Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Department of Social Psychology

Institute of Psychology

Faculty of Psychology and Movement Science

Spontaneous State Inferences

Felix Kruse

Dissertation submitted to Universität Hamburg in partial fulfillment of the requirements for the academic degree Dr. rer. nat.

March 31, 2022

Felix Kruse

Spontaneous State Inferences

Dissertation submitted to Universität Hamburg in partial fulfillment of the requirements for the academic degree Dr. rer. nat., March 31, 2022

Date of Disputation:	June 13, 2022
Head of Committee:	Prof. Dr. Martin Spieß
1. Dissertation Reviewer:	Prof. Dr. Juliane Degner
2. Dissertation Reviewer:	Prof. Dr. Leonel Garcia-Marques
1. Disputation Reviewer:	Prof. Dr. Jan Wacker
2. Disputation Reviewer:	Prof. Dr. Martin Spieß

Felix Kruse
Universität Hamburg
Institute of Psychology
Department of Social Psychology
Von-Melle-Park 5
20146 Hamburg

Abstract

People are social animals. In order to successfully navigate through our social environment, we routinely form impressions of others. One way to do so is observing and interpreting others' behavior. Research on spontaneous trait inferences (STIs) has revealed that people use behavioral information to spontaneously infer dispositional information (i.e., traits) about actors. Dispositional inferences have since then been central and ubiquitous in social psychological theorizing about person perception – at times even coined *fundamental attribution error* (Gawronski, 2004). While previous research has provided substantial evidence that trait inferences can and do occur, it has neither demonstrated that traits are the only or even the dominant category of inference, and neglected considerations of the occurrence of other spontaneous person inferences. My dissertation research aims at bridging this theoretical and empirical gap, and demonstrating that people can and do infer *mental states* from behavior at least as much as they infer personality traits.

This dissertation consists of three related series of experiments. In Chapter 3, I present a series of four pre-registered experiments focused on providing proof of concept with regard to spontaneous state inferences. We employed two classical experimental paradigms repeatedly used in STI research, a false recognition task in Experiments 1-3 and a probe recognition paradigm in Experiment 4, and behavioral statements that potentially allowed for either a trait- or a state-inference (Experiments 1 and 2) or both inferences simultaneously (Experiments 3 and 4). Results document significant effects of both, trait and state inferences with moderate effect sizes (d_z = 0.22 - 1.60). We thus replicated the established trait-inference effect and, more importantly, provided first empirical evidence for the occurrence of spontaneous state inferences. In Experiment 5, we further demonstrated that participants ascribe differential predictive value to traits and states, indicating a functional distinction between state and trait inferences.

In the further two chapters, we aimed at providing a better understanding of process characteristics of state inferences.

In Chapter 4, we aimed at investigating whether the mechanisms underlying state and trait inferences can be understood as mere activation of behavior-trait associations as compared to actual actor-inferences. We observed that both trait and state inferences were comparable with regards to their inferential nature.

In Chapter 5, we first discuss if and to what extent person inferences from behavior rely on automatic processes. In a series of two pre-registered experiments, we investigate the efficiency of inference processes as one exemplary automaticity feature, by testing whether they occur under conditions of limited cognitive resources. Our experiments provide evidence that working memory load reduced effect sizes of simple state and trait inferences to a comparable degree. When using more complex behavioral statements that allow for simultaneous inferences of both traits and states, however, high working memory load eliminated inference effects.

Our findings have strong implications for theory building regarding the underlying mechanisms and processes of first impression formation in person perception. We have demonstrated that state inferences represent a viable alternative to trait inferences. They occur spontaneously and simultaneously with trait inferences. Thus, we need to reconsider our understanding of social inferences and the processes underlying the occurrence of biases such as the correspondence bias. Current theories do not account for spontaneous inferences of psychological states, let alone the occurrence of multiple, simultaneous spontaneous inferences. I present a framework that integrates available evidence into a coherent understanding of the social inference process and speculate on mechanisms not yet supported by empirical data. Our findings contribute towards expanding these theories to paint a more complete picture of the psychological processes involved in perceiving others' behavior.

Acknowledgements

First, I would like to thank you, Juliane, for introducing me to the world of academia, for your encouragement and your support over the years. Thank you for helping me shake off the insecurities that come with being a first-generation academic.

Thank you, Anna, Felix, Anna, and Carsten, for sharing parts of this journey with me. I could not have done it without your moral support.

Several people helped a great deal by providing comments and helpful advice. Thank you, Leonel, for your support right from the beginning, a very warm welcome in Lisbon and the opportunity to share my work with your Cognition in Context lab. Thank you, Klaus, for encouraging me to publish our first studies in a high-ranking journal. Thank you, Marine and Mirka, for your patience in helping me grasp multilevel modeling. And a big thank you to my colleagues and the staff at the Institute of Psychology in Hamburg, whose helpfulness has constantly been outstanding.

A number of outstanding research assistants have supported this dissertation by never refusing tedious tasks involving stimulus development and study testing - thank you, Tania, Sven, Ricardo, Nicoleta, Louis, Marlene, & Silke.

In this dissertation, I briefly mention some of the experiments conducted by my Master's Thesis students, Antonia and Eva. Thank you for the dedication and scrutiny with which you engaged in our collaborations.

Special thanks to my former roommate Marvin, who volunteered to appear in this dissertation as the example stimulus Marvin who laughed at a joke - a fitting description. Thank you for all your support, Marvin.

Danke, Mama & Papa.

Contents

1 Introduction

> *Felix, qui potuit rerum cognoscere causas*
>
> — **Virgil (29 BC)**

One of the best things in life, in my opinion, is encountering other people. And there are plenty of these encounters: Unless you live in a small, rural community, you will come across new people almost on a daily basis. Our minds have developed strategies to deal with these situations. For instance, if you came across Marvin, and observed him laugh about a joke, you might draw the conclusion that Marvin is a jolly person - someone who is chronically prone to be happy and cheerful. In other words, you may infer one of Marvin's personality traits as the cause of his behavior, which may be beneficial to you as it allows you to assess the situation and to adjust your own behavior accordingly. Research suggests that people form these trait inferences spontaneously, that is, unconsciously and without intent (but see Chapter 5). While there are good reasons why people might infer others' personality traits, or - in other words - draw *spontaneous trait inferences* (STIs), why should spontaneous impressions be limited to these stable dispositions? If I yawned and told you that I was going to bed, why should you not infer that I was *tired*, without necessarily over-generalizing and assuming that I am a *lazy* person? Mental states are a ubiquitous part of our lives. We frequently query each other about our respective states ("How are you?"), we explain our behavior in terms of states ("I was too tired to go out on the weekend"), and are quite adapt at recognizing emotional states from facial expressions. And yet the existing evidence on impression formation from behavior suggests that we appear to constantly over-generalize other peoples' behavior and infer their personality traits from single instances of behavior. This dissertation project revolves around the question of whether people might spontaneously infer states from behavior (they do), whether these spontaneous state inferences may occur simultaneously with other inferences (they do), and how these insights warrant changing our understanding of the social inference process and require further specifying the existing theories on impression formation from behavior.

