be the proposition that match m_n lights, and let W_n be the proposition that match m_n is wet. Let S be $S_1 \wedge S_2 \wedge \ldots$, so S says that each match is struck. Then F_n is $S \wedge ((W_n \wedge \neg L_n) \wedge (W_{n+1} \wedge \neg L_{n+1}) \wedge \ldots)$. So F_n says that each match is struck, but every match from n onwards is wet and does not light.

Note that for all n, F_n contains F_{n+1} as a conjunct. So in this version of the case, the dispositions ascribed in (D.+) are simply dispositions to believe conjuncts of conjunctive information received, and therefore clearly permissible. So let us turn to the dispositions ascribed in (D.−), and let us consider the instance $F_2 \Rightarrow \neg F_1$. Note first that since Matt believes each match to be in an environment maximally conducive to its lighting upon being struck, learning that the first match is struck (S_1) would give Matt good reason to believe that match 1 lights (L_1). So it seems rational for Matt to believe that L_1 upon learning that S_1. Now F_2 is the conjunction of S_1 with some information exclusively about other matches, believed by Matt to be causally isolated from match 1. None of this additional information seems in any way to undermine the support that S_1—match 1 is struck—offers for L_1—match 1 lights.[13] So it also seems rational for Matt to believe that L_1 upon learning that F_2. But now note that L_1 logically entails $\neg F_1$, since F_1 contains $\neg L_1$ as a conjunct. So since it seems clearly rational for Matt to form the belief that L_1 upon learning F_2, and L_1 logically entails $\neg F_1$, it also seems clearly rational for Matt to retain the belief that $\neg F_1$ upon learning that F_2. In other words, learning that F_2 not only provides no reason for Matt to give up the belief that $\neg F_1$, it gives Matt additional support for that belief. Parallel considerations apply with equal force to the other instances of (D.−). I conclude that at least in this more complicated variant, the beliefs and dispositions to revise we have ascribed to Matt are jointly rationally permissible.

Now consider the infinite disjunction $F_1 \vee F_2 \vee \ldots$, and let us use F to abbreviate it. Assuming that the proposition that F is also a possible update, how should Matt be disposed to revise his beliefs upon learning that F? It is clear that there has to be *some* number n such that it is permissible for Matt to give up the belief that $\neg F_n$ upon learning that F. After all, if Matt were to retain each belief that $\neg F_n$ and add the belief that F, the resulting belief system would be inconsistent. We can also say something more specific, it seems to me. For it is hard to see how giving up $\neg F_n$ could be permissible for Matt for the case of, say, $n = 17$ but not for $n = 1$. So it also seems safe to assume that it is permissible for Matt to give up the belief that $\neg F_1$ upon learning that F.

We may summarize the central results of this section as follows: There is some permissible doxastic state which extends $\mathcal{D}$ and which, for some n, includes the disposition to give up the belief that $\neg F_n$ upon learning that F. In particular, there is some permissible doxastic state extending $\mathcal{D}$ and including the disposition to give up the belief that $\neg F_1$ upon learning that F.

In the next section, I will show that these results conflict with the AGM approach. Before that, let me address a kind of dismissive attitude towards these scenarios that some readers may be tempted to adopt. Clearly, both the domino- and the match-example are somewhat unrealistic. There are no infinite sequences of domino stones, and no infinite collections of matches in causal isolation from one another. So what, one might therefore ask, if our theory of belief revision has implausible implications with respect to such bizarre and silly cases? What matters, surely, is how belief systems relevantly similar to our own may be rationally revised, and the problematic kinds of doxastic states do not seem very similar to our own!

In response, it should be noted, firstly, that the specific subject matter of the above examples is of course not essential to the problem that they give rise to. All we need to generate that problem is an instance of the *general structure* exhibited by the cases of the dominos and the matches. So the objection can succeed only if all instances of this structure are silly. But that is not so. As Fine ([8, p. 36]) points out, one way

[13]Perhaps one might object that even though the matches are assumed to be causally isolated from one another, since according to F_n, the fates of matches n and onwards are so similar, it is still rational to suspect some kind of systematic explanation, which could then also suggest that earlier matches suffered the same fate. But it is not even necessary to suppose that events in the different regions are similar in this way. All we need is that S_n always says that some 'trigger'-event occurred, that L_n says that the corresponding standard result occurred, and that W_n says that some corresponding 'blocker'-condition obtained. (Cf. [7, p. 225], see also [8, p. 35, fn. 1].)

to obtain more realistic instances is by considering, instead of infinite sequences of objects, infinite sequences of values of some quantity capable of continuous change, or at least taken by the agent to be so capable. Thus, we may consider an agent's beliefs concerning the flight of a missile believed to possess an automatic mechanism for correcting any deviations from its intended path (the example is Fine's). The propositions F_1, F_2, ... are now to the effect that the missile deviated by 1 inch off course, that the missile deviated by 1/2 inch off course, ... Since any deviation occurs in a continuous way, upon learning F_n the agent will believe F_m whenever $m \geq n$. But they may rationally retain the belief $\neg F_m$ whenever $m < n$, taking the mechanism to have prevented any greater deviation.

Another idea, more promising in our context of belief revision than in Fine's context of counterfactuals, is to construct an example using actual infinite sequences of abstract objects, such as the sequence of the natural numbers. What we would need is an example of a property with respect to which an ideal agent might initially believe that no number has it, and be disposed, upon learning that n has the property, to form the belief that m has it for all $m \geq n$, and to retain the belief that m does not have it for all $m < n$. Indeed, we might approximate the structure of the match example by letting F_n say that (a) for each n, attempts have been made to prove that n has the property, and (b) a proof has been found for each $m \geq n$. The supposition that this kind of situation could arise for some complicated number-theoretic property does not appear problematically unrealistic.

Secondly, the objection overestimates the role that infinity plays for the problem. As we shall shortly see, the relevant assumptions of the intensional approaches yield highly counter-intuitive results even in application to related, finitary contexts. Roughly speaking, the role of infinity is only to turn counter-intuitive results into contradictory ones. Relatedly, the approach I shall eventually propose deviates from its intensional rivals even in finitary contexts, and may be argued to be superior to them even on the basis of considering only finitary contexts.

5 Against the Possible Worlds Approach

We shall now show that the AGM approach is incompatible with the results of the previous section. To start, let us assume Matt is disposed to give up the belief that $\neg F_1$ upon learning that F:

$$(X_1) \quad F \not\Rightarrow \neg F_1$$

This is rationally incompatible, given AGM, with

$$(1) \quad F_1 \Rightarrow F_2$$

$$(2) \quad F_2 \Rightarrow \neg F_1$$

To see this, note first that under the AGM approach, for any propositions P and Q, $P \Rightarrow Q$ holds iff $B * P$ entails Q, i.e. iff $B * P \subseteq Q$. So (X_1) implies that $B * F \not\subseteq \neg F_1$. By the definition of revision in terms of the plausibility ordering, $B * F$ comprises exactly the maximally plausible F-worlds, so (X_1) requires that some maximally plausible F-world be an F_1-world. Call that world w. By (1), every maximally plausible F_1-world is an F_2-world, so w is also an F_2-world. By (2), every maximally plausible F_2-world is a $\neg F_1$-world, so among the F_2-worlds, some world v must be more plausible than w. But every F_2-world is also an F-world, so v is a more plausible F-world than w, contrary to the assumption that w is a maximally plausible F-world. Since we found it to be rationally permissible for Matt to satisfy (X_1), this is a problem.

Moreover, by similar reasoning we can show that *any* instance of

$$(X_n) \quad F \not\Rightarrow \neg F_n \tag{1}$$

is rationally incompatible, under the possible worlds approach, with (D.+) and (D.−), for there is no ordering of the worlds that satisfies the conditions (1)–(4) on plausibility orderings and that validates the dispositions in (D.+) and (D.−). In particular, any such ordering that respects (D.+) and (D.−) is such that there is no maximal F-world. For suppose w is an F-world, and let m be some number such that w is an F_m-world. Suppose for contradiction that w is a maximal F-world. Then in particular, w is a maximal F_m-world. By (D.+), every maximal F_m-world is also an F_{m+1}-world. By (D.−), no maximal F_{m+1}-world is an F_m-world. So w is not a maximal F_{m+1}-world, and hence not a maximal F-world after all.

Alternatively, as Fine shows ([7, pp. 244ff]), we can also derive all instances of $F \Rightarrow \neg F_n$ from (D.+) and (D.−) using only the following inference rules, all of which are valid under the possible worlds approach:

Substitution	$P \Rightarrow Q \ / \ P' \Rightarrow Q$	[if P and P' are logically equivalent]
Entailment	$/ \ P \Rightarrow Q$	[if P logically entails Q]
Transitivity	$P \Rightarrow Q,\ P \wedge Q \Rightarrow R \ / \ P \Rightarrow R$	
Conjunction	$P \Rightarrow Q,\ P \Rightarrow R \ / \ P \Rightarrow Q \wedge R$[14]	
Disjunction	$P \Rightarrow R,\ Q \Rightarrow R \ / \ P \vee Q \Rightarrow R$[15]	

The complete proof of this result is fairly long and complicated, so I shall refrain from reproducing it here. To see where the reasoning of the proof might best be resisted, and thus which rule might best be given up, it is more helpful to present it in more informal terms. And since the match-case is rather complex and hard to think about, in commenting on the various steps, I will use the dominos-example again. Let me first explain why giving up $\neg F_1$ in response to F would involve a violation of the rules. We may divide the reasoning into three main steps.

The first step is an application of *Substitution*, taking us from $F_2 \Rightarrow \neg F_1$—an instance of (D.−)—to $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$. In the domino case, this says that given that Dom is disposed to retain the belief that the first stone stands given the information that the second fell, he must also be disposed to retain that belief given

[14] Fine also has an infinitary version of this rule, allowing us to infer $P \Rightarrow Q_1 \wedge Q_2 \wedge \ldots$ from $P \Rightarrow Q_1, P \Rightarrow Q_2, \ldots$. Using this rule we could show that conforming to the rules would lead, in the case at hand, to Matt's believing an outright contradiction upon learning F. But it seems bad enough if Matt ends up with an unsatisfiable belief system, accepting an infinite disjunction while rejecting each disjunct. For this result we only need the finitary rule.

[15] As Fine points out, we actually require only a weaker rule with the added condition that P and Q be logically exclusive. The difference is not essential for present purposes, so for simplicity, I've here stated the stronger one.

the information that either the first and second, or not the first but the second stone fell.

The second step is to infer from $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ that $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$: Dom must also retain the belief that the first stone stands upon learning that the either the first stone fell or not the first but the second fell. The justification for this is that Dom is disposed to form the belief that F_2 given the information that F_1. Because of this, for Dom, learning F_1 and learning $F_1 \wedge F_2$ effectively come to the same thing, and the same is then true for learning $F_1 \vee (\neg F_1 \wedge F_2)$ and $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2)$.

The third step is to infer from $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ that $F \Rightarrow \neg F_1$. Here, the idea may be described as follows. Learning F presents Dom with a choice: he needs to pick *some* stone s_n as the left-most stone for which to give up the belief that $\neg F_n$. Now $F_1 \vee (\neg F_1 \wedge F_2)$ says that either s_1 or s_2 is the first stone to fall. So learning $F_1 \vee (\neg F_1 \wedge F_2)$ presents Dom with a related choice: he needs to pick some stone $s_n \in \{s_1, s_2\}$ as the left-most stone for which to give up the belief that $\neg F_n$. Now the point is that if Dom does not pick s_1 among the options s_1 and s_2, he cannot rationally pick s_1 among the options $s_1, s_2, \ldots$ Put another way, that $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ means that Dom prefers the scenario in which s_2 is the left-most stone to fall to the scenario in which s_1 is the left-most stone to fall. But giving up $\neg F_1$ in response to F would mean *not* preferring any alternative scenario to the scenario with even s_1 falling. So in particular, it would mean not preferring a scenario s_2 as the left-most stone to fall to the scenario with even s_1 falling.

So if Dom is to conform to the above rules, he must retain $\neg F_1$ upon learning F, and hence conclude that one of the other stones fell, i.e. $F_2 \vee F_3 \vee \ldots$ But the same considerations that prevent him from giving up $\neg F_1$ to accommodate F also prevent him from giving up F_2 to accommodate $F_2 \vee F_3 \vee \ldots$, so in the end he is prevented from giving up any $\neg F_n$. Resisting this final part of the argument seems hopeless. As mentioned before, it simply beggars belief that general constraints of rationality should prevent Dom, in the case at hand, from giving up $\neg F_1$ in response to F, while allowing him to give up, say, $\neg F_{17}$.

Applying AGM theory proper to the examples is not completely straightforward, since the examples involve infinite (conjunctions and) disjunctions, and AGM, strictly speaking, is concerned only with finitary propositional languages. Still, we may consider a trivial extension of AGM to languages with infinite conjunction and disjunction, in which we simply retain all the usual postulates. In this extension of AGM, the above rules can all be derived, and thus the proof that Dom and Matt won't be allowed to give up any belief of the form $\neg F_n$ can be carried out.

But we can also adjust the example so as to do without any infinitely long sentences. Instead, we may replace each infinite conjunction and each infinite disjunction used in our argument by a propositional letter, interpreted as expressing the same proposition as the infinitary sentence it replaces. If it is objected that these propositions may not be graspable by finite thinkers, we can instead let the propositional letters express the universal quantifications corresponding to the infinite conjunctions and the existential quantifications corresponding to the infinite disjunctions.