As I will attempt integrating our findings with existing theories, it will first be necessary to provide a quite thorough overview over them, and their historical development. In Chapter

2, I thus review available theorizing and past research on STIs. I outline how the idea of spontaneous impressions evolved from early research on attribution and explicit impression formation. I will set forth my concerns with limitations of past research, and argue that our research provides answers to some of the thus far unresolved challenges in the interpretation of the results of past research.

In Chapter 3, I present our first series of experiments, which serve as a proof-of-concept: We observed consistent evidence across four pre-registered experiments using two different experimental paradigms that people spontaneously infer mental states when perceiving others' behavior. Our results even suggest that people may draw inferences about both states and traits *simultaneously*. We published these results in the *Journal of Personality and Social Psychology* (Kruse & Degner, 2021). This proof-of-concept is the most important contribution of our research, because it raises important questions about our understanding of how people form impressions of others' behavior. I discuss these implications in Chapter 6.

Chapter 4 addresses cognitive processes underlying spontaneous state inferences. The methods we used in Chapter 3, while established methods in STI research, allow for the potential alternative interpretation that not inferential processes, but mere associative processes are at play. Instead of *attributing* states to an actor, participants may merely *associate* actor and state, the same way they might associate an actor and an apple, if they were presented together. Thus, in a series of two experiments, we provide evidence that, while part of the effect size can be explained by associative processes, inferential processes are indeed at play in spontaneous state inferences.

In the final set of experiments, we investigated whether spontaneous state inferences possess characteristics of automatic processes beyond spontaneity. In Chapter 5, I describe the first experiments we conducted investigating the efficiency of spontaneous state inferences (do they also occur if perceivers spare little cognitive resources?), and discuss the more challenging results of these studies.

Finally, I discuss the implications the research presented in this dissertation has for our understanding of the social inference process, integrate them into a final conclusion, and provide an outlook for future research.

2 Perceiving Others' Behavior

> *Social psychology since Heider's time has perhaps overemphasized the importance of other people's traits as a form of intentionality, to the relative exclusion of other types of person causes (temporary goals, moods, etc.) that could explain the reasons for their behavior.*
>
> — **Moskowitz (2005, p. 237)**

In this chapter, I will review relevant research on attribution theory and explicit impression formation from behavior. I will outline how research on spontaneous trait inferences (STI) developed from there and review the literature on STI. I will conclude this chapter by outlining some of the limitations of past STI research.

2.1 Attribution

The social world we live in is extremely complex. So complex, in fact, that to a human perceiver, it is virtually impossible to process all the information about other people available to us. If you walked the streets and encountered a person laughing, you would probably infer that the person is *jolly* - and save yourself the trouble of analyzing all the information that is available in this small encounter. In other words, you go beyond the information available to you, and try to make sense of what you see - a laughing person. The person laughs, so they probably are a jolly person. Attributing a cause to an observed event helps people in dealing with the complexity of the world around them (e.g., Kelley, 1967). Moreover, inferring person characteristics allows predicting future behavior (e.g., Dunning et al., 1990; McCarthy & Skowronski, 2011b; Nussbaum et al., 2003), which allows us to form expectations about the other person, so that we can adjust our own behavior accordingly. This may reduce uncertainty (e.g., Heider, 1958; Trope & Gaunt, 2000), and contribute to maintaining the illusion of controllability (Biner et al., 1995; Bruner, 1957).

The idea of complexity reduction by ascribing causes for behavior, or inferring the "so-called dispositional properties of the world" (Heider, 1958, pp. 79 - 80), has directed social psychological research since its very beginning. Asch (1946, p. 48) began his seminal paper on impression formation as follows:

> We look at a person and immediately a certain impression of his character forms itself in us. A glance, a few spoken words are sufficient to tell us a story about a highly complex matter. We know that such impressions form with remarkable rapidity and with great ease. Subsequent observations may enrich or upset our view, but we can no more prevent its rapid growth than we can avoid perceiving a given visual object or hearing a melody.

In this brief introduction, Asch anticipated much of the social inference research to come. He emphasized that people may form impressions of others' characters, and may do so with some degree of automaticity - that is, unintentionally, unconsciously, uncontrollable, and effortlessly (see Chapter 5 for a more thorough discussion). A great number of studies conducted over the last decades support Asch's theorizing.

While Asch focused explicitly on a person's character as prime inferential goal, Heider (1958) discussed two main forces that could account for people's actions: those emanating from the person and those emanating from the situation. As person forces, he named not only the stable character, disposition, or essence of the person, but explicitly included temporary states, motives, goals, and sentiments. Despite his acknowledgement of these temporary factors, Heider was nevertheless convinced that people would tend to favor conclusions about a person's personality, and speculated about what might compel people to generalize so readily. In his opinion, "behavior engulfs the field" (Heider, 1958), meaning that another person's behavior draws attention to the person instead of to the context in which it occurs, and thus makes the person much more salient than the situation. This disproportionate salience of behavior and actor makes it more difficult to process situational factors that might strongly impact the person's behavior: For example, while it is quite easy to see a person blushing and stuttering while giving a presentation, it is more difficult and requires more effort to see the situational pressure exerted by an auditorium of critical attendees. What is salient might thus capture perceivers' attention, and perceivers are likely to attribute causal roles to what they perceive as salient (e.g., Taylor & Fiske, 1975). Heider (1944, 1958) further argued that given people's adversity towards uncertainty and doubt, they would prefer explanations that provide stable, clear, and concrete causes, thus reducing complexity and uncertainty. Attributing the cause of behavior either to the person's stable personality or to the person's mental and motivational states would provide the clearest, simplest, most certain type of explanation, according to Heider.

Jones and Davis (1965) proposed the first systematic model of such dispositional inferences in impression formation, the theory of *correspondent inference*. Jones and Davis posited that only intentional actions were informative about actors. They understood intentional actions as those in which the actor is both aware of the effects his or her actions have and has the ability to achieve those effects. They further posited that, only if an action was perceived as intentional, a dispositional inference would be drawn. If perceivers were aware, for example, that an actor unintentionally performed a certain behavior (e.g., a person falling over after being pushed), no correspondent inference should be drawn.

Shortly thereafter, in the seminal "Castro-Study", Jones and Harris (1967) provided first empirical evidence in support of the theory of correspondent inference, with an important addition: They observed the occurrence of correspondent inferences even when participants were explicitly informed about situational constraints. They presented participants with one of two essays about Fidel Castro, either ostensibly written by another student, along with information that the essay's author had been either required to write a pro-Castro vs. anti-Castro essay, or had been free to choose which stand they took in their essay. Jones and Harris then asked the readers of the essays to rate the writers' personal attitude towards Castro. They found that perceivers judged writers' attitudes as correspondent to the attitudes expressed in their essays, *even when* they were informed that the writers had been required to write a certain essay. So not only was situational information less salient as compared to an actor and their behavior, as Heider (1958) posited, but situational information was ignored in the impression formation process, even when it was explicitly made available to the participants.

In the seminal "Quizshow-Study", Ross et al. (1977) developed another paradigm that provided evidence in support of correspondent inference: Participants were asked to take part in an ostensible quiz show. One participant was assigned to be a contestant in the quiz. Another participant was assigned to observe the behavior of the contestant as they attempted to answer difficult questions. The questions were written by a third participant playing the role of the quiz master, who was instructed to think of difficult questions of their area of expertise. This task is relatively easy for the quiz master, but relatively difficult for the contestant. Contestants were likely to perform poorly on this task, independently of their individual intelligence. Thus, participants were likely to observe a contestant failing to answer between 30 and 50 percent of the quiz master's questions. Ross and colleagues found that participants rated the quiz master as more intelligent than the contestant: They did not consider the situational constraints inherent in the task, and attributed contestants' failure to answer the questions to their dispositions instead. In conclusion, people appear to be biased towards attributing behavior to stable dispositions when asked to form impressions of others, even in cases where perceivers are aware of situational constraints that invalidate the dispositional inference. Ross (1977) went as far as to label the phenomenon *fundamental attribution error* (but see Gawronski, 2004).

Correspondent inferences have been characterized as a "remarkably robust and easily replicated phenomenon" (Jones, 1979, p. 107), or even "the most robust and ubiquitous finding in the domain of interpersonal perception" (Jones, 1990, p. 164), if not "a candidate for the most robust and repeatable finding in social psychology" (Jones, 1990, p. 138). And to date, the correspondence bias effect has indeed been replicated repeatedly and consistently (e.g., Bauman & Skitka, 2010; Klein et al., 2018). Some moderating conditions have been identified. For instance, perceivers' ability for perspective taking (Hooper et al., 2015) and their mindfulness (Hopthrow et al., 2017) have been shown to moderate the effect size of the correspondence bias (for a more extensive review, see Gawronski, 2004).

Generally, however, until the early 1980s, impression formation was thought to be an effort-dependent process and perceivers were considered "lazy", and would thus only exert such effort when needed: "A person's interests [. . .] determine when he will become motivated to make attributions at all" (Kelley & Michela, 1980, p. 473). This idea of the "lazy perceiver" integrated well with the perspective of the *cognitive miser* (Taylor, 1981), which regarded resource preservation as a central motive for humans. Perceivers were thus assumed to only form impressions of others when they were motivated to do so, for instance if they anticipated to interact with a person in the future, and if they spared the cognitive resources to engage in an effortful impression formation process. Kelley (1967, 1971, 1972, 1973) developed a systematic model of attribution. He assumed that, as *naïve scientists*, perceivers would systematically analyze behavior on three dimensions: consistency (does the actor show the behavior at different times?), distinctiveness (does the actor show the behavior in different situations?), and consensus (does the actor's behavior differ from the behavior of people around her/him?). Kelley posited that perceivers would engage in a systematic analysis of covariance between these factors, if they had access to multiple instances of a person's behavior in which they could search for the respective information.

McArthur (1972) exemplified the covariance analysis using the example "John laughs at the comedian". This behavior may be caused by the person (John), the stimulus (the comedian), the circumstances (the comedy club on that night), or a combination of these factors. If John is the only person laughing at the comedian (low consensus), he laughs at the comedian at other comedy clubs (high consistency), and he laughs at other comedians (low distinctiveness), then it is likely something in John that caused him to laugh. If everyone is laughing at the comedian (high consensus), John laughs at the comedian at other comedy clubs (high consistency), and he does not laugh at other comedians (high distinctiveness), then it is likely that something about the stimulus (the comedian) caused John to laugh. If everyone is laughing at the comedian (high consensus), John doesn't laugh at the comedian at other comedy clubs (low consistency), and he laughs at other comedians at the club (low distinctiveness) then it is likely something about the circumstance (the comedy club on that night) that caused John to laugh. This systematic covariance analysis is only possible if perceivers have access to multiple instances of behavior. When perceiving a single or initial instance of behavior, according to Kelley, perceivers would need to analyze the given configuration of possible causal factors, and in the presence of multiple causal factors, weigh the influence of each of the factors. If a factor that sufficiently explains the behavior is encountered, the other factors are disregarded, the so-called *discounting principle*. Frequently, these sufficient explanations are personality traits. However, more recent research does not support the assumption that perceivers actually engage in a systematic search for covariance, even if this information is made available to them (Lalljee et al., 1984).

In summary, early research on impression formation concluded that perceivers appear to attribute the causes of behavior to personality traits as the result of a deliberate process in which perceivers exert effort to analyze possible causes of observed behavior. But does it make sense to construe all impression formation as such deliberate processes? Most contemporary psychologists would beg to differ, but when Winter and Uleman (1984) first investigated whether people would also draw inferences spontaneously, that is, in the absence of deliberate intention to form an impression, they were surprised by their findings. They observed that in fact, perceivers inferred traits *spontaneously* from written behavioral statements. They defined *spontaneous inferences* as inferences occurring without intention to engage in an effortful impression formation process as described above. To the contrary, trait inferences seemed to have been drawn quite readily. Winter and Uleman sparked a long and fruitful line of research on *spontaneous trait inferences*, that I will review in the following section.

2.2 Spontaneous Trait Inferences

Fig. 2.1.: Typical stimulus used in the false recognition paradigm.

Spontaneous trait inferences (STI) describe impression formation processes that occur "without intentions or instructions, at the encoding stage of processing behavioral information" (Winter & Uleman, 1984, p. 237). Investigating such inferences poses an interesting challenge to researchers. In order to ensure that inferences drawn from a behavior are *spontaneous*, directly asking participants about their impression is not an option - this would prompt participants to engage in a deliberate impression formation process, which may or may not differ from what happens when impressions are formed spontaneously. Typically in research on spontaneous trait inferences, people are presented with brief verbal descriptions of behaviors that imply, but do not explicitly mention a trait, without any intention to form impressions or to infer traits. Participants are then instructed to memorize behaviors for a subsequent memory test, or to familiarize themselves with "stimuli that will be used later in the study". For example, they might read, "Marvin laughed about the joke", which implies that he is jolly. The challenge is to find out when people make such inferences, without asking them directly. While quite a few indirect experimental paradigms have been developed and used successfully, in the following, I will focus on describing one of these methods in detail, and discuss how it allows conclusions regarding spontaneous inference processes in person perception. An overview over other available methods can be found in Bott et al. (2022).