Since the dispositions ascribed in (D.+) and (D.−), under this modification, concern the same propositions as before—or perhaps quantificational counterparts—they are no less reasonable than before. So we still find that there can be no maximally plausible F-worlds. Since our background language now has a propositional letter true in exactly the F-worlds, it follows that no ordering of the worlds can satisfy ($\leq$1)–($\leq$3) together with the weakened version of ($\leq$4). And so we can infer by the mentioned equivalence that no AGM-revision operation can accord with (D.+) and (D.−) under the finitary replacement.

Most of the derivation given by Fine also still goes through under this modification. Infinitary sentences are involved only in the third step of the argument as described above, in which we infer $F \Rightarrow \neg F_1$ from

$$(3) \qquad F_1 \vee \left(\neg F_1 \wedge F_2\right) \Rightarrow \neg F_1$$

Formally, the way the derivation works is this. By *Entailment*, we also have

$$(4) \qquad \neg F_1 \wedge \neg F_2 \wedge \left(F_3 \vee F_4 \vee \dots\right) \Rightarrow \neg F_1$$

By *Disjunction*, we obtain

$$(5) \qquad \left(F_1 \vee \left(\neg F_1 \wedge F_2\right)\right) \vee \left(\neg F_1 \wedge \neg F_2 \wedge \left(F_3 \vee F_4 \vee \dots\right)\right) \Rightarrow \neg F_1$$

Now the point is that this big disjunction is logically equivalent to $F = F_1 \vee F_2 \vee \dots$, so that by *Substitution* we may infer $F \Rightarrow \neg F_1$.

But now let F and F^3 be propositional letters expressing the proposition that some stone fell, and that some stone other than the first two fell, respectively. Then *Entailment* and *Disjunction* also give us

$$(4') \qquad \neg F_1 \wedge \neg F_2 \wedge F^3 \Rightarrow \neg F_1$$

$$(5') \qquad \left(F_1 \vee \left(\neg F_1 \wedge F_2\right)\right) \vee \left(\neg F_1 \wedge \neg F_2 \wedge F^3\right) \Rightarrow \neg F_1$$

Now the antecedent in (5′) is not logically equivalent to F, so we cannot infer $F \Rightarrow \neg F_1$ simply by an application of *Substitution*. But it is very plausible to assume that it must always be rationally permissible for the agent to treat $\left(F_1 \vee \left(\neg F_1 \wedge F_2\right)\right) \vee \left(\neg F_1 \wedge \neg F_2 \wedge F^3\right)$ as equivalent to F in his dispositions to revise. So we may simply make it a further non-logical *assumption* of the case, in addition to (D.+) and (D.−), that the agent's dispositions satisfy this condition. Given this assumption, we may then infer $F \Rightarrow \neg F_1$, and similarly for all other $\neg F_n$. In this way, even without the use of infinitary sentences, we obtain examples of ideally rational doxastic states that violate some of the AGM principles.

6 Against Intensionalist Responses

We saw that the doxastic states described, in virtue of satisfying (D.+) and (D.−), yield a violation of the condition ($\leq$4) of the possible worlds approach, requiring each set of possible worlds—or each expressible set of worlds in case of the weakened version—to have a maximally plausible member. One obvious idea for responding to the problem while retaining much of the original framework is therefore to drop this condition. There is even a precedent for this move for the case of counterfactuals, as the counterpart to ($\leq$4) in this setting is the so-called *limit assumption*, famously rejected by Lewis.

Without ($\leq$4), we can no longer define the revision by update P as the set of the maximally plausible P-worlds. How are we to define it instead? Lewis's proposal for truth-conditions for counterfactuals is of no help. Lewis takes $P \mathrel{\Box\!\!\rightarrow} Q$ to be true iff Q is true in all sufficiently close P-worlds, i.e. iff by restricting attention more and more to ever closer P-worlds, eventually we will be left only with Q-worlds. The simplest way to see that this won't help is to note that the Lewis-style truth-conditions for $P \Rightarrow Q$ actually validate all the above inference rules.[16]

A natural idea at this point is that the new belief state, in cases where there are no maximally plausible updates, should simply contain *all* the update-worlds.[17] At first glance, this may look attractive. It allows (D.+) and (D.−) to hold, while also allowing that the agent gives up $\neg F_1$ upon learning F, since some F-worlds are F_1-worlds. But at second glance it becomes clear that this suggestion throws out the baby with the bathwater. For the proposal does not allow our agent to have *any* beliefs, post-revision, save for those entailed by the update F. For example, our agent is not allowed to believe, post-revision, that if F_1 then F_2, since it is compatible with the truth of F that $F_1 \wedge \neg F_2$. But it is clearly rational in our scenario to retain the belief that F_2 if F_1, and so the proposal still misclassifies rational doxastic states as irrational.

Perhaps, then, we might give up on the idea that every rational revision function must be definable in terms of a plausibility ordering. Instead, we might say merely that any rational revision function must *conform*, in some suitable sense, to a plausibility ordering, and allow that there may be more than one revision function conforming to a given plausibility ordering. A natural first suggestion would be to take a revision function to conform to a plausibility ordering iff it maps any update P to the set of maximally plausible P-worlds if that set is non-empty, and to some upwards closed non-empty subset of P if not, where a subset P' of P is upwards closed iff P' includes every P-world that is more plausible than some world in P'.[18]

In terms of the inference rules employed in Fine's derivation, this proposal invalidates the *Disjunction* rule. In particular, it leads to the rejection of the inference from (3) and (4′) to (5′):

$$(3) \qquad F_1 \vee \left(\neg F_1 \wedge F_2\right) \Rightarrow \neg F_1$$

$$(4') \qquad \neg F_1 \wedge \neg F_2 \wedge F^3 \Rightarrow \neg F_1$$

$$(5') \qquad \left(F_1 \vee \left(\neg F_1 \wedge F_2\right)\right) \vee (\neg F_1 \wedge \neg F_2 \wedge F^3) \Rightarrow \neg F_1$$

[16]There is a rule that is invalidated by adopting the Lewis-style truth-conditions, namely the infinitary version of the conjunction rule (cf. [7, p. 225]). As mentioned before, this rule is not required for our purposes.

[17]This corresponds to the idea considered by Fine [7, p. 228f] of taking $P \mathrel{\Box\!\!\rightarrow} Q$ to be true iff Q is true in all the closest *and all the stranded* P-worlds, where a P-world is stranded iff there is no closest world closer than it. Fine's most important objection against the proposal is analogous to my criticism in the main text.

[18]Probably, one should then impose some further constraints on how the choices of subsets for different updates have to relate. For instance, any world in the revision by $F_2 \vee F_3 \vee \ldots$ should probably also be included in the revision by $F_1 \vee F_2 \vee F_3 \vee \ldots$

In terms of the AGM postulates, the proposal invalidates the postulate of *Superexpansion*, which says that the result of revising with a proposition P, conjoined with Q, entails the result of revising with $P \wedge Q$. To see how this fails, note that the result of revising with F, under the present proposal, is compatible with F_1, and remains so when conjoined with F_2. At the same time, since $F \wedge F_2$ is logically equivalent with F_2, the result of revising with $F \wedge F_2$ is not compatible with F_1, since the belief that $\neg F_1$ is retained in the revision by F_2.

Although an improvement over the previous attempts, this strategy is still unsatisfactory. For the proposal to be adequate, two conditions must be satisfied. Firstly, the complete extensions of the doxastic state that it classifies as permissible must really be so. Secondly, it must classify every permissible extension of the state as permissible. With respect to both conditions, there are good reasons to be skeptical.

Regarding the first condition, the problem is that the rejected applications of *Disjunction* and *Superexpansion* are intuitively very plausible. In the case of *Disjunction*, we assume that upon learning that $F_1 \vee (\neg F_1 \wedge F_2)$—all matches are struck, but all matches from the first or the second onwards are wet and do not light—Matt retains the belief that $\neg F_1$, and thus excludes the possibility that the first match is wet and does not light. He also retains that belief, obviously, upon learning that $\neg F_1 \wedge \neg F_2 \wedge F^3$. How can it then be rational for Matt not to retain the same belief—and thus to allow for the possibility that the first match is wet—upon learning the disjunction of these two pieces of information?

The case of *Superexpansion* seems even more compelling. We take for granted that Matt gives up the belief that $\neg F_1$—and so allows for the possibility that the first match is wet and does not light—upon learning that all matches are struck, but all matches from some match onwards are wet and do not light. But then how can it be rational to retain the belief that $\neg F_1$—and thus exclude the possibility that the first match is wet and does not light—upon receiving the same information, with the addition that either match 1 or match 2 is the first match to be wet and fail to light?

Regarding the second condition, there are strong reasons to think that there are other permissible extensions of the doxastic state than those envisaged under the present proposal. For instance, it seems very plausible that it should be permissible for Matt's doxastic state to be such that

$$(6) \quad F_1 \vee \ldots \vee F_{100} \not\Rightarrow \neg F_{99}$$

That is, it should be permissible for Matt to be disposed to give up the belief that $\neg F_{99}$ upon learning that $F_1 \vee \ldots \vee F_{100}$. For consider what $F_1 \vee \ldots \vee F_{100}$ says. It says that all matches are struck, and that for some match m_k among the first 100, all matches from m_k onwards are wet and do not light. It would seem quite bizarre for Matt, upon receiving this information, to retain the belief that $\neg F_{99}$, and thus to conclude that m_k must have been m_{100}, i.e. that it must have been match 100 that is the first in the sequence to be wet and fail to light. It certainly does not seem as though having the dispositions in (D.+) and (D.−) *requires* Matt to respond in this way to the information that $F_1 \vee \ldots \vee F_{100}$.[19]

[19] Note that this problem arises in exactly the same way in a finitary version of the example, as the assumption that there are infinitely many matches does no work here. This is way I said at the end of the previous

Similarly, it seems that it should be permissible for Matt to be such that

$$(7) \quad F_1 \vee F_2 \not\Rightarrow \neg F_1$$

That is, it should be permissible for Matt to be disposed to give up the belief that $\neg F_1$ upon learning that all matches are struck, and either all the matches, or all matches from the second onwards, are wet and do not light.

These intuitions are in conflict with the principle of intensionality. For under the interpretation given in the match case, F_1 contains F_2 as a conjunct, and so $F_1 \vee F_2$ is logically equivalent to F_2—and upon learning that F_2, by (D.−), Matt is disposed to retain the belief that $\neg F_1$. Likewise, all of $F_1, \ldots, F_{99}$ contain F_{100} as a conjunct, so $F_1 \vee \ldots \vee F_{100}$ is logically equivalent to F_{100}—and upon learning that F_{100}, by (D.−), Matt is disposed to retain the belief that $\neg F_{99}$. Let us see, then, where we can get by dropping the assumption of intensionality and trying to accommodate these intuitions.

7 Towards a Hyperintensional Solution

We begin by sketching a general method for revising one's beliefs that Matt might be seen to follow and that would lead to his conforming to the intuitions just observed. Both (6) and (7) concern how Matt revises his beliefs by a disjunctive piece of information. A very natural idea is that he does this by forming the disjunction of the results of revising his beliefs by each disjunct. Thus, if we write B for Matt's initial beliefs and $*$ for his revision function, the idea is that $B * (F_1 \vee \ldots \vee F_{100}) = (B * F_1) \vee \ldots \vee (B * F_{100})$, and $B * (F_1 \vee F_2) = (B * F_1) \vee (B * F_2)$.[20] If so, since $B * F_{99}$, for example, does not entail that $\neg F_{99}$, then neither does $B * (F_1 \vee \ldots \vee F_{100})$, in line with (6). And likewise since $B * F_1$ does not entail that $\neg F_1$, then neither does $B * (F_1 \vee F_2)$, in line with (7).

Borrowing a term from Fine ([8, p. 52]), we may call this the method of *wayward revision*, since it involves revising, one by one, by each disjunct of the update, i.e. by each way for the update proposition to be true. (And here, as in Fine, waywardness is considered a good thing.) Revising in this way means that every disjunct of the update is accommodated by the agent in the sense that there is some way for the revised belief system to be true under which that disjunct of the update is true. In other words, for each disjunct Q of the update P, according to the wayward revision by P, it might be that Q. Now to adopt the view that it might be that Q on some occasion for revision—even if one's beliefs previously excluded the possibility that Q—is to treat the occasion as telling one that it might be that Q. In this sense, the method of wayward revision seems to depend on a principle about updates that we

section that the role of infinity is merely to turn counter-intuitive results—like this one—into contradictory ones.

[20] This assumes that belief systems are among the kinds of things to which the operation of disjunction can be applied. This is unproblematic if belief systems are identified with propositions, and slightly less straightforward when belief systems are identified with sets of sentences, though it is clear enough how the notion of disjunction should be extended from sentences to sets of sentences. Still, we shall always be thinking of belief systems as propositions.

may roughly express like this:

(M) A situation with update $P \vee Q$ is a situation telling the agent that it $might$ be that P, and that it $might$ be that Q.

Whenever a pair of situation and update satisfy (M) with respect to all disjuncts of the update, I shall say that the update is *mighty* in that situation. The fully general claim (M) is then that updates are always mighty. A central feature of the approach to belief revision that I want to propose is that it endorses principle (M).

Why should one endorse that principle? One consideration in favour of (M)—not the only one—is that it makes sense of our above described intuitions: our intuitive verdicts regarding the rational ways to revise by updates such as $F_1 \vee \ldots \vee F_{100}$ or $F_1 \vee F_2$ in our puzzle cases seem to arise from a tacit assumption that the updates are mighty in the situations under considerations.