In the false recognition paradigm (Todorov & Uleman, 2002), participants are first asked to memorize a series of statements describing individual actors' behaviors that imply, but not explicitly mention a trait, along with the respective actor's name and portrait (see Fig. 2.1, for an example). To avoid that participants engage in deliberate impression formation, they are instructed to memorize the information for a memory test. The rationale of this paradigm is that if participants spontaneously infer the traits implied by the statements, they are encoded in memory along with the information that participants deliberately memorize about the actors, such as the actor's name, their portrait, as well as information provided in the actor's statement. In our example, "Marvin laughed about the joke", the implied trait "jolly" might be inferred and encoded, along with the name "Marvin", his portrait, and the idea that he laughed at a joke. After memorizing a number of stimuli in rapid succession, participants are presented with the image of an actor and their name paired with a probe word. They are instructed to indicate whether the probe word had occurred in the respective actor's statement or not. In the crucial conditions, these probes are the traits implied by the statements presented along with the actors. Of course, the correct answer in this case is no - the trait did not occur in the actor's statement, it was merely implied by it. However, if participants did indeed infer the trait from the statement as they read it, this will make it more difficult for them to correctly reject the probe word presented to them. Participants are likely to fall prey to a source-monitoring failure (Johnson et al., 1993), in which the familiarity of the trait can be misattributed to a prior presentation. In other words, participants fail to discriminate between the inference and the information actually presented. In a control condition, traits implied by a *different* actor's statement are presented as probes, to control for familiarity effects that are not bound to the respective actors. The mean difference between participants' errors (*false recognitions*) in these two conditions is the dependent measure. If participants show more false recognitions in the critical condition, as compared to the control condition, researchers infer that a spontaneous trait inference has occurred.

Plenty of evidence documenting the occurrence of STIs has been provided over the past 35 years. It would go beyond the scope of this dissertation to review all the available evidence (others have done so, see Bott et al., 2021; Uleman et al., 2012). I will thus focus on two important issues, namely whether STIs represent actual inferences about actors, as opposed to categorizations of behavior, and whether they actually do happen when perceiving behavior.

First, it remained unclear whether STIs did indeed represent person inferences, as compared to mere categorizations of behavior. In other words, if you infer that Marvin is jolly, do you think that *Marvin* is a *jolly person* or that *laughing* is a *jolly activity*? The false recognition paradigm described above already suggests that inferences are linked to the actors: Traits are presented with only the actor's portraits and names as cues and compared to traits presented with different but familiar actors. This suggests that STIs are actually bound to the actor, not only to the behavior. Participants, when presented with Marvin's portrait, activate the trait *jolly* which they had inferred earlier, causing an increased probability of false recognitions, as compared to control trials. An alternative explanation is plausible: If participants used traits only to categorize behavior, the same pattern of results could be observed as a result of the following processes: Participants categorize the behavior in terms of the trait, in our example, they infer that *laughing* is a *jolly activity*. When presented with the portrait, it prompts recall of the behavior, which in turn activates the trait and might lead to higher false recognition rates, without proving the occurrence of an inference about the actor. This alternative explanation, however, could be refuted. Todorov and Uleman (2002, Experiment 5) showed that even if participants were unable to recognize the behavioral statements, and thus could not have recalled them, they were more likely to falsely recognize traits paired with the respective actor, as compared to randomly paired traits. STIs can thus confidently be considered *person inferences*. Further evidence for the inferential nature of STIs is discussed thoroughly in Chapter 4.

If we now assume that STIs do in fact represent person inferences, another interesting question is whether it really is the behavior that prompts these trait inferences. It is conceivable that single words contained in trait-implying statements bear a strong semantic association with the implied trait. For instance, the word *joke* in our example might be so strongly linked to the trait *jolly* that it alone activates the trait, instead of an understanding of the complete behavioral episode. Orghian et al. (2019) raised this concern and provided a method to investigate it. They presented participants with two versions of a behavioral statement, one in the original order (*Marvin laughed at the joke*), and one rearranged with roughly the same words (e.g., *Marvin read a joke about someone laughing*). Note that while both statements contain words that might be linked to the implied trait *jolly*, only the statement in the original order actually implies it, whereas the other does not. If participants showed indications of STIs for the rearranged statements, they would have to be activated by single words, and not the behavior. Orghian and colleagues were able to show that for their sample of statements, no word-based activation effects occured, which suggests that it is more likely that really the behavior had prompted the formation of STIs.

We need to address one more important aspect of STIs, namely the question of when they occur. We interpreted the false recognition effects as evidence for trait inferences during encoding of behavioral statements. Another explanation is conceivable, namely that STIs only occur in the recognition phase, when participants are presented with the implied traits paired with portraits and then retrospectively connect them to the statement. This would mean that, when participants are presented with behavioral statements, no inference is drawn. Instead, when they are later presented with the implied trait paired with the portrait, they might recall the actual behavioral statement and see the implied trait as fitting the actor's behavior. This, in turn, could make the implied trait seem familiar and lead them to fail to reject the implied trait. For the false recognition paradigm, it is unlikely that participants are able to consciously recall all of the statements, as typically, a large number of them are presented for only a very brief period of time (in our adaptation of the paradigm, 36 statements are presented for 5-6s each). Todorov and Uleman (2002, Experiments 5 and 6) demonstrated false recognition effects even when participants were presented with as many as 120 behavioral statements for 5s each. When they presented participants with only the portraits and asked them to recall the behavioral statement corresponding to the portrait, participants recalled an average of 12% of behaviors. However, we cannot assume trait inferences as the only possible reason for false recognitions, but must assume that in some trials at least, a confound with these retrieval-based effects may occur.

In sum, in paradigms such as the false recognition paradigm, which require participants to store behavioral information in long-term memory, inference effects do seem to occur during encoding, but might be overestimated due to processes at retrieval. This account integrates well with the meta-analytical observation that long-term memory based measures typically yield larger effect sizes ($d_z = 0.68$), as compared to on-line measures ($d_z = 0.38$; Bott et al., 2022), such as the probe recognition paradigm (see Section 3.4.3, for a detailed description).

2.3 Models of Impression Formation from Behavior

The occurrence of STIs provided evidence that perceivers are by no means *lazy* - they appear to constantly engage in impression formation processes, even without explicit goals of doing so. Perceivers' readiness to infer traits might explain why people appear to be biased towards dispositional attributions. In fact, several of the early models of impression formation from behavior provide a framework for understanding dispositional attribution as failed corrections of initial spontaneous trait inferences. Below, I will briefly review the most relevant models.