Now it might be objected that it is a mistake to let oneself be guided by these intuitions, since they are simply owed to certain *pragmatic* effects. The thought might be spelled out as follows: The update is supposed to capture the total information received by the agent in the relevant situation. To say that an agent receives the information that $P \vee Q$ pragmatically conveys that, in the situation in question, the agent is given some reason to allow for the possibility that Q. For suppose the agent is given no such reason. Then it will normally be wrong to say that the total information received is that $P \vee Q$, since the agent will then also have received the information that P, which is normally stronger than the information that $P \vee Q$. The exception is if, as in our examples, the propositions that P and that $P \vee Q$ are logically equivalent, since Q is of the form $P \wedge R$. But in such cases it will still be highly misleading to say that the information received is that $P \vee Q$, since it is hard to see what the point could be of presenting the information in this disjunctive form except to convey that the agent is given some reason to allow for the possibility that Q. Still, that the agent is given such a reason is *merely* pragmatically conveyed by the statement that the total information they received is that $P \vee Q$. It is not, or so the objection goes, part of the semantic content of that statement.

The objection misses the point. For all I want to argue here, it may well be that as a sentence of ordinary English, an instance of 'the total information the agent received is that P or Q' does not semantically imply that the agent is given reason to allow that it might be that P, and that it might be that Q. But our ultimate goal here is not to analyse ordinary discourse about people receiving information, it is to develop an adequate theory of belief revision, i.e. to adequately capture the general rationality constraints on doxastic states. As part of this, we require some means to pair occasions for revision with propositions—which we call the updates—in such a way that only dynamically equivalent occasions are assigned the same proposition. A rough and ready informal characterization of a suitable pairing uses talk of what the agent learns, or what information they receive. But in developing our theory of belief revision, we may have occasion to clarify or refine that rough characterization in certain ways. How this should be done depends more on the theoretical requirements of a theory of belief revision, and less on the available readings of the relevant locutions in ordinary English.

What I wish to claim is, firstly, that we *can* pair occasions for revision with propositions as their updates in such a way that (i) only dynamically equivalent situations are paired with the same update, and (ii) updates are always mighty. Secondly, I claim that for the purposes of theorizing about rational belief revision, it is *beneficial* to characterize occasions for revision in terms of these mighty updates. The distinction between pragmatic and semantic implications has little bearing on these claims. In defence of these claims, I will develop a conception of updates as mighty which is based on the framework of truthmaker semantics (Section 8), formally characterize a class of permissible doxastic states within the truthmaker framework and show that they satisfy versions of all the usual AGM postulates save for intensionality (Section 9), and finally highlight what I take to be the important general advantages, apart from our puzzle cases, of the resulting approach and especially the conception of updates as mighty (Section 10).[21]

8 Mighty Truthmaker Updates

To begin, let me make two important initial clarifications regarding the notion of the update which are independent of any issues around mightiness or hyperintensionality. The first is that I take the update to represent the information the agent *takes* themselves to obtain in the given situation, or perhaps better: the information the agent *treats* the situation as providing them with. In particular, if there is also a distinct notion of what information a situation *really* provides a given agent with, whether or not the agent regards and treats the situation accordingly, then that is not what I intend to capture in the update. An example may help to make this clearer. Suppose I have the kind of visual experience that would normally lead to me coming to know that my neighbour is walking towards my house. The experience is caused in the appropriate sort of way by my neighbour walking towards my house, my visual system is as it should be, and so on. But suppose further that I have misleading evidence to take my visual system to be compromised, and thereby to doubt the veridicality of my experience. In one sense, perhaps, this is a situation in which I receive the information that my neighbour is walking towards my house—it is just that circumstances are such as to (rationally) prevent my uptake of that information. But in the sense I intend, this is not a situation in which I receive the information that my neighbour is walking towards my house. For it is not a situation which I treat as giving me this information. Conversely, a situation in which someone tells me that P, and I trust the speaker, would be a situation in which, in the intended sense, I receive the information that P, even if the speaker is actually lying, and it is false that P.

A second, in some ways complementary clarification is that I take the update to represent what the subject treats *the situation*—on its own, as it were—as telling them. Consider a version of the previous scenario in which I have no doubts about

[21] The truthmaker approach to belief revision sketched here is being developed in much more detail in joint work (in progress) of the author with Kit Fine, Steve Yablo, and Daniel Rothschild. My discussion of the idea in this paper has been influenced by, and has greatly benefited from, our joint work. At the same time, it should not be assumed that my collaborators would agree with everything I say here.

my visual system and accordingly come to the belief that my neighbour is walking towards my house. Suppose further that my wife previously told me that my neighbour is away on holiday, leading me to conclude that my wife was mistaken. In one sense, perhaps, I might be said to treat the situation as providing me with the information that my wife was mistaken. But this seems to be a case in which, *in the course of revising* my belief system in the light of the new information, I come to acquire this belief. It is not a case in which the relevant belief is part of the information I treat the situation *on its own* as providing me with.

Both these stipulations are reasonable independently of the questions of mightiness and hyperintensionality. Unless we make the first stipulation, it is doubtful that rationality requires the agent to come to believe the update.[22] Unless we make the second stipulation, we lose the distinction between the interpretation of a situation by an agent on the one hand and the resulting adjustment of their previous beliefs on the other.

We may thus think of the process of belief revision as divided into two stages. The first stage consists of the agent *interpreting* the situation in which they find themselves, and deciding what to take it as telling them. The second consists of the agent revising their beliefs in light of what they've taken the situation to tell them. The role of the update is to represent the outcome of stage one. In explaining our conception of the update, what we need to explain is therefore what it says about how the agent interprets the given situation that we are assigning to it a particular update.

The conception of updates I wish to propose is intended to render them mighty, so that by assigning to a situation the update $P \vee Q$, we are saying, among other things, that the agent interprets the situation as telling them that it might be that P, and that it might be that Q. The condition of the situation telling the agent that it might be that P here is to be understood in a specific, comparatively demanding way. In a weak sense, we might say that the situation tells the agent that it might be that P whenever the situation, as interpreted by the agent, does not—actively and by itself, as it were—exclude the possibility that P. A more natural interpretation of the condition is more demanding. It requires, we might say, that the situation explicitly presents it as a possibility that P, that it being the case that P would (at least) help account for the situation, or that it being the case that P would (at least) partially constitute the truth of what the agent takes the situation to tell them.

The distinction is difficult to define in independent, non-metaphorical terms, but it is clear and familiar enough. An example may help to illustrate the idea. Suppose my neighbour has twin sons, Bob and Bill. Suppose further that I see someone walking towards my house, and that I see them well enough to be able to tell that it is definitely either Bob or Bill, but I can't tell which. So I take the situation to tell me,

[22] This requirement is implicit in the rule of *Entailment*, and captured in the AGM postulate *Success*; cf. Stalnaker [34] for a similar approach to justifying the *Success* postulate. (I do not mean here to exclude the possibility of fruitfully theorizing about belief revision on the basis of a different conception of the update, not subject to the requirement that the agent takes themselves to come to know the update. But this would constitute a more radical departure from the AGM tradition than I wish here to consider. In the literature, approaches of this sort often go under the label of *non-prioritized* belief revision; for a brief introduction see [19, Section 6.3].)

among other things, that Bob or Bill is coming over. Consequently, some propositions are *incompatible* with the situation as I interpret it, such as any proposition to the effect that both Bob and Bill are away on holidays. Some propositions are *merely compatible* with the situation as I interpret it, such as the proposition that it is sunny in Ohio. And some propositions are explicitly presented as possiblities by the situation, such as the proposition that Bob is coming over, and the proposition that Bill is coming over. These are propositions we might describe as (partially) accounting for the situation I find myself in, as I interpret it, as propositions whose truth would partially constitute the truth of what I take the situation to tell me. Let us call propositions in this final category *explicit possibilities* of the situation (under the agent's interpretation[23]), and those in the former category merely *implicit* possibilities.[24]

The distinction between explicit and implicit possibilities is relevant to how an agent may rationally revise their beliefs. If a proposition is an explicit possibility in a situation, then the situation provides some *reason* for the agent to allow for the possibility of the proposition's being true, even if their original belief system excludes that possibility. Thus, in the example, even if I initially believed both Bob and Bill to be away on holiday, the situation provides some reason for me to allow for the possibility that Bob is coming over, and it provides some reason for me to allow for the possibility that Bill is coming over. But if a proposition is a merely implicit possibility, then the situation does not give the agent reason to allow for the possibility that it is true. If in our example I originally believed it not to be sunny in Ohio, then the situation provides no grounds whatsoever to subsequently allow for the possibility of it being sunny in Ohio.

Crucially, the condition that the situation tells the agent that it might be that P in (M) is to be understood as requiring that the proposition that P is an *explicit* possibility in the situation. So under a conception of updates as mighty, to say that the update in a given situation is $P \vee Q$ is to say, among other things, that the agent is given some reason, in that situation, to allow for the possibility that P, and to allow for the possibility that Q.

We can now argue that if updates are mighty, they *must* be individuated in a hyperintensional way. For assuming intensionality, any given update P can also be written as $P \vee \big(P \wedge Q\big)$, for arbitrary Q. Assuming mightiness, it follows that in any situ-

[23] This qualification will henceforth usually remain tacit.

[24] The distinction between what I have called explicit and implicit possibilities in a situation may be compared to von Wright's distinction between the strong and weak permissions of a system of norms (cf. [36, p. 90]), where an action is weakly permitted iff it is compatible with the system of norms, and strongly permitted iff it is actively singled out, as it were, as permitted by the system of norms. The difficulties in capturing these distinctions within an intensional framework are likewise parallel. Fine [12] proposes a truthmaker semantics for statements of permission that is sensitive to the distinction, and captures it in much the same way that I propose below. In Section 8 of that paper, Fine also adresses the problem of deontic updating and notes the connection to belief revision. The approach to deontic updating Fine sketches is related to the approach to belief revision to be described below, but with a simple mereological construction taking the place of the transition relation invoked below. Related approaches to deontic updating are also pursued by Yablo [37] and Yablo & Rothschild [31], who likewise draw the connection to belief revision.

ation with update P, the agent is told that it might be that $P \wedge Q$, and hence that it might be that Q, for arbitrary Q. Whatever P is, there will be few if any such situations. Conversely, it seems most situations will not be representable by a mighty intensional update. It seems safe to conclude, therefore, that a conception of updates as mighty requires a hyperintensional way of individuating updates, in particular one that allows us to distinguish between pairs of the form P and $P \vee (P \wedge Q)$.[25]

I propose that we model updates as propositions as conceived within the framework of truthmaker semantics.[26] Within this theory, propositions are characterized not (merely) in terms of the possible worlds at which they are true, but in terms of the possible *states* which *make* them true.[27] Informally, a possible state may be thought of as a (proper or improper) *part* or *fragment* of a possible world, but officially the notion is a primitive of the theory. States are taken to be ordered by part-whole ($\sqsubseteq$), and some states $s_1, s_2, \ldots$ are said to be *compatible* if there is a possible state that contains all of them as parts. It is assumed that there is always a smallest state to contain some given states $s_1, s_2, \ldots$, which we call their *fusion* $\bigsqcup\{s_1, s_2, \ldots\} = s_1 \sqcup s_2 \sqcup \ldots$[28] We may recover a notion of a possible world as the notion of a maximal possible state, i.e. a possible state that contains every state it is compatible with.

An exact truthmaker of a proposition is a state that is not only modally sufficient for the truth of the proposition, but also *responsible* for it. Thus, the state of it being sunny in New York is not an exact truthmaker of the proposition that 2+2=4. In addition, to be an exact truthmaker of a proposition, a state must be *wholly relevant* to the truth of the proposition. Thus, the state of it being sunny and cold in New York is not an exact truthmaker of the proposition that it is sunny in New York, since it contains an irrelevant part—the state of it being cold in New York—and therefore fails to be wholly relevant. The condition of being wholly relevant renders exact truthmaking non-monotonic: a given state may exactly verify, i.e. be an exact truthmaker of, a given proposition, without some bigger state also exactly verifying the same propo-

[25] Since AGM is based on an intensional conception of the update, it would seem to follow from this that AGM updates cannot be considered mighty. On the other hand, one might argue that the AGM method of revision does reflect a conception of updates as mighty. The reasoning is this. As will become clearer in the next section, regarding an update as mighty means that for each disjunct P of the update, *absent special reasons to the contrary*, P must be accommodated as a possibility. Now under the AGM account, the agent must accommodate a disjunct P unless they consider no P-worlds to be among the most plausible update-worlds. To the extent that this is a special reason not to accommodate P, AGM revision embodies a view of the update as mighty. Indeed, one might think this is exactly what goes wrong in our puzzle cases. AGM lets us retain $\neg F_1$ upon revising by F_2 only if we have special reasons to discard the F_1-worlds among the F_2-worlds. So in this way, the update F_2 is treated as mighty, and as identical to $(F_1 \wedge F_2) \vee F_2$. But we can be in a situation where it is fine just by default to accept F_2 and retain $\neg F_1$, because F_1 is merely compatible with the update F_2, and not an explicit possibility.

[26] A semantics of this sort was first formulated by Bas van Fraassen [35]. In recent years, the approach and its various applications have been further developed by Fine and others. Fine's [10, 11] offer the best general presentation of the theory. The following brief introduction is indebted to these works.

[27] A formally precise presentation of the framework is given in Appendix A. For many applications of truthmaker semantics—including, I believe, some applications related to belief revision—, it is useful also to allow for a multiplicity of impossible states. For our present concerns, however, impossible states are not essential, though it will be convenient to assume that there is a single impossible state.

[28] When $s_1, s_2, \ldots$ are incompatible, this will be the impossible state.

sition. Relatedly, if a state s is an exact truthmaker of some proposition, then we may conclude that the proposition is in some good sense *about* the whole state s (though not in general *only* about s).[29]

This understanding of truthmaking suggests a particular account of the truthmakers of disjunctions and conjunction: A state makes a disjunction true iff it makes one of the disjuncts true, and it makes a conjunction true iff it is the fusion of truthmakers of the conjuncts.[30,31] Under this account, we can make the required distinction between P and $P \vee (P \wedge Q)$. Any fusion of a truthmaker of P and a truthmaker of Q is a truthmaker of $P \vee (P \wedge Q)$, but since a truthmaker of Q will not in general be relevant to the truth of P, such a fusion will not in general be a truthmaker of P. In particular, we can distinguish between, for example, the logically equivalent F_2 and $F_1 \vee F_2$ in the match example. For by the clause for disjunction, every exact truthmaker of F_1 will be a truthmaker of $F_1 \vee F_2$. But since F_1 by definition has $W_1 \wedge \neg L_1$ as a conjunct, by the clause for conjunction, any such truthmaker will contain a part that makes true $W_1 \wedge \neg L_1$, the proposition that the first match is wet and does not light. That state will be irrelevant to the truth of F_2, and therefore no truthmaker of F_1 will be an exact truthmaker of F_2.[32]

[29]For much more on the relation of (non-monotonic) truthmaking to the notion of aboutness or *subject matter*, see Steve Yablo's [38] and Fine's [11, 13].

[30]There is also an alternative, *inclusive* clause for disjunction, in which the fusion of truthmakers of each disjunct is also considered a truthmaker. In some applications of truthmaker semantics it is preferable to work with the inclusive conception of disjunction, but as we shall see shortly, for the present application there are specific reasons not to do so.

[31]Readers may wonder about the case of negation. The simplest approach is to associate any given proposition with both a set of exact truthmakers, and a set of exact falsitymakers, and to let negation 'flip' the two sets. For now, since none of the AGM postulates involves negation, we may set negation to one side. (Negation does of course play an important role in the relation between AGM-style revision and another important AGM-operation, namely contraction, which corresponds to the mere removal of a belief. There are important questions about the treatment of contraction and similar operations under a truthmaker approach, as well as about the matter of negation, but discussion of these will have to wait for another occasion.)

[32]The central feature of the truthmaker framework is thus its use of a concept of *relevant* truthmaking, which makes it possible to capture various relationships of relevance between propositions. Relatedly, the distinctive features of the truthmaker-based approach to belief revision developed here can also be put in terms of relevance. On the conception of the update as mighty, the update $P \vee (P \wedge Q)$ is relevant to a prior belief in $\neg Q$, whereas the corresponding update P need not be so relevant. In particular, in our example, the update $F_2 \vee (F_1 \wedge F_2)$, but not the update F_2, is regarded as relevant to the belief that $\neg F_1$, and so the agent is permitted to be disposed to give up that belief in processing the former update while not being so disposed with regard to the latter update. The claim that AGM is not appropriately sensitive to the matter of which existing beliefs a given update is relevant to has also been made by earlier authors; most notably by Parikh [28], whose proposal for extending AGM by a relevance axiom has been the subject of extensive discussion and refinements, cf. e.g. [23, 27]. A proper comparison of the present approach with this tradition or other 'relevantist' criticisms of AGM is beyond the scope of this paper, but it may be worth mentioning two significant points of difference. Most of the work in the tradition initiated by Parikh embraces intensionality and accordingly does not adopt a conception of the update as mighty. That tradition also tends to follow a syntactically driven approach to understanding relevance (an exception is [29], providing a system-of-spheres semantics for Parikh's relevance axiom), whereas the present approach is chiefly driven by semantic concepts and considerations. It would be very interesting to study the relation between these approaches more deeply. One might try, for example, to formulate a suitably hyperintensional version of the relevance axiom and investigate whether it may be satisfied under some version of the present approach.

We can now say which truthmaker proposition we take to be the update on a given occasion for revision. First, note that a division between explicit and implicit possibilities can also be made at the level of states. A state is an (at least) implicit possibility if it is compatible with what the agent takes the situation to tell them, and it is an explicit possibility if it also partially constitutes the truth of, i.e. partially *makes* true, what the agent takes the situation to tell them.[33] Among the explicit possibilities, we may then further distinguish between those that *merely partially* make true what the situation tells the agent, and those that *fully* make true what the situation tells the agent. In our example, what I take the situation to tell me is perhaps not exhausted by the claim that either Bob or Bill are coming over. Perhaps I also see that Bob or Bill—whoever it happens to be—is wearing a black sweater and a red hat. Let us call explicit possibilities that fully make true what the situation tells the agent *complete*, and the others *incomplete*. The *truthmaker-update*—short: tm-update—in a given situation, as interpreted by the agent, is then the set of all and only the situation's complete explicit possibilities.[34]

Since the tm-update includes *only* explicit possibilities, a situation with tm-update P tells the agent, for each state $s \in P$, that s *might* obtain, and hence for each disjunct Q of P, that it *might* be the case that Q. Since the tm-update comprises *every* complete explicit possibility, moreover, a situation with update P tells the agent that it *must* be the case that P: the situation is taken by the agent to rule out any scenario in which it is not the case that P. We may summarize the point by saying that tm-updates are both *musty* and *mighty*.

By way of comparison, consider how an intensional conception of the update might be obtained. The obvious answer would seem to be as follows. Given an agent's interpretation of a situation, we divide the possible worlds into two exclusive and exhaustive categories. To the first belong those worlds that are compatible with the situation, under the agent's interpretation, and to the second belong the others. The *possible worlds update*—short: pw-update—is the set of the former worlds. Then pw-updates are certainly also musty: given that every world that is compatible with the situation is included in the update, we can conclude that the situation tells the agent that one of the update-worlds must obtain. But in contrast to tm-updates, which are musty and mighty, pw-updates are merely musty. For as we saw above, intensional updates cannot be mighty in the demanding sense in which tm-updates are.

Note that under our conception of tm-updates, assuming as given two situations with logically equivalent tm-updates P and Q that differ with respect to their truthmakers, there is nothing mysterious about why these situations can be dynamically inequivalent even assuming the agent knows the updates to be logically equivalent. That the agent knows that P and Q are logically equivalent means they know that it

[33] Note that 'partial' here means *part of* rather than *has as part*. Thus, by a partial truthmaker I mean something which is *part of* a truthmaker rather than something which has a truthmaker as a part.

[34] Note that this set is plausibly not closed under fusion. For instance, Ben's coming over and Bob's coming over may each be explicit possibilities without Ben and Bob both coming over being one. That is the reason why I think that in the application to belief revision, we need to allow for truthmaker propositions that fail to be closed under fusion, and relatedly to opt for the non-inclusive clause for disjunction, on which fusions of verifiers of the disjuncts are not automatically verifiers of the disjunction; cf. Footnote 30 above.

is absolutely impossible for P to be true without Q being true as well, and vice versa. A situation with tm-update P is one in which the agent takes themselves to learn that P. Knowing P to be equivalent to Q, they will also conclude that Q. Similarly in a situation with update Q. But *how* the belief that P, or the belief that Q, may appropriately be incorporated in these situations depends also on what the situations tell the agent about what *might* be the case. Given that P and Q have different truthmakers, situations with tm-updates P and Q will differ in this regard, and may therefore differ with respect to their range of rational responses.[35]

9 Revision

Given the proposed conception of updates as sets of truthmakers, how can we characterize the rationally permissible ways to revise a belief system by an update? First, we need to decide how to model belief systems within our revised setting. Although the issue calls for extended discussion, for present purposes we may adopt a policy of keeping this as simple as possible, and of minimizing deviation from the AGM approach, so that we may see how much, or how little, of that approach we are forced to give up to accommodate the problem cases. We shall therefore continue to model a belief state by the set of possible worlds at which it is true. Thus, the update will be the only source of hyperintensionality under the resulting approach.[36]

In imposing rationality constraints on doxastic states, we follow a similar strategy as the possible worlds approach in that we demand that the revision function be definable in a certain way. We suggested above that Matt might plausibly be seen to revise by disjunctions by disjoining revisions by the disjuncts. Within the truthmaker framework, a disjunct of an update is any subset of the update, and the disjuncts of the update which are not themselves disjunctive are the subsets with exactly one truthmaker as member. So the suggestion is, in effect, to take the revision by an update to be the disjunction of the revisions by the individual truthmakers of the update. In this way, we obtain what we called the wayward revision of a belief system by an update.

Under certain circumstances, however, it may be rationally permissible for an agent to deviate from the method of wayward revision. The idea is that one may take a situation to tell one that it might be that P, and at the same reasonably hold that one knows better, as it were—that information one possesses independently of the given occasion for revision, and that is not undermined by the new information obtained, may justify one in continuing to exclude the possibility that P, even if the situation

[35]We might compare the situation to the one in approaches to revision using belief bases, which are sets of sentences not (normally) closed under logical consequence. There, a distinction is made between, roughly speaking, sentences an agent believes to be true purely because they follow logically from other sentences the agent believes and sentences an agent believes to be true on (partly) independent grounds. The view is that rational revision is sensitive to this difference, and different but logically equivalent belief bases may rationally be revised differently. Just as in our case, the view is fully compatible with a view of agents as logically omniscient. See e.g. [18, pp. 17ff].

[36]That being said, I suspect that an ultimately more satisfactory approach may be obtained by also embracing hyperintensionality with respect to the belief system and representing an agent's beliefs by their exact truthmakers rather than all the verifying worlds.

on its own is taken to explicitly present P as a possibility. First of all, one might so interpret a situation as to assign it the update $P \vee Q$, where P but not Q is compatible with one's previous beliefs. In the Bob-and-Bill case, for example, I might take the situation to tell me that Bob might be coming (P) and Bill might be coming (Q), when my beliefs are compatible with the former but not the latter possibility. In such a case, it is permissible for me to disregard the revision by Q and simply select as my new belief system the revision by P.[37] Second of all, even if all disjuncts of the update are incompatible with the agent's current beliefs, those beliefs may exclude the revisions by some disjuncts much more firmly than others, and it may then be rational for the agent to disregard the latter. In the context of the dominos, a plausible example might be the update to the effect that either all the stones fell, or exactly the odd-numbered stones fell. Given the setup of the case, any world verifying the second disjunct might seem a so much more remote possibility than worlds verifying the first disjunct that it may justifiably be disregarded. This suggests a modification of the simple method of revision, whereby revisions by disjunctive updates are constructed by first forming the disjunction of the revisions by each disjunct, and then applying a "plausibility filter", discarding those disjuncts that are regarded as sufficiently less plausible than others. Just like we did under the possible worlds approach, therefore, we may appeal to a plausibility ordering of the worlds, and let $B * P$ comprise only the most plausible worlds in the wayward revision of B by P.

It needs to be emphasized, however, that while from a formal perspective the plausibility orderings used here are just like those used in AGM, their representational role is quite different, and much less central to the overall account. In particular, under the present approach, an agent may consider two initially excluded worlds equally plausible and yet, after a rational revision, continue to exclude one of them, while no longer excluding the other. Indeed, as will become clearer below, this is exactly what allows us to deal in an intuitively satisfactory way with the puzzle cases.

It remains to characterize the rationally acceptable ways to revise a belief state by a single truthmaker. A natural idea is to once more take a leaf out of Fine's semantics for counterfactuals (cf. [7, pp. 236ff]), and to postulate a *transition* relation that encodes, roughly speaking, how each of the various worlds in the belief state may be adjusted upon revision by any given input state.[38] We write $s \rightarrow_b w$ to say that world w is a revision of world b by state s, and define the *wayward* revision $B \circ P$ of belief state B by update P as $\{w : p \rightarrow_b w \text{ for some } p \in P \text{ and } b \in B\}$. The final revision is then obtained by applying the plausibility filter. Where X is a set of worlds, we

[37] Indeed, it is standardly assumed that this is not only permissible but mandatory. Specifically, the AGM postulate of *Vacuity* demands that no beliefs be given up in incorporating information compatible with the agent's current beliefs.

[38] Although in its use of a transition relation, the present approach thus maintains a strong parallel to Fine's semantics for counterfactuals, it should be noted that there is no counterpart in the latter to our use of plausibility orderings. Roughly speaking, while I propose to divide the work done by plausibility orderings under the possible worlds approach between plausibility orderings and a transition relation, Fine proposes to let transition do all the work of the similarity ordering in the possible worlds analysis of counterfactuals. I suspect that by using a similarity ordering in the account of counterfactuals, much as we use a plausibility ordering here, we might be able to avoid the difficulties for Fine's semantics raised by Embry [4].

let $g(X)$ be the set of the maximally plausible members of X, and define $B * P$ as $g(B \circ P)$.[39]

Note that the revision operations of the usual possible worlds account constitute a special case of our revision operations, which corresponds to the condition that a world w is a revision of another world b by a consistent state s iff w contains s as part. Then the wayward revision of any belief state is simply the set of worlds at which the update is true, and the final, filtered revision is the set of the maximally plausible update-worlds. Thus, the way our present account improves on the possible worlds account, is by allowing the transition relation to narrow our focus from the start on some subset of the update-worlds, and to do so in a way sensitive to the exact truthmakers of the update.