Several theoretical accounts have attempted to describe the processes involved in impression formation from behavior, and to account for the bias towards dispositional attributions. Gilbert et al. (1988) divided person perception from behavior into three sequential processes: categorization, characterization, and correction. They posited that in order to form an impression, perceivers would first categorize an actor's behavior, that is, for example, identify the distinct pattern of muscular contraction as laughing (Frank & Ekman, 1996), to then characterize the actor performing such behavior in terms of a trait, in our example, a *jolly* person. This initial trait inference could then be corrected by considering possible situational constraints that might have caused the behavior. For instance, we might realize that the laughing person is walking out of a comedy club, in which he saw a *funny* comedian - a situation that might explain his laughing behavior just as well. Gilbert and colleagues further posited that while the first two steps occur relatively automatically, the correction in the last step is optional as it requires deliberate effort, motivation, and opportunity. Perceivers would thus need to be motivated to engage in a deliberate correction process, and invest the cognitive resources to identify possible situational constraints. Finally, they may correct or confirm their initial trait inference. Gilbert et al. argued that the additional motivation and resources that need to come into play for a trait inference to be corrected may be a possible explanation why perceivers appear to be biased towards dispositional attribution.

Fig. 2.2.: Ambiguous face: Anger - Happiness (adapted from Trope, 1986, p. 245)

Trope (1986) suggested that situational information would already be considered in the categorization of behavior, for instance to disambiguate behavior whose categorization might vary depending on the context. A screaming person might, for instance, be angry, if they are attacked, or happy, if their favorite team scored a goal (Fig. 2.2).

While Gilbert and Trope both assumed that situational information may supplement dispositional inferences, Krull (1993) devised a model that included initial inferences about both traits and situations. Depending on their processing goals, perceivers might initially infer either a trait of the actor or a property of the situation. If, for example, a perceiver ran into our smiling person, they might be interested in what the person is like, in order to form expectations of their behavior, and spontaneously infer the trait *happy*. If, however, a perceiver was about to watch the movie the smiling person just watched, they might be more interested in what they could expect from that movie, and less concerned about the person.

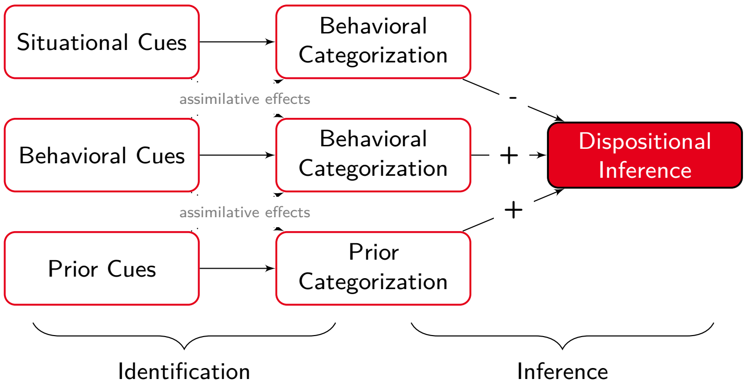

Fig. 2.3.: Two-stage model of dispositional inference. Adapted from Trope (1986).

They might then infer that the person is smiling because the movie is *funny*. Subsequently, perceivers could engage in a correction process, to either adjust their initial inference for person vs. situational factors. Krull and Erickson (1995) summarized the model (see Figure 2.4).

The presented models on impression formation from behavior document peoples' tendency to attribute the cause of others' behavior to stable dispositions or situational properties. While the explanations for the occurrence of correspondence biases that have been offered so far do seem quite convincing, they remain theoretical. In the following, I will discuss evidence that the existing models fail to acknowledge and that does not integrate with these accounts. I will discuss how the focus on spontaneous trait inferences may have lead scholars in the field to overlook the possibility that other types of spontaneous social inferences occur just as well - and review the existing evidence on some of these other inferences. I will argue that, in sum, the evidence suggests that the available models are underspecified and our understanding of the processes underlying correspondence biases is severely limited, and finally describe the empirical contribution the current research has to add to it.

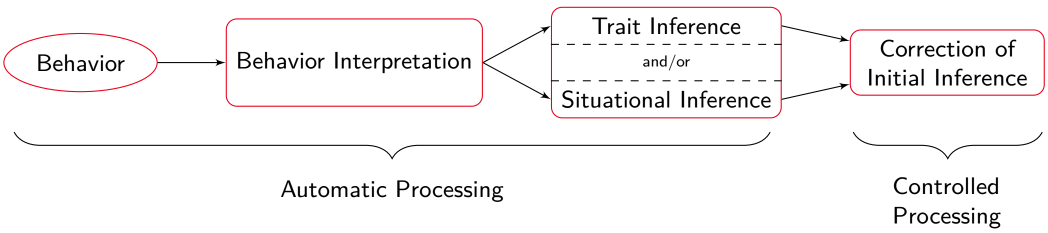

Fig. 2.4.: Social Inference Model. Adapted from Krull and Erickson (1995).

A common explanation for the occurrence of corrensponderce biases, also frequently used in social psychology lectures and text books, goes as follows: Upon perceiving a behavior, observers infer a trait, with some degree of automaticity. More often than not, they do not engage in the effortful process of considering context information that might suggest a situational cause for the observed behavior, and could result in a corrected inference. Then, when observers are asked to attribute behavior, a trait is likely to be named due to its availability (Moskowitz, 2005, p. 309; Ferreira et al., 2012, p. 10). Over the course of the last 35 years, a large amount of evidence supporting spontaneous trait inferences has been provided. The assumption that trait inferences are ubiquitous and that first impression formation is dominated by (more or less) fundamental biases in favor of dispositional inferences, has become established textbook content in social psychology (e.g., Aronson et al., 2019; Hewstone et al., 2016; Smith et al., 2014). However, closer inspection of the available empirical evidence from the large body of studies on spontaneous trait inferences suggests that this may be a premature conclusion, (a) because many studies allow only limited generalization to impression formation from behavior in general and (b) because studies have barely ever tested spontaneous trait inferences against other person inferences that may occur with the same level of spontaneity, thus lacking discriminant validity. We will look at these limitations in more detail below.

2.3.1 Limited Generalizability

When reviewing the STI research conducted during the last three decades for a meta-analysis (Bott et al., 2022), we noted that stimulus materials in many studies seem to be created such that they strongly promote trait inferences. First, behavioral statements often describe rather unusual, rare, or extreme behaviors (e.g., “The farmer paints a swastika on the synagogue wall”, “The barber loses 20 lbs. in 6 weeks on a new diet”; Winter & Uleman, 1984; but see Levordashka & Utz, 2017, for a recent exception). Unexpected or untypical behaviors that violate consensus (Kelley, 1967) and/or desirability (Jones & Davis, 1965) are more diagnostic of actors’ dispositions because they can hardly be explained by situational causes. Dispositional trait inferences are thus not only more appropriate attributions for such behaviors, but probably the only sensible attributions possible. STIs observed with such stimulus materials, however, may not be generalizable for impression formation based on less extreme and more mundane behaviors, that are much more frequent in everyday life (e.g., Levordashka & Utz, 2017).