To illustrate the idea, we sketch a truthmaker model of Dom's doxastic state in the dominos example.[40] For simplicity, we let our worlds be built up purely from states of the form f_n—stone n falls—and $\overline{f_n}$—stone n does not fall. Then Dom's initial belief has just one verifier, the state $b = \bigsqcup\{\overline{f_n} : n \in N\}$. Its revision by the proposition that F_2, with its sole verifier f_2, will comprise exactly the maximally plausible worlds w with $f_2 \rightarrow_b w$. Its revision by the proposition that $(F_1 \wedge F_2) \vee F_2$, with is two verifiers f_2 and $f_1 \sqcup f_2$, will comprise exactly the maximally plausible worlds w with either $f_2 \rightarrow_b w$ or $f_1 \sqcup f_2 \rightarrow_b w$. We capture the fact that Dom takes the falling of any stone to lead to the falling of every subsequent stone by letting $f_1 \sqcup f_2 \rightarrow_b w$ hold if and only if $w = \bigsqcup\{f_n : n \in N\}$. The fact that Dom takes the falling of any stone not to support the falling of any previous stone is captured by letting $f_2 \rightarrow_b w$ hold if and only if $w = \overline{f_1} \sqcup \bigsqcup\{f_n : n \geq 2\}$. More generally, we capture Dom's dispositions to respond to a proposition of the form F_n by letting $f_n \rightarrow_b w$ hold iff $w = \bigsqcup\{\overline{f_m} : m < n\} \sqcup \bigsqcup\{f_m : m \geq n\}$. To accommmodate the fact that Dom is disposed to make room for the possibility that F_1 upon learning that F, or learning that $(F_1 \wedge F_2) \vee F_2$, we may stipulate that all *regular* worlds other than b are equally plausible, where a world is regular iff it is of the form $\bigsqcup\{\overline{f_m} : m < n\} \sqcup \bigsqcup\{f_m : m \geq n\}$ for some n.

Thus, in revising by F_2, the world with all stones falling is excluded. But it is *not* excluded because it is less *plausible* than the other F_2-worlds. Instead, it does not even come up for consideration at the stage at which the plausiblity filter is applied, because it is not among the worlds that are revisions of b by f_2. Why is it not among those worlds? Because the state f_2 of the second stone falling is taken by Dom to provide no grounds for replacing the state $\overline{f_1}$ of the first stone standing by f_1. Such

[39] It is worth mentioning here that by appealing to the mereological as well as modal profile of states, it is possible to define in logical terms certain *defaults* for transition and plausibility. For instance, we might say that by default, a world w transitions to another v upon revision by a state s iff v is maximal among the s-containing worlds with respect to its mereological overlap with w. For plausibility, a natural default is to take all worlds incompatible with the current belief system to be equally plausible. In this way, the truthmaker approach allows us to give a purely logical characterization of a non-trivial operation of belief revision. This idea will be studied in detail in the previously mentioned joint work.

[40] A proper definition of such a state, and a proof that it satisfies the assumptions of the case as well as the constraints proposed below, is given in Appendix C.

a reason to change the relevant state obtains only for the other, later stones in the sequence.

So we do not wish to hold that a world w transitions to another v upon revision by s whenever v contains s—this would render our account intensional, and equivalent to AGM. But there are a number of weaker constraints that we may plausibly impose. In particular, for any consistent state s and any world $b \in B$, we shall require that[41]

1. there is some world w with $s \rightarrow_b w$,
2. if $s \rightarrow_b w$ then w is a world,
3. if $s \rightarrow_b w$ then $s \sqsubseteq w$,
4. if $s \sqsubseteq b$ then $s \rightarrow_b b$,
5. if $s \rightarrow_b w$ and $r \sqsubseteq w$, then $s \sqcup r \rightarrow_b w$
6. if $s \sqcup t \rightarrow_b w$, then $s \rightarrow_b v$ for some $v \leq w$

These constraints, together with the familiar assumptions about plausibility orderings, ensure that filtered revision satisfies natural counterparts of all the basic AGM postulates *except* for *Intensionality*. They also ensure that under the natural interpretation of $\Rightarrow$ in terms of filtered revision, all of Fine's rules from Section 5 are valid with the exception of the intensionalist rule of *Substitution*.[42]

The situation is more complicated with respect to the postulates of *Superexpansion* and *Subexpansion*. These are usually stated in a form in which they relate revisions by conjunctions to revisions by their conjuncts. *Superexpansion* then says that $(B * P) \wedge Q$ entails $B * (P \wedge Q)$, and *Subexpansion* says that if $B * P$ is compatible with Q, the converse entailment also holds, so that $B * (P \wedge Q)$ entails $(B * P) \wedge Q$. Now as we have noted before, within an intensional framework, the relation between a conjunction and its conjuncts is simply the relation between a proposition and a proposition entailed by it, and thus the same as the relation between a proposition and a disjunction in which it is a disjunct. As a result, we can also formulate versions of *Superexpansion* and *Subexpansion* that relate revisions by disjunctions to revisions by their disjuncts. These versions will be equivalent to the usual ones under the assumption of intensionally individuated updates, but they will not be equivalent within our hyperintensional framework. We may thus distinguish between the following four principles:

Superexpansion$(\wedge)$	$(B * P) \wedge Q$ entails $B * (P \wedge Q)$
Subexpansion$(\wedge)$	$B * (P \wedge Q)$ entails $(B * P) \wedge Q$, if $B * P$ and Q are compatible
Superexpansion$(\vee)$	$(B * (P \vee Q)) \wedge P$ entails $B * P$
Subexpansion$(\vee)$	$B * P$ entails $(B * (P \vee Q)) \wedge P$, if $B * (P \vee Q)$ and P are compatible

[41] Most of these constraints are similar or identical to ones that Fine imposes on transition relations in his semantics for counterfactuals; cf. [7, pp. 239ff].

[42] Proofs, here and below, are again delegated to Appendix A.—It is worth mentioning that for the basic AGM postulates, it is sufficient to impose conditions (1)–(4). (5) and (6) are only required for the versions of the supplementary postulates given below, and for Fine's rule of *Transitivity*.

It turns out that conditions (1)–(5) on transition relations, together with the conditions on plausibility orderings, ensure that *Superexpansion(*$\wedge$*)* and *Subexpansion(*$\vee$*)* are satisfied. *Superexpansion(*$\vee$*)* and *Subexpansion(*$\wedge$*)* are not in general satisfied.

This is a good thing, though. For as I show in Appendix B, there is no way to do so, given the other principles and constraints, without the account collapsing again into AGM and thereby validating *Intensionality*. Moreover, we can construct compelling counter-examples to these postulates on the basis of our example cases. For simplicity, consider the dominos case again. For *Superexpansion(*$\vee$*)*, let P be the proposition that F_2 and Q the proposition that $F_1 \wedge F_2$. Then as we have argued, it is permissible for $B * P$ to rule out that F_1 while $B * \big(P \vee Q\big)$—and then also $\big(B * \big(P \vee Q\big)\big) \wedge P$—does not, and therefore fails to entail $B * P$, in violation of *Superexpansion(*$\vee$*)*. For *Subexpansion(*$\wedge$*)*, let P be as before and let Q be the proposition that $\big(F_1 \wedge F_2\big) \vee F_2$. So P says that the second stone fell, and Q says that the second, or the first and the second stone fell. $B * P$ then says that the first stone stands, but the second stone and all subsequent ones fell. This is of course compatible with Q; indeed, it entails Q. $P \wedge Q$ is equivalent, even in terms of its truthmakers, to Q. So $B * \big(P \wedge Q\big) = B * Q$. But given our assumptions, $B * Q$ makes room for the possibility that all stones fell, and so it cannot entail $B * P$, which does not allow for that possibiliby. A fortiori, $B * Q$ then does not entail $\big(B * P\big) \wedge Q$, in violation of *Subexpansion(*$\wedge$*)*.

10 The Advantages of Mightiness

The results of the previous sections show that a viable, hyperintensional theory of rational belief revision can be developed within the framework of truthmaker semantics and on the basis of a conception of the update as mighty. Moreover, we saw that this kind of approach allows us to give a very natural account of what is going on in our puzzle cases, which is much more in line with an intuitive assessment of these cases than any account that could be given within an intensional framework. In this final section of the paper, I want to briefly indicate at a more general and abstract level some of the further advantages of the proposed approach and in particular the use of mighty updates.

The central requirement on a conception of the update is that dynamically inequivalent situations always be assigned distinct updates. As we have seen, there are logically equivalent tm-updates that represent dynamically inequivalent situations. At first glance, if the tm-updates associated with a pair of dynamically inequivalent situations are logically equivalent, it would seem that the pw-updates associated with those situations must be identical. This would show that pw-updates are plainly incapable of capturing the relevant features of occasions for revision. So the question arises how, if at all, intensionalists can avoid this conclusion.

It will be useful to consider a concrete example. Suppose I have hurt my ankle playing football. I take it to be nothing serious but go to the doctor just in case. After examining me, she tells me: ‘Your ankle is sprained, or sprained and broken’. I trust the doctor and see no reason to suspect her to try to mislead me. So I take the situation to tell me that my ankle must be sprained, and that it might in addition be broken. It is then reasonable for me to give up my belief that my ankle is not broken.[43] Now consider a version of the situation in which the doctor tells me simply: 'Your ankle is sprained'. Again, I trust the doctor and see no reason to suspect her to be anything less than fully perspicuous in sharing her opinion of my ankle. So I take the situation to tell me that my ankle is sprained, and I do not take it to tell me that my ankle might be broken. It is then reasonable for me to retain my belief that my ankle is not broken.

Clearly, we have a pair of dynamically inequivalent occasions for revision. Moreover, it seems plausible that under the specified interpretations of the situations, they are to be assigned logically equivalent tm-updates. The tm-update in the first situation—call it the *sprained/broken* scenario—might plausibly be taken to be the truthmaker proposition that my ankle is sprained, or sprained and broken. In the second situation—call it the *sprained* scenario—the tm-update is plausibly taken to be the truthmaker proposition that my ankle is sprained. These propositions, of course, are logically equivalent. Now if the corresponding pw-updates are simply the sets of worlds in which these tm-updates are true, then the two situations are assigned the same pw-update, in spite of their dynamic inequivalence. Can the intensionalist plausibly deny the claim that these are the pw-updates?

A natural idea is to point out that the update is supposed to capture the *total* information received by the agent, and that the updates we specified do not satisfy this condition. For example, in the first scenario, I presumably also obtain the information that the doctor assertorically utters the sentence 'Your ankle is sprained, or sprained and broken', and perhaps I obtain the information that the doctor is not convinved that my ankle is not broken. And in the second scenario, I obtain the information that the doctor utters 'Your ankle is sprained' instead, and perhaps take the situation to also tell me that the doctor confidently rules out that my ankle is broken. If we enrich the updates given above by these further bits of information, then the updates assigned to the two situations will not be logically equivalent.[44]

In order to properly evaluate this response, we need to get clearer about the requirement that the update represent the total information received by the agent in the situation under consideration. On the one hand, it is uncontroversial that we need some form of such a completeness requirement: we simply cannot determine the rational responses to a situation purely on the basis of the fact that *part* of what the agent learns is that P, without being told what else the agent learns. On the other

[43] Note that nothing I have said about this scenario depends on it being part of the semantic content of the doctor's utterance that my ankle might be broken. It is perfectly consistent with what I say that this is merely a pragmatic implication. But since I trust the doctor and assume that she is not trying to mislead me, I take on board not only the semantic but also the pragmatic implications of what she says.

[44] While this seems to be the most natural response, it is perhaps not the only possible response. A more comprehensive and detailed examination of the options available here is beyond the scope of this paper, but let me mention one alternative strategy, hinted at by Wolfgang Spohn [32, Section 6], when discussing a somewhat similar example. The idea is to maintain that in the situation in question, the appropriate response by the agent consists not simply in a revision by some given update, but in a sequence of belief change operations, first simply removing my previous beliefs about the health of my ankle, including the belief that my ankle is not broken, and then revising with the proposition that my ankle is sprained. This suggestion may yield the right results in our example, but absent plausible general principles telling us what situations call for what combinations of operations, the response appears objectionably ad hoc.

hand, a naïvely strict interpretation of the completeness requirement gives rise to severe methodological difficulties. For under such an interpretation, in more or less any realistic situation a doxastic agent might find themselves in, the total information received will be unmanageably rich and complex. For a start, as long as the agent has their eyes open, they would seem to receive, at any point in time, a very rich body of visual information that it is not even feasible to express in words. So if the update needs to capture the total information received in this very demanding sense, we lose the ability to test any proposed theory of belief revision by applying it to realistic scenarios and working out its implications.

In practice, belief revision theorists do not attempt to specify anything like an update that would be complete in this demanding sense. Nor, it might be added, do they normally attempt to fully specify anything like a realistic complete initial belief state that is to be revised, or a complete revised belief state. How is this practice to be justified? To a rough approximation, a natural idea is as follows. First of all, in considering examples, we usually limit attention to the evolution of a certain subset of an agent's beliefs, such as their beliefs concerning the status of certain domino stones or matches, or the whereabouts of their neighbour's twins. We specify those initial beliefs, and tacitly stipulate that in the kind of situation to be considered, any other beliefs the agent might have are irrelevant to how the subset we are considering can rationally be revised. With regard to the update, a related policy is in place: the update is assumed to be complete in the sense of encoding all the information received that is relevant to how the part of the agent's belief system under consideration may rationally be revised. What the example of my injured ankle helps bring out is that the truthmaker approach and the intensional AGM approach differ greatly with respect to how, and how easily, the demands of relevant completeness may be met.

Under the truthmaker approach, we can adequately model the example by specifying my initial beliefs about the health of my ankle, and by taking the updates in the two scenarios to be as described above—that my ankle is sprained in the *sprained* scenario, and that my ankle is sprained, or sprained and broken in the *sprained/broken* scenario. Given the assumptions of the example, there seems to be no reason to take these updates to be relevantly incomplete. Under the possible worlds approach, we need to work with a much more complicated model of the situation. In order to capture all the relevant differences about the information received, we have to incorporate in the updates information about which sentences were uttered, or perhaps about which beliefs the doctor holds or does not hold concerning my ankle. To make room for the fact that I can reasonably give up the belief that my ankle is not broken in the *sprained/broken* scenario while retaining the same belief in the *sprained* scenario, we should then say that, roughly speaking, I consider worlds in which the doctor's diagnosis is correct to be more plausible than ones in which it is mistaken.

At least in this kind of case, the truthmaker approach thus affords a simpler, more direct, and more elegant representation of the case. But more significantly, it also allows us to straightforwardly capture intuitive rational constraints that cannot be captured under the alternative, more complicated model. To see this, note that how I am disposed to revise in the *sprained* scenario imposes constraints on how I may rationally be disposed to revise in the *sprained/broken* scenario. In particular, it seems that my beliefs about the health of my ankle should be strictly weaker in the sprained/broken scenario than in the sprained scenario. Under the truthmaker approach, this constraint follows given the logical relationship between the associated tm-updates. For these constitute a pair of the form P and $P \vee (P \wedge Q)$, and we can show using the principles of *Success*, *Consistency* and *Subexpansion(*$\vee$*)* that $B * P$ entails $B * (P \vee (P \wedge Q))$ whenever P is consistent.[45] But it is hard to see how a similar result could be obtained on the basis of a representation of the situations in terms of the associated pw-updates.

Appendix A: Doxastic States in a Truthmaker Framework

The basic structure of truthmaker semantics is that of a *state-space*, which is a special kind of a partially ordered set. Recall that a partial order on a set S is a two-place relation $\sqsubseteq$ that is reflexive—$s \sqsubseteq s$ for all $s \in S$—, transitive—$s \sqsubseteq t$ and $t \sqsubseteq u$ implies $s \sqsubseteq u$ for all $s, t, u \in S$—, and anti-symmetric—$s \sqsubseteq t$ and $t \sqsubseteq s$ implies $s = t$ for all $s, t \in S$. Call a partial order $\sqsubseteq$ on a set S *complete* iff for every subset T of S, a least upper bound exists, i.e. there is an element $s \in S$ such that $t \sqsubseteq s$ for all $t \in T$, and $s \sqsubseteq u$ whenever $t \sqsubseteq u$ for all $t \in S$. By designating a certain subset of the states as the set of *possible* or *consistent* states, we obtain a modalized state-space.

Definition 1 A *modalized state-space* is a triple $(S, S^{\Diamond}, \sqsubseteq)$ such that

1. S is a non-empty set,
2. $\sqsubseteq$ is a complete partial order on S, and
3. $S^{\Diamond}$ is a non-empty subset of S such that $s \in S^{\Diamond}$ whenever $s \sqsubseteq t$ and $t \in S^{\Diamond}$.

Informally, the members of S are the states, $\sqsubseteq$ is the parthood-relation, and $S^{\Diamond}$ is the set of possible, or consistent states. For $T \subseteq S$, we write $\bigsqcup T$ for the least upper bound of T, which we also call the *fusion* of the members of T, and often write as $t_1 \sqcup t_2 \sqcup \ldots$ when $T = \{t_1, t_2, \ldots\}$. We call a *(possible) world* any possible state that contains every state it is compatible with, and we denote their set by S^w. A modalized state-space is called a *W-space* iff every possible state is part of a possible world, and it is called *topsy* if it contains only one impossible state. This will then be the one state, written $\blacksquare$, which contains every state. Throughout this section and the next, we shall be working within some fixed, topsy W-space $\mathcal{S} = (S, S^{\Diamond}, \sqsubseteq)$.

A *proposition* P is any non-empty subset of S. A proposition is consistent iff one of its members is, and propositions P and Q are compatible iff some member of P is compatible with some member of Q. The conjunction $P \wedge Q$ of propositions P and Q is $\{p \sqcup q : p \in P \text{ and } q \in Q\}$. Note that this is non-empty even if P and Q are incompatible, in which case $P \wedge Q$ is $\{\blacksquare\}$.[46] The disjunction $P \vee Q$ is $P \cup Q$.

[45] By *Success*, $B * (P \vee (P \wedge Q))$ entails $P \vee (P \wedge Q)$, and hence P. If P is consistent, so is $P \vee (P \wedge Q)$ and thus by *Consistency*, $B * (P \vee (P \wedge Q))$, so $B * (P \vee (P \wedge Q))$ is compatible with P. By *Subexpansion(*$\vee$*)*, it then follows that $B * P$ entails $B * (P \vee (P \wedge Q))$.

[46] The main reason for countenancing the impossible $\blacksquare$, in the present context, is to avoid technical inconveniences such as having to allow for empty propositions.

A proposition P is said to *(loosely*[47]*) entail* ($\models$) a proposition Q iff every world containing a truthmaker of P contains a truthmaker of Q .

We now turn to the task of defining the class of permissible doxastic states. We do this using a notion of a *coherent pair* of a *plausibility* ordering and a *transition* relation.

Definition 2 A *plausibility ordering* is a two-place relation $\leq$ on $S^w \cup \{\blacksquare\}$ satisfying the following conditions, where $g(X)$:=$\{w \in X : w \leq v$ for all $v \in X\}$ for all $X \subseteq S^w \cup \{\blacksquare\}$:

(P-Connectedness)	$w \leq v$ or $v \leq w$
(P-Transitivity)	$w \leq v$ and $v \leq u$ implies $w \leq u$
(P-Limit)	$g(X) \neq \emptyset$ if $\emptyset \subset X \subseteq S^w \cup \{\blacksquare\}$
(P-Inconsistency)	$\blacksquare \leq s$ implies $s = \blacksquare$

These are exactly the conditions imposed under the possible worlds approach, except for the added clause dealing with $\blacksquare$. Given any plausibility ordering, we often use B to refer to $g(S^w)$, since this is the set of worlds at which the agent's beliefs are true.

Definition 3 A *transition relation* is a three-place relation $\rightarrow$ on S subject to the following conditions:

(T-Success)	$s \rightarrow_u t$ implies $t \sqsupseteq s$
(T-Completeness)	if $u \in S^w$ and $s \rightarrow_u t$ then $t \in S^w \cup \{\blacksquare\}$
(T-Consistency)	if $s,\ u \in S^\Diamond$ then $s \rightarrow_u t$ for some $t \in S^\Diamond$
(T-Vacuity)	if $s \sqsubseteq u$ then $s \rightarrow_u u$
(T-Incorporation)	if $s \rightarrow_u t$ and $r \sqsubseteq t$ then $s \sqcup r \rightarrow_u t$

Definition 4 Let $\leq$ and $\rightarrow$ be a pair of plausibility ordering and transition relation. The operations of *wayward revision* $\circ$ and *(filtered) revision* $*$ induced by $\leq$ and $\rightarrow$ are defined as follows, for P a non-empty subset of S:

$$B \circ P := \{t \in S : p \rightarrow_b t \text{ for some } p \in P \text{ and } b \in B\}$$
$$B * P := g(B \circ P)$$

Definition 5 A pair of a plausibility ordering $\leq$ and a transition relation $\rightarrow$ is *coherent* iff, whenever $b \in B$:

(PT-Existence)	$s \rightarrow_b t$ for some $t \in S$
(PT-Link)	if $w \in B \circ \{s \sqcup t\}$ then $v \leq w$ for some $v \in B \circ \{s\}$

[47] There are several other, narrower relations of entailment that may be defined within the truthmaker framework. For present purposes, however, we may confine attention to this one.

Theorem 1 *Let $*$ be the revision function induced by some coherent pair of plausibility ordering and transition relation. Then*

(*R-Success*)	$B * P \models P$
(*R-Vacuity*)	$B * P \models B$ *if B is compatible with P*
(*R-Inclusion*)	$B \wedge P \models B * P$
(*R-Consistency*)	$B * P$ *is consistent if P is*
(*R-Superexpansion*$(\wedge)$)	$(B * P) \wedge Q \models B * (P \wedge Q)$
(*R-Subexpansion*$(\vee)$)	$B * P \models (B * (P \vee Q)) \wedge P$ *if* $B * (P \vee Q)$ *is compat. w. P*

Proof Note that $B * P \subseteq S^w \cup \{\blacksquare\}$ for all non-empty P. So to establish that any revision entails some proposition Q, we need to show that every truthmaker of the revision contains a truthmaker of Q as part.

(R-Success): If $s \in B * P$, then $p \rightarrow_b s$ for some $p \in P$ and $b \in B$, and by (T-Success), $s \sqsupseteq p$.

(R-Vacuity): Suppose B is compatible with P, and let $b \in B$ be compatible with $p \in P$. Since b is a world, it follows that $b \sqsupseteq p$, so by (T-Vacuity) $p \rightarrow_b b$, and hence $b \in B \circ P$. Since $b \in g(S^w)$, it follows that $b \leq v$ for all $v \in B \circ P$ and hence $b \in B * P$. But then $v \leq b$ for all $v \in B * P$, so $B * P \subseteq B$ and hence $B * P \models B$.

(R-Inclusion): Note that by (PT-Existence), $B * P$ is non-empty. Suppose $s \in B \wedge P$, and let $b \in B$ and $p \in P$ be such that $s = b \sqcup p$. If $s = \blacksquare$, s contains every state, and hence some verifier of $B * P$. If s is consistent, then b is compatible with p, so since $b \in S^w$, $b \sqsupseteq p$, and thus $s = b$. By (T-Vacuity), $p \rightarrow_b b$, $b \in B \circ P$. Since $b \in B = g(S^w)$, $b \leq w$ for all $w \in B \circ P$, hence $s = b \in B * P$.

(R-Consistency): Suppose $s \in P$ is consistent and let $b \in B$. By (T-Consistency), $v \in S^w$ for some v with $s \rightarrow_b v$, so $B \circ P$ has a consistent member. By (P-Inconsistency), it follows that $g(B \circ P)$ has some (indeed, only) consistent members.

(R-Superexpansion($\wedge$)): Note first that by (PT-Existence), $B * (P \wedge Q)$ is non-empty. Now suppose $s \in (B * P) \wedge Q$. Then $s = t \sqcup q$ for some $t \in B * P$ and $q \in Q$. If $s = \blacksquare$, s contains every state, and so some verifier of $B * (P \wedge Q)$. If s is consistent, then t is a possible world, and hence contains q as part. Since $t \in B * P$, $p \rightarrow_b t$ for some $p \in P$ and $b \in B$. By (T-Incorporation), also $p \sqcup q \rightarrow_b t$, and hence $t \in B \circ (P \wedge Q)$. Now consider any $v \in B \circ (P \wedge Q)$. Let $p' \in P$ and $q' \in Q$ be such that $v \in B \circ \{p' \sqcup q'\}$. Then by (PT-Link) there is some $u \leq v$ with $u \in B \circ \{p\}$ and hence $u \in B \circ P$. Since $t \in B * P$, it follows that $t \leq u$ and hence by (P-Transitivity) $t \leq v$. So $t \in g(B \circ (P \wedge Q)) = B * (P \wedge Q)$, as desired.

(R-Subexpansion($\vee$)): Suppose $B * (P \vee Q)$ is compatible with P. Suppose $s \in B * P$. As before, the case in which s is inconsistent is easy. Suppose instead $s = w \in S^w$ and $p \rightarrow_b w$ with $p \in P$, and $b \in B$. We wish to show that $w \in (B * (P \vee Q)) \wedge P$. Since $w \sqsupseteq p$ and $p \in P$, it suffices to show that $w \in B * (P \vee Q)$. Since $w \in B * P$, $w \in B \circ P$ and hence $w \in B \circ (P \vee Q)$. It remains to show that $w \leq v$ for all $v \in B \circ (P \vee Q)$. Now note that since $B * (P \vee Q)$ is compatible with P, there is some $w' \in g(B \circ (P \vee Q))$ with $w' \leq v$ for all $v \in B \circ (P \vee Q)$, and $w' \sqsupseteq p'$

for some $p' \in P$. It then suffices to show that $w \leq w'$. Now either $w' \in B \circ P$ or $w' \in B \circ Q$. If $w' \in B \circ P$, then $w \leq w'$ is immediate from $w \in g(B \circ P)$. So suppose $w' \in B \circ Q$, so $q \rightarrow_b w'$ for some $q \in Q$. Since $w' \sqsupseteq p'$, by (T-Incorporation), $p' \sqcup q \rightarrow_b w'$. Then by (PT-Link), $u \in B \circ \{p'\}$, so $u \in B \circ P$, and hence $w \leq u$. By (P-Transitivity), $w \leq w'$, as desired. □

(R-Success), (R-Vacuity), (R-Inclusion), and (R-Consistency) are the obvious counterparts in our (semantic) setting to the (syntactically formulated) AGM postulates of *Success*, *Vacuity*, *Inclusion* and *Consistency*. The postulate of *Closure* serves mainly to ensure intensionality with respect to belief states, which is guaranteed under our account by the identification of belief states with the set of possible worlds at which they are true. The *Intensionality* postulate, of course, does not hold. Within our semantic setting, the only valid version of this principle is the triviality that $B * P = B * Q$ if $P = Q$. Under a syntactic formulation of the theory, though, we would have the non-trivial principle that $K * \alpha = K * \beta$ if α and β are *exactly equivalent*, i.e. have the same exact truthmakers.[48]

Moreover, as expected, all Finean rules from Section 5 except for the intensionalist rule of *Substitution* are valid under the obvious interpretation of $\Rightarrow$.