Secondly, stimulus materials often contain temporal markers indicating repeated behaviors and thus temporal consistency (e.g., “I *usually* crawl out of bed at noon . . . ”, “I attend my church *twice a week* . . . ”; Carlston & Skowronski, 1994 [emphases added]). Repeated behaviors are generally more likely to be related to dispositional causes, thus traits. After all, traits tend to be defined as habitual patterns of behavior, thought, and emotion (e.g., Hamaker et al., 2007). Indeed, high temporal consistency of behavior has been demonstrated to increase trait inferences in deliberate impression formation (e.g., Olcaysoy Okten & Moskowitz, 2018) and also appears to increase spontaneous trait inferences - at least with people who are generally less likely to form trait inferences (i.e., people with liberal ideology; Olcaysoy Okten & Moskowitz, 2019, Experiment 1). Thus, the use of behavioral statements including such consistency markers may have led to an overestimation of the strength and frequency of spontaneous trait inferences from behavior.

Finally, we also found a number of behavioral statements that additionally contain expressions of inner dispositions such as attitudes, beliefs, or values (e.g.,"I *hate* animals. Today . . . I saw this puppy. So I kicked it out of my way.", "I *like* to write short stories and poetry, and I spend much of my free time doing this writing . . . "; Carlston & Skowronski, 1994 [emphases added]). If behavioral statements contain verbal expressions that explicitly assert attitudes and preferences, they elicit a corresponding implicational schema (e.g., Reeder & Brewer, 1979) and guide perceivers' attention towards stable inner psychological dispositions of actors as potential causes of behavior, therefore making trait inferences more likely (Semin & Fiedler, 1988).

In sum, it appears that stimulus materials in many STI studies have provided more than just behavioral descriptions. The inclusion of further trait-implying information in the stimulus materials of previous research, thus "trait-tailored" stimulus materials (Malle & Holbrook, 2012), makes the observation of spontaneous trait inference effects less compelling. To say the least, we cannot conclude from these findings that spontaneous trait inferences are ubiquitous inferences that people generally draw from observations of behavior.

2.3.2 Limited Discriminant Validity

Additionally, we currently lack empirical support for the notion that trait inferences are indeed the only or even dominant person inference perceivers spontaneously draw from actors' behaviors. Most studies on STIs have exclusively used trait words as their dependent measures, thus constraining participants' responses to trait inferences. For participants in such a study, it is thus impossible to demonstrate any other person inference from behavior. If we only prompt participants on trait inferences, we are only able to observe trait inferences, while other possible person inferences remain undetected. For example, studies on deliberate impression formation that used a higher variety of response options or open response formats not only document a wide array of person inferences but also question the presumed inevitability or priority of trait inferences (e.g., Malle & Holbrook, 2012). In this line, the few available studies on spontaneous impression formation that provided participants with alternative - goal - probes in their measurements demonstrated that people also draw spontaneous goal inferences when processing others' behaviors (e.g., Hassin et al., 2005; Olcaysoy Okten & Moskowitz, 2019; Van Overwalle et al., 2012). If studies do not allow for alternative spontaneous inferences and do not directly compare them to spontaneous trait inferences, the available research lacks the discriminant validity needed for supporting the conclusion that STI are ubiquitous or dominant person inferences in impression formation (Uleman, 2005).

In summary, we believe that the available evidence for STIs is limited due to (a) non-representative behavioral stimuli, which prevents generalization to the behaviors that people observe in their daily lives, and (b) non-representative dependent variables, which prevents the conclusion that traits are the dominant, let alone the only person inference people spontaneously draw.

2.4 Multiple Spontaneous Inferences

Spontaneous trait inferences have been investigated intensively, and the evidence in their support has repeatedly been called ubiquitous (Schneid, Carlston, et al., 2015; Uleman, 2005). Unfortunately, it seems that the ubiquity of evidence for the occurrence of spontaneous trait inferences is often misconceived as evidence for the ubiquity of trait inferences in impression formation from behavior. For example, many established textbooks in social psychology refer only to trait inferences when discussing spontaneous person perception (e.g., Hewstone et al., 2016; Myers & Twenge, 2018). This representation, however, overlooks two important aspects: On the one hand, researchers have repeatedly discussed and empirically documented multiple spontaneous inference activation, especially simultaneously occurring person and situation attributions. Ham and Vonk (2003) investigated whether people would draw both trait and situational inferences from behavior descriptions. They presented participants with statements that implied both, a trait and a situational property. For instance, the statement "Marvin laughed at the joke" may imply that Marvin is a *jolly* person, and/or that the joke he read was *funny*. Immediately afterwards, participants were asked to indicate whether a probe word had been explicitly mentioned in the statement.

The answer in these critical trials is no, of course, but participants took more time to reject the probe word when it had been implied by the statement presented just before the probe, as compared to non-implied control words. Ham and Vonk interpreted these results as indicating that participants had inferred a trait or situational property from the statement. More importantly, they were able to show that between participants, both traits and situational properties were spontaneously inferred from the same statements. Todd et al. (2011) provided evidence that spontaneous inferences about traits and situational properties may even occur *simultaneously*, by manipulating trait vs. situational probes within-participants (e.g., Lupfer et al., 1990; Todd et al., 2011). On the other hand, person dispositions as attributions of behavior can encompass more than only personality traits. For example, studies on deliberate impression formation that used a higher variety of response options or open response formats document a wide array of person inferences (ranging from intentions and desires to values and beliefs) and question the sometimes presumed inevitability or priority of trait inferences (e.g., Malle & Holbrook, 2012). In this line, a number of studies on spontaneous impression formation demonstrated that people draw various types of both short-lived and stable inferences: spontaneous goal inferences (e.g., Hassin et al., 2005; Olcaysoy Okten & Moskowitz, 2019; Van Overwalle et al., 2012), spontaneous motivational inferences (e.g., Reeder, 2009a, 2009b; Reeder et al., 2001), and spontaneous emotional inferences (Diergarten & Nieding, 2016). Thus, when making sense of others' behavior, perceivers are able to spontaneously draw multiple inferences: I may infer that Marvin is a *jolly person*, while simultaneously inferring that the *joke* may be very *funny*, and/or that Marvin may have the current goal of *being entertained* and having a good time.

In light of these more recent findings, it seems clear that spontaneous impressions of others are not limited to traits, but encompass a variety of person inferences. Evidence for some of these inferences already exists in the literature. We argue that these may not be limited to goals and motivations, but more generally include inferences about an actor's current mental states, such as affective or emotional states (e.g., happy, surprised, anxious, ashamed), cognitive and attentional states (e.g., focused, interested, distracted, bored) and physiological states (e.g., hungry, full, sick, tired).

2.4.1 Spontaneous State Inferences

Psychological theorizing and research has separated traits from states in a number of different ways. A common theoretical account defines traits as stable, inter-individual differences in peoples' proneness, tendency, style, or disposition to behave, feel, or think in certain ways (e.g., Hamaker et al., 2007). States also describe person dispositions for thoughts, feelings, or behaviors. However, states refer to *intra*-individual differences that reflect peoples' continuous adaptation to situational demands. Thus, whereas traits are typically conceptualized by their relative stability, consistency, and in-variance over time and across situations, states are characterized by their relative instability, inconsistency, and variance in response to temporal or situational variation (Hamaker et al., 2007). In our studies, we operationalize traits and states as variable in their relative temporal stability. Although by no means a binary criterion, traits tend to endure over time, while states are rather short-lived phenomena, and should thus influence predictions about future behavior in different ways.

It is conceivable that peoples' lay personality theories contain a similar differentiation of traits and states that in turn affects spontaneous impression formation and behavior attribution. For example, a person attribution as conceptualized in Kelley's seminal attribution model (1967) locating the cause of the behavior entirely within the person and their stable invariant characteristics may be represented by a trait inference. Nevertheless, a person-by-entity attribution or a person-by-situation attribution that locates the cause of observed behavior in person characteristics in their interaction with external and unstable causes may be represented by a state inference. Of course, Kelley's attribution model presumed elaboration, information search, and processing in order to arrive at a conclusion about the causes of observed behavior. However, there are several theoretical and empirical perspectives that support the assumption of spontaneous state inferences as well. First, some theoretical accounts of spontaneous impression formation actually take state inferences as a given (e.g., Uleman, 2005), at times even regard them as the "default mode" of understanding behavior (Korman & Malle, 2016; Malle & Holbrook, 2012). Heider (1958) originally assumed in his considerations of peoples' naive analyses of action that person inferences for causal attributions of behavior include contemplation of transitory person states such as fatigue and mood, attitudes and needs, or social and legal status. Similarly, in their considerations of correspondence inferences, Jones & Davis (1965) argue that trait inferences about actors rely on inferences about actors' intentional states. A similar argument is brought forward in Reeder's multiple inference model (MIM; Reeder, 2009a, 2009b). It is striking that although several of the most seminal theoretical accounts of interpersonal impression formation and attribution have addressed state inferences, research on spontaneous impression formation has not given these inferences much necessary attention.[1] Indeed, the need to investigate

[1] This has been repeatedly attributed to early misinterpretations of Heider's considerations by equating person attributions with trait attributions (e.g., Malle, 2008)

whether and to what extent people spontaneously infer states from behavior has been stated repeatedly (e.g., Lillard & Skibbe, 2006; Uleman et al., 2008). Mental states have seen some more attention in research on deliberate impression formation. There is indeed empirical evidence that people infer mental states from behavior when explicitly prompted to form impressions (e.g, Ames, 2004) or when asked to write about their impression of others (McClung & Reicher, 2018), and readily explain intentional actions in terms of beliefs, desires, values and internal states (e.g., Malle, 2004; Malle & Holbrook, 2012; Olcaysoy Okten & Moskowitz, 2019, Experiment 3). Related research in the field of developmental psychology has repeatedly documented children's and adults' ability to deliberately infer and use others' affective and cognitive states, termed *mentalizing* or *theory of mind* (ToM; e.g., Hamlin et al., 2013; Ruffman, 2014; Scott & Baillargeon, 2017). However, no compelling evidence for the (non-)automaticity of process characteristics underlying ToM-effects inferences has been provided yet (Apperly et al., 2006) and we lack empirical research investigating if and to what extent any of the above listed state inferences occur spontaneously (e.g., Lillard & Skibbe, 2006).

There is another field of research, which may provide important insights for the current research question, namely the field of text comprehension. Given that most STI research relies on written statements about others' behavior (but see Fiedler & Schenck, 2001; Fiedler et al., 2005), theorizing and research on text comprehension may provide helpful insights to the question which inferences people draw spontaneously from written behavior statements. There are indeed several theoretical accounts of text comprehension that describe spontaneous inferences about mental states from text (e.g., Cook & O'Brien, 2017; Graesser et al., 1994; Kintsch, 1998), and empirical evidence for their occurrence has been provided, at least for emotional states (e.g., Gernsbacher et al., 1992).

There are good reasons why mental state inferences could and should occur *spontaneously* in the social inference process. First, research on person perception has already documented that people are generally able to process situational information when forming impressions (e.g., Reeder et al., 2001). Trope's (1986) model of impression formation and attribution even assumes that situational information is initially processed and used in order to identify observed behavior (i.e., situational inducement). Moreover, when situational constraints are highly salient, people are less likely to infer traits (Jones & Davis, 1965). For example, observing a person crying at a funeral would most likely lead observers to infer that this is a sad situation without necessarily drawing the conclusion that the crying person is generally depressive. Yet, we cannot assume that situations directly *cause* behavior. It is thus not the funeral that makes the person cry: we have to assume that situations or their appraisals impact people who in turn respond with behavior. Previous research has indeed documented that people process others' mental states when thinking about eliciting situations (e.g., Thornton & Tamir, 2020). In our above example, it is the emotional *state* of sadness caused by the sad situation that leads to the crying behavior and it is reasonable to assume that most observers would draw the same conclusion.

Second, it has been repeatedly argued that the presumed dominance of trait inferences in impression formation results from their high functional value: Knowing a person's dispositions allows predicting their future behavior (e.g., Heider, 1958; Hoffman et al., 1981; McCarthy & Skowronski, 2011b). However, drawing state inferences can be equally informative. On the one hand, inferring others' mental states does enable a perceiver to derive situation-specific behavioral expectations and tailor their own behavior accordingly (e.g., Thornton & Tamir, 2020). On the other hand, state inferences can also signal that the current observation may not warrant predicting future behavior, or that behavior predictions should be limited to the very short-term, thus preventing erroneous over-generalizations about others.

Finally, recent theorizing in personality psychology actually defines the person-descriptive aspects of traits as density distributions of states (i.e., whole trait theory; Fleeson & Jayawickreme, 2015; Jayawickreme et al., 2019). In this understanding, a trait-ascription is related to a person's frequent manifestations of trait-related states. For example, a person who is described with the trait label *shy* is expected to *feel and act in in shy ways* frequently and in many situations, an idea originating with Mischel (1968). Applying whole trait theory to person perception suggests that in order to ascribe a trait to an actor, observers might need to recognize the actor's current state, infer that the actor experiences this state frequently and in many situations, and thus generalize that this current state is a representative manifestation of an underlying trait.

In summary, there are several theoretical and empirical reasons that support our notion that perceivers may spontaneously infer mental states from behaviors. When referring to mental states, we explicitly include *any* behavior-related person condition that is temporally or situationally limited. We thus propose a wider range of inferences than the considerations of intentionality, desire, and belief proposed by Malle (e.g., Malle, 2005; Malle & Holbrook, 2012), and further include affective, emotional, cognitive, attentional, and physiological states.

2.4.2 The Current Research

The goal of the current research was to investigate the spontaneous occurrence of state inferences in impression formation from behavior. We therefore conducted a series of nine experiments using different stimulus materials and established experimental paradigms. In the first set of two experiments, we employed single-implication behavioral descriptions that allow for either unambiguous trait inferences (e.g., "I gave the homeless man five euros" - generous) or unambiguous state inferences (e.g., "When my sister and her husband exchanged rings, I just couldn't hold back the tears" - touched). We used a false recognition paradigm (Todorov & Uleman, 2002), collected data online (Experiment 1), and compared results with data collection in the lab (Experiment 2). We then developed a second set of dual-implication behavioral stimuli that allowed for simultaneous trait and state inferences (e.g., "Vanessa read the book until late at night" - studious, interested) which we employed in a false recognition paradigm (Experiment 3) and a probe recognition paradigm (Study 4; Todd et al., 2011). To foreshadow results: All four studies provided robust evidence for the spontaneous and simultaneous occurrence of state and trait inferences. Finally, in Experiment 5, we explored if and to what extend people functionally distinguish between state and trait inference in impression formation, focusing on perceived predictability of future behavior (Chapter 3). In Chapters 4 and 5, we explore further process characeristics of spontaneous state inferences. All experiments were preregistered with the Open Science Framework (Chapter 3: https://osf.io/v5j78; Chapter 4: https://osf.io/795q8/; Chapter 5: https://osf.io/pz2r4/) where we also provide open access to materials, raw data, and analyses codes. We report how we determined our sample sizes, all data exclusions, all manipulations, and all measures in the description of each experiment and the Appendix. The faculty's local ethics committee approved all procedures (protocol number: 2018_180, spontaneous person inferences).

2.4.3 Material Generation

For the current research endeavor, it was paramount to create stimulus materials that allow for a valid differentiation between state and trait inferences without favouring either inference a priori. We therefore conducted a pilot study and a series of pretests in order to develop one set of single-implication behavioral statements that exclusively allow for either trait or state inferences (to be used in Experiments 1, 2, and 8) and another set of dual-implication behavioral statements that allow for simultaneous trait and state inferences (to be used in Experiments 3 - 7 and 9). We summarize this extensive preparatory research here and provide detailed descriptions and results of all pretests in the Appendix (p. 137).

Given that our experiments relied on established experimental paradigms that use adjective probes, we first conducted a pilot study to establish a set of person-describing adjectives that people use unequivocally to either refer to traits or states. We briefed a sample of participants ($n = 55$) about trait and state concepts and asked them to rate the relative stability vs. variability of a list of 323 adjectives. Based on the resulting ratings, we chose those 69 adjectives that were rated as the most stable as trait adjectives (e.g., introverted, smart, ambitious) and the 72 most variable as state adjectives (e.g., bewildered, thirsty, disgusted) as basis for further stimulus generation. In a next step, we generated a set of 196 statements in German language that described single behaviors implying these selected states or traits (without explicitly mentioning the respective adjectives). We ensured that behavioral statements focused on relatively mundane behaviors and did not contain any additional linguistic markers that might elicit trait or state inferences. We submitted these statements to a pretest, in which we asked participants to name a person-describing adjective that came to mind when reading a statement. For Experiments 1, 2, and 8, we selected only statements that reached a consensus score of 50% or higher[2]. The selected trait-implying and state-implying statements did not differ significantly in agreement rate or statement length (see Appendix, p. 144).

For Experiments 3 - 7 and 9, we created a second set of dual-implication behavioral statements in English language that described more ambiguous behaviors potentially allowing for simultaneous trait and state inferences. For statement creation, we relied on a large sample of crowd-workers. We provided participants with a list of person describing states selected from our pilot study and asked them to think of a behavior that could be indicative of this state while at the same time being attributable to a person's trait (see Appendix, p. 151, for detailed instructions). Based on these responses, we selected a set of 288 statements that we submitted to a further pre-test. We first briefed an independent sample of participants ($N = 171$) about the trait and state concepts as stable vs. variable actor characteristics and asked them to name both a state and a trait that may come to mind when reading a behavioral statement. Again, we only selected statements to be included in Experiments 3 - 7 and 9, that reached a consensus score of 50% or higher for both, trait and state, with the

[2]Note that in some cases the adjectives named by participants were not identical with the pre-tested adjectives on which we had based stimulus generation. Separate analyses, however, revealed no significant effects on the pattern of results.

further constraint that mean consensus scores for traits and states did not differ across all selected stimuli (see Appendix, p. 152).

3 Spontaneous State Inferences

3.1 Experiment 1

In the first experiment, our goal was to establish whether people show indications of spontaneous state inferences when presented with unambiguous state-implying behavioral statements and compare these to trait inferences. We adapted the established false recognition paradigm (Todorov & Uleman, 2002). In this paradigm, participants are first asked to memorize a series of statements describing individual actors' behaviors paired with acotrs' names and portraits. In a later recognition phase, they are presented with the images of the actors and indicate whether a probe word occurred in the statement or not. In the crucial conditions, these probes are adjectives implied by the presented behavior. Spontaneous trait inferences are inferred from higher false recognition rates in the implied-trait condition (e.g., erroneously responding "yes" to an adjective that was implied but not presented in a behavioral statement) as compared to a control condition. In the present experiment we implemented the paradigm using behavioral statements either unambiguously implying states or traits in order to (a) establish whether significant indicators of state inferences can be observed and to (b) provide a first effect size comparison between state and trait inferences. We used a similar approach as Levordashka and Utz (2017), presenting behavioral statements to appear like ostensible posts on a social media platform.

3.1.1 Method

3.1.1.1 Sample Size Determination

Based on a minimum effect size of interest of Cohen's $d_z = 0.20$ for crucial one-tailed within-sample *t*-tests comparing between implied and implied-other conditions (see *Procedure*) for the trait and the state condition, both conditions required a minimum of $n = 156$ participants, respectively, to provide enough statistical power ($1 - \beta = .80$), calculated using G*Power 3.1 (Faul et al., 2007). However, to avoid unnecessary spending