Theorem 2 *Let $*$ be the revision function induced by some coherent pair of plausibility ordering and transition relation. For any propositions P, Q, let $P \Rightarrow Q$ hold iff $B * P \models Q$. Then*

$$
\begin{array}{ll}
(\textit{R-Entailment}) & P \Rightarrow Q \text{ whenever } P \models Q \\
(\textit{R-Transitivity}) & \textit{If } P \Rightarrow Q \text{ and } P \wedge Q \Rightarrow R \text{ then } P \Rightarrow R \\
(\textit{R-Conjunction}) & \textit{If } P \Rightarrow Q \text{ and } P \Rightarrow R \text{ then } P \Rightarrow Q \wedge R \\
(\textit{R-Disjunction}) & \textit{If } P \Rightarrow R \text{ and } Q \Rightarrow R \text{ then } P \vee Q \Rightarrow R
\end{array}
$$

Proof (R-Entailment) and (R-Conjunction) are immediate from the definition of $\Rightarrow$ and the fact that the $\models$-consequences of a proposition are closed under conjunction. (R-Disjunction) is immediate from the observation that $g(X \cup Y) \subseteq g(X) \cup g(Y)$.

(R-Transitivity): Assume $B * P \models Q$ and $B * (P \wedge Q) \models R$. If P is inconsistent, $P \Rightarrow R$ follows immediately given (Success). So suppose P is consistent. Then $B * P \subseteq S^w$. So let $w \in B * P$, and suppose $p \rightarrow_b w$ with $p \in P$ and $b \in B$. We need to show that $w \sqsupseteq r$ for some $r \in R$. Since $B * P \models Q$, we have $w \sqsupseteq q$ for some $q \in Q$. Then $p \sqcup q$ is consistent. By (T-Incorporation), $p \sqcup q \rightarrow_b w$, so $w \in B \circ (P \wedge Q)$. Now let $v \in B \circ (P \wedge Q)$, and let $p' \in P$ and $q' \in Q$ be such that $v \in B \circ \{p' \sqcup q'\}$. By (PT-Link), $u \in B \circ \{p'\}$ and hence $u \in B \circ P$ for some $u \leq v$. But since $w \in B * P$, $w \leq u$ and hence $w \leq v$. So $w \in B * (P \wedge Q)$. Since $B * (P \wedge Q) \models R$, $w \sqsupseteq r$ for some $r \in R$, as desired. □

[48] On the logic of this equivalence relation, see [3, 9, 24].

Appendix B: Expansion and Collapse

We now show that the validity of *Subexpansion(*$\wedge$*)* or *Superexpansion(*$\vee$*)* would collapse our account into intensional AGM. Some additional notation will be helpful. For any proposition Q, let Q^w be the set of Q-worlds, i.e. $\{w \in S^w : w \sqsupseteq q$ for some $q \in Q\}$. The AGM revision of B by P is then simply $g(P^w)$. Now consider the following constraint on the connection between plausibility orderings and transition relations:

$$\big(\textit{PT-Expansion}\big) \quad g\big(\{p\}^w\big) \subseteq B \circ \{p\}$$

It turns out that whenever (PT-Expansion) is satisfied, the resulting revision function is an AGM revision function.[49]

Proposition 1 *Let* $\leq$ *and* $\rightarrow$ *be a coherent pair of plausibility ordering and transition relation that satisfy (PT-Expansion). Then* $B * P = g\big(P^w\big)$ *whenever* P *is consistent.*

Proof Let P be consistent. Then $B * P = B * \big(P \cap S^{\Diamond}\big)$ and $g\big(P^w\big) = g\big(\big(P \cap S^{\Diamond}\big)^w\big)$, so we may assume without loss of generality that $P \subseteq S^{\Diamond}$.

Suppose first that $w \in g\big(P^w\big)$. Let $p \in P$ be such that $w \sqsupseteq p$. Then $w \in g\big(\{p\}^w\big)$, so by (PT-Expansion), $w \in B \circ \{p\}$ and hence $w \in B \circ P$. Suppose $v \in B \circ P$. Then $v \in P^w$, so since $w \in g\big(P^w\big)$, $w \leq v$. It follows that $w \in B * P$.

Suppose now that $w \in B * P$, so $w \in g\big(B \circ P\big)$. Then $w \in P^w$. Now let $v \in P^w$. Pick a world $u \in g\big(P^w\big)$, so $u \leq v$, and let $p' \in P$ be such that $u \sqsupseteq p'$ and hence $u \in g\big(\{p\}^w\big)$. By (PT-Expansion), $u \in B \circ \{p\}$ and hence $u \in B \circ P$. Since $w \in g\big(B \circ P\big)$, it follows that $w \leq u$ and hence $w \leq v$. So $w \in g\big(P^w\big)$, as desired. □

Any violation of (PT-Expansion), however, yields a violation of both *Subexpansion(*$\wedge$*)* and *Superexpansion(*$\vee$*)*.

Proposition 2 *Let* $\leq$ *and* $\rightarrow$ *be a coherent pair of plausibility ordering and transition relation. Let* $w \in S^w$, $p \sqsubseteq w$, *and* $w \leq v$ *for all* $v \in B \circ \{p\}$, *but* $w \notin B \circ \{p\}$. *Then both Subexpansion(*$\wedge$*) and Superexpansion(*$\vee$*) are invalid.*

Proof For *Subexpansion(*$\wedge$*)*, set $Q = \{p, w\}$ and $P = \{p\}$. It suffices to show that (i) $w \notin B * P$, (ii) $B * P$ is compatible with Q, and (iii) $w \in B * \big(P \wedge Q\big)$. But (i) is immediate from the assumption that $w \notin B \circ \{p\}$. For (ii), note that since p is consistent, by (T-Consistency) and (T-Completeness), for some world u, $u \in B * P$, and by (T-Success), $u \sqsupseteq p$, so $B * P$ is compatible with Q. For (iii), note that $B \circ \big(P \wedge Q\big) = B \circ \{p, w\} = \big(B \circ \{p\}\big) \cup \big(B \circ \{w\}\big)$. Since w is a world, by (T-Success) and (T-Consistency), $w \rightarrow_b u$ implies $u = w$ for all $b \in B$, so $B \circ \{w\} = \{w\}$. It follows that $w \in B \circ \big(P \wedge Q\big)$. Moreover, for all $v \in B \circ \big(P \wedge Q\big)$, either $v = w$ or $v \in B \circ \{p\}$. By assumption, either way we have $w \leq v$, and hence $w \leq v$ for all $v \in B \circ \big(P \wedge Q\big)$, so $w \in B * \big(P \wedge Q\big)$.

[49] Modulo some irrelevant differences resulting from the subtly different treatment of inconsistent updates using ■.

For *Superexpansion($\vee$)*, set $Q = \{w\}$ and $P = \{p\}$. As before, $w \notin B * P$. So it suffices to show that $w \in B * \big(P \vee Q\big)$, since then by $w \sqsupseteq p$ also $w \in \big(B * \big(P \vee Q\big)\big) \wedge P$. But $P \vee Q = \{p, w\}$, and we already showed above that $w \in B * \{w, p\}$. □

It is natural to wonder if the collapse may be avoided by invalidating *Subexpansion($\vee$)* and *Superexpansion($\wedge$)* instead of *Subexpansion($\wedge$)* and *Superexpansion($\vee$)*. But this is not so. Indeed, (PT-Expansion) implies (PT-Link), the principle required to establish *Subexpansion($\vee$)* or *Superexpansion($\wedge$)*. For assume $s \sqcup t \to_b w$. If $s \sqcup t$ is inconsistent, $w = \blacksquare$ and $w \leq v$ for any $v \in S^w$. If $s \sqcup t$ is consistent, w is a world containing $s \sqcup t$ and hence $w \in \{s\}^w$. By (PT-Expansion) $g\big(\{s\}^w\big) \subseteq B \circ \{s\}$. By (P-Limit), $g\big(\{s\}^w\big)$ is non-empty, so let $v \in g\big(\{s\}^w\big)$. Then $v \leq w$ and $v \in B \circ \{s\}$ as required.

Appendix C: A Space for Dominos

Finally, we construct a model of a doxastic state that satisfies the assumptions of the domino case and whose revision function is obtained from a coherent pair of plausibility ordering and transition relation. Let $f_1, f_2, \ldots$ be a countable infinity of sentence letters, and let L be the corresponding set of literals, i.e. the set including exactly the sentence letters f_n as well as their negations, which we write as $\overline{f_n}$. Say that a subset s of L is consistent iff for all n, at most one of f_n and $\overline{f_n}$ is a member of s. Let $S^\diamondsuit = \{s \subseteq L : s \text{ is consistent}\}$ and $S = S^\diamondsuit \cup \{L\}$. It is straightforward to show that $\mathcal{S} = \big(S, S^\diamondsuit, \sqsubseteq\big)$, with $\sqsubseteq$ interpreted as the subset-relation, is a topsy W-space, the set of worlds S^w being the set of the maximal consistent subsets of L.

We now define a plausibility ordering $\leq$ on the worlds. First, let $b = \{\overline{f_n} : n \in N\}$. Say that a world w is *regular* if for some n, $w = \{\overline{f_m} : m < n\} \cup \{f_m : m \geq n\}$, and irregular if not regular and distinct from b. Then for $w, v \in S^w$, we let $w \leq v$ iff (a) $w = b$, or (b) w is regular and $v \neq b$, or (c) w is irregular and v is irregular or identical to $\blacksquare$, or (d) $w = v = \blacksquare$. It is readily verified that $\leq$ satisfies the conditions of (P-Connectedness), (P-Transitivity), (P-Limit), and (P-Inconsistency).

Next, we define a transition relation $\to$. It will be sufficient to specify revisions of b by any state s. Indeed, for each state s we shall always specify a *unique* revision of b by s. If s is $\blacksquare$, we let $s \to_b t$ iff $t = \blacksquare$. If s is consistent, then for each n, we consider the largest $m \leq n$, if any, for which either f_m or $\overline{f_m}$ is a member of s, and we include f_n in our output state if s contains f_m, and $\overline{f_n}$ otherwise. More precisely, say that s is n-positive iff (a) $f_m \in s$ for some $m \leq n$, and (b) $f_m \in s$ for m the greatest number $\leq n$ for which either $f_m \in s$ or $\overline{f_m} \in s$. If s is not n-positive, then it is n-negative. Then let $\phi\big(s, n\big) = f_n$ if s is n-positive and $\overline{f_n}$ otherwise, and for consistent s, let $s \to_b t$ iff $t = \{\phi\big(s, n\big) : n \in N\}$.

Proposition 3 *$\leq$ and $\to$ are a coherent pair of plausibility ordering and transition relation on $\mathcal{S}$.*

Proof We skip the straightforward proof that $\leq$ is a plausibility ordering.

(T-Success): if $s = \blacksquare$, then $s \rightarrow_b t$ implies $t = \blacksquare$ and hence $t \sqsupseteq s$. If s is consistent, then $s \rightarrow_b t$ implies $t = \{\phi(s, n) : n \in N\}$. Suppose $f_n \in s$. Then s is n-positive and hence $f_n \in t$. Suppose $\overline{f_n} \in s$. Since s is consistent, $f_n \notin s$. So s is not n-positive, and hence $\overline{f_n} \in t$. So $s \sqsubseteq t$, as required.

(T-Completeness): $s \rightarrow_b t$ implies that either $t = \blacksquare$ or $t = \{\phi(s, n) : n \in N\}$. By construction, for each n, $\{\phi(s, n) : n \in N\}$ has either f_n or $\overline{f_n}$ as a member, so $\{\phi(s, n) : n \in N\} \in S^w$.

(T-Consistency): $s \rightarrow_b t$ implies $t = \{\phi(s, n):n \in N\}$ given that s is consistent. By construction of $\{\phi(s, n):n \in N\}$, it follows that for all n, at most one of f_n and $\overline{f_n}$ are members of $\{\phi(s, n):n \in N\}$, so t is consistent.

(T-Vacuity): If $s \sqsubseteq b$, then s contains no f_n as member. By construction, then neither does $\{\phi(s, n) : n \in N\}$, so $\{\phi(s, n) : n \in N\} = b$.

(T-Incorporation): Assume $s \rightarrow_b t$ and $r \sqsubseteq t$. We need to show that $s \sqcup r \rightarrow_b t$. Suppose first that s is inconsistent. Then $s \sqcup r = s$, and $s \sqcup r \rightarrow_b t$ follows immediately. Suppose then that s is consistent, so $t = \{\phi(s, n):n \in N\}$. By definition of $\rightarrow$, $s \sqcup r \rightarrow_b \{\phi(s \sqcup r, n) : n \in N\}$, so it suffices to show that for all n, s is n-positive iff $s \sqcup r$ is.

Suppose first that $s \sqcup r$ is n-positive. So for m the greatest number $\leq n$ such that $s \sqcup r$ contains either f_m or $\overline{f_m}$, we have $f_m \in s \sqcup r$. Since $s \sqcup r \sqsubseteq t = \{\phi(s, n):n \in N\}$, it follows that $\phi(s, m)$ is f_m, so s is m-positive. But since m is the greatest number $\leq n$ for which $s \sqcup r$ contains either f_m or $\overline{f_m}$, it follows that there can be no number k between m and n for which s contains $\overline{f_k}$, and hence it follows that s is n-positive also.

Suppose now that $s \sqcup r$ is n-negative. Then either (a) there is no $m \leq n$ with $f_m \in s \sqcup r$, or (b) we have $\overline{f_m} \in s \sqcup r$ for m the greatest number $\leq n$ for which either $f_m \in s \sqcup r$ or $\overline{f_m} \in s \sqcup r$. If (a), then there is no $m \leq n$ with $f_m \in s$, so s is n-negative. If (b), then since $s \sqcup r \sqsubseteq t = \{\phi(s, n):n \in N\}$, it follows that $\phi(s, m)$ is $\overline{f_m}$, so s is m-negative. But since m is the greatest number $\leq n$ for which $s \sqcup r$ contains either f_m or $\overline{f_m}$, it follows that there can be no number k between m and n for which s contains f_k, and hence it follows that s is n-negative also.

(PT-Existence): note that $g(S^w) = \{b\}$, and $s \rightarrow_b \{\phi(s, n) : n \in N\}$ for all $s \in S^\diamondsuit$, and $s \rightarrow_b \blacksquare$ for $s = \blacksquare$.

(PT-Link): Suppose $v \in B \circ \{s \sqcup r\}$, so $s \sqcup t \rightarrow_b v$. We need to show that $s \rightarrow_b u$ for some $u \leq v$. To that end, we establish

1. if $s \sqcup t \rightarrow_b b$, then $s \rightarrow_b b$
2. if $s \sqcup t \rightarrow_b w$ for some regular w, then $s \rightarrow_b v$ for some regular v
3. If $s \sqcup t \rightarrow_b w$ for some irregular, consistent w, then $s \rightarrow_b v$ for some consistent v

For (1), note that if $s \sqcup t \rightarrow_b b$, then $s \sqcup t \sqsubseteq b$ by (T-Success), so $s \sqsubseteq b$, so $s \rightarrow_b b$ by (T-Vacuity). For (2), we prove the contrapositive. Suppose that $s \rightarrow_b v$ and $v = \{\phi(s, n):n \in N\}$ is irregular. By definition of irregularity there are $m < k$ with both f_m and $\overline{f_k}$ members of v. By definition of ϕ and the fact that $b = \{\overline{f_n}:n \in N\}$ it follows that for some $m' < k'$, both $f_{m'}$ and $\overline{f_{k'}}$ are members of s, and hence of $s \sqcup t$. By (T-Success), both $f_{m'}$ and $\overline{f_{k'}}$ are members of w, which rules out w being regular. For (3), note that if $s \sqcup t \rightarrow_b w$ with w consistent, then $s \sqcup t$ is consistent, hence so is s, and so by consistency so is v with $s \rightarrow_b v$. □

We now show that our doxastic state satisfies the assumptions of the domino case under their obvious interpretation. Since some of these assumptions concern negated propositions, we will move to a bilateral conception of propositions as a pair of set of truthmakers and a set of falsitymakers. We shall take the revision of a belief state by a bilateral proposition to be simply the revision by the set of truthmakers, so our overall account of revision is not changed.

More precisely, we call a *bilateral proposition* $\mathbf{P}$ any pair of unilateral propositions. The first (second) coordinate of $\mathbf{P}$ is denoted by $\mathbf{P}^+$ ($\mathbf{P}^-$) and comprises the truthmakers (falsitymakers) of $\mathbf{P}$. Let $\mathbf{P} \wedge \mathbf{Q} = (\mathbf{P}^+ \wedge \mathbf{Q}^+,\ \mathbf{P}^- \vee \mathbf{Q}^-)$, $\mathbf{P} \vee \mathbf{Q} = (\mathbf{P}^+ \vee \mathbf{Q}^+,\ \mathbf{P}^- \wedge \mathbf{Q}^-)$, and $\neg\mathbf{P} = (\mathbf{P}^-,\ \mathbf{P}^+)$. $\mathbf{P}$ is said to be *exhaustive* iff every $w \in S^w$ contains either a member of $\mathbf{P}^+$ or a member of $\mathbf{P}^-$ as a part, and it is said to be *exclusive* iff no $w \in S^w$ contains both a member of $\mathbf{P}^+$ and a member of $\mathbf{P}^-$ as a part. Both properties can be shown to be preserved under the boolean operations, and it can also be shown that the logic of loose entailment over exclusive and exhaustive propositions is classical (cf. [10, pp. 665ff]).

Now let $\mathbf{F}_n = (\{f_n\},\ \{\overline{f_n}\})$ and $B = g(S^w)$. Note that $\mathbf{F}_n$ is always exclusive and exhaustive. Let $\mathbf{P} \Rightarrow \mathbf{Q}$ hold iff $B * \mathbf{P}^+ \models \mathbf{Q}^+$, and let $\Rightarrow \mathbf{Q}$ hold iff $B \models \mathbf{Q}^+$. Then

Proposition 4 *The belief state and revision function induced by $\leq$ and $\rightarrow$ satisfy the assumptions of the domino case:*

$$(B) \qquad \Rightarrow \neg\mathbf{F}_n f \text{ or all } n$$
$$(D.+) \quad \mathbf{F}_n \Rightarrow \mathbf{F}_m f \text{ or any } m \geq n$$
$$(D.-) \quad \mathbf{F}_n \Rightarrow \neg\mathbf{F}_m foranym < n$$

Proof (B) is immediate from the facts that $B = \{b\}$ and the definition $b = \{\overline{f_n} : n \in N\}$.

For (D.+) and (D.−), note that f_n is m-positive iff $m \geq n$, so $f_n \rightarrow_b t$ iff $t = \{\overline{f_m}:m < n\} \cup \{f_m:m \geq n\}$. So $B * \mathbf{F}_n^+ = \{\{\overline{f_m} : m < n\} \cup \{f_m : m \geq n\}\}$. So the sole truthmaker of $B * \mathbf{F}_n^+$ contains a truthmaker of $\mathbf{F}_m$ whenever $m \geq n$, as required for (D.+), and a truthmaker of $\neg\mathbf{F}_m$ whenever $m < N$, as required for (D.−). □

Acknowledgements For helpful comments on earlier drafts of this paper and/or discussion of the ideas here developed, I would like to thank Singa Behrens, Alexander Dinges, Martin Glazier, Stephan Leuenberger, Stefan Roski, Pierre Saint-Germier, Daniel Rothschild, Josh Schechter, Benjamin Schnieder, Moritz Schulz, Peter Verdée, Bruno Whittle, Steve Yablo. Special thanks are due to Kit Fine for numerous extended exchanges about these topics. My work on this paper was done as part of the Emmy Noether project *Relevance*, funded by the *Deutsche Forschungsgemeinschaft* (Grant number KR 4516/2-1); I am very grateful for that support.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

1. Alchourrón, C. E., P. Gärdenfors, & D. Makinson (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, *50*, 510–530. https://doi.org/10.2307/2274239.
2. Berto, F. (2019). Simple Hyperintensional Belief Revision. *Erkenntnis*, *84*, 559–575. https://doi.org/10.1007/s10670-018-9971-1.
3. Correia, F. (2016). On the Logic of Factual Equivalence. *The Review of Symbolic Logic*, *9*(1), 103–122. https://doi.org/10.1017/S1755020315000258.
4. Embry, B. (2014). Counterfactuals Without Possible Worlds? A Difficulty for Fine's Exact Semantics for Counterfactuals. *The Journal of Philosophy*, *111*, 276–287. https://doi.org/10.5840/jphil2014111522.
5. Fagin, R., & Halpern, J. Y. (1988). Belief, awareness, and limited reasoning. Artificial Intelligence, 39–76. https://doi.org/10.1016/0004-3702(87)90003-8.
6. Fermé, E., & Hansson, S. O. (2018). Belief Change. Introduction and Overview. Springer. https://doi.org/10.1007/978-3-319-60535-7.
7. Fine, K. (2012a). Counterfactuals without Possible Worlds. *The Journal of Philosophy*, *109*(3), 221–246. https://doi.org/10.5840/jphil201210938.
8. Fine, K. (2012b). A Difficulty for the Possible Worlds Analysis of Counterfactuals. *Synthese*, *189*, 29–57. https://doi.org/10.1007/s11229-012-0094-y.
9. Fine, K. (2016). Angellic Content. *Journal of Philosophical Logic*, *45*, 199–226. https://doi.org/10.1007/s10992-015-9371-9.
10. Fine, K. (2017a). A theory of truthmaker content I: conjunction, disjunction and negation, (Vol. 46 pp. 625–674). https://doi.org/10.1007/s10992-016-9413-y.
11. Fine, K. (2017b). A Theory of Truthmaker Content II: Subject Matter, Common Content, Remainder and Ground. *Journal of Philosophical Logic*, *46*(6), 675–702. https://doi.org/10.1007/s10992-016-9419-5.
12. Fine, K. (2018). Compliance and Command II, Imperatives and Deontics. *The Review of Symbolic Logic*, *11*(4), 634–664. https://doi.org/10.1017/S1755020318000059.
13. Fine, K. (ms). Review of Steve Yablo's Aboutness. Unpublished manuscript.
14. Gärdenfors, P. (1986). Belief Revisions and the Ramsey Test for Conditionals. *Philosophical Review*, *95*, 81–93. https://doi.org/10.2307/2185133.
15. Gärdenfors, P. (1987). Variations on the Ramsey Test: More Triviality Results. *Studia Logica*, *46*, 319–325. https://doi.org/10.1007/BF00370643.
16. Gärdenfors, P. (1988). Knowledge in Flux. Modeling the Dynamics of Epistemic States. MIT Press.
17. Grove, A. (1988). Two Modellings for Theory Change. *Journal of Philosophical Logic*, *17*, 157–170. https://doi.org/10.1007/BF00247909.
18. Hansson, S. O. (1999). A Textbook of Belief Dynamics. Theory Change and Database Updating. Kluwer.
19. Hansson, S. O. (2017). Logic of Belief Revision. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.* https://plato.stanford.edu/archives/spr2022/entries/logic-belief-revision/.
20. Huber, F. (2013). Belief Revision I: The AGM Theory. *Philosophy Compass*, *8/7*, 604–612. https://doi.org/10.1111/phc3.12048.
21. Jago, M. (2014). *The Impossible*. London: Oxford University Press.
22. Katsuno, H., & Mendelzon, A. O. (1991). Propositional knowledge base revision and minimal change. *Artifical Intelligence*, *52*, 263–294. https://doi.org/10.1016/0004-3702(91)90069-V.
23. Kourousias, G., & Makinson, D. (2007). Parallel Interpolation, Splitting, and Relevance in Belief Change. *Journal of Symbolic Logic*, *72*(3), 994–1002. https://doi.org/10.2178/jsl/1191333851.

24. Krämer, S. (202x). Truthmaker Equivalence. To appear in a volume for Alan Weir.
25. Lewis, D. K. (1973). *Counterfactuals*. Oxford: Blackwell.
26. Makinson, D. (1990). The Gärdenfors Impossibility Theorem in Nonmonotonic Contexts. *Studia Logica*, *49*, 1–6. https://doi.org/10.1007/BF00401549.
27. Makinson, D. (2009). Propositional Relevance through Letter-Sharing. *Journal of Applied Logic*, *7*(4), 377–387. https://doi.org/10.1016/j.jal.2008.12.001.
28. Parikh, R. (1999). Beliefs, belief revision, and splitting languages. *Logic, Language, and Computation*, *96*(2), 266–278.
29. Peppas, P., Chopra, S., & Foo, N. (2004). Distance semantics for relevance-sensitive belief revision. In *Principles of knowledge representation and reasoning: proceedings of the ninth international conference (KR2004)* (pp. 319–328).
30. Rantala, V. (1982). Impossible world semantics and logical omniscience. Acta Philosophica Fennica, 18–24.
31. Rothschild, D., & Yablo, S. (202x). Permissive Updates. Forthcoming. In F. Faroldi, & F. van de Putte (Eds.) *Outstanding Contributions to Logic: Kit Fine*. Springer.
32. Spohn, W. (2014). AGM, Ranking Theory, and the Many Ways to Cope with Examples. In S. O. Hansson (Ed.) *David Makinson on Classical Methods for Non-Classical Problems* (pp. 95–118). Springer. https://doi.org/10.1007/978-94-007-7759-0_6.
33. Stalnaker, R. (1968). A Theory of Conditionals. In N. Rescher (Ed.) *Studies in Logical Theory*, (Vol. 2 pp. 98–112). Blackwell. https://doi.org/10.1007/978-94-009-9117-0_2.
34. Stalnaker, R. (2009). Iterated Belief Revision. *Erkenntnis*, *70*(2), 189–209. https://doi.org/10.1007/s10670-008-9147-5.
35. van Fraassen, B. (1969). Facts and Tautological Entailment. *Journal of Philosophical Logic*, *66*(15), 477–487. https://doi.org/10.2307/2024563.
36. von Wright, G. H. (1963). *Norm and Action*. Chicago: Routledge & Kegan.
37. Yablo, S. (2011). A Problem about Permission and Possibility. In A. Egan, & B. Weatherstone (Eds.) *Epistemic Modality* (pp. 270–294). Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199591596.003.0010.
38. Yablo, S. (2015). *Aboutness*. Princeton: Princeton University Press.
39. Özgün, A., & Berto, F. (2020). Dynamic Hyperintensional Belief Revision. *The Review of Symbolic Logic*, 1–46. https://doi.org/10.1017/S1755020319000686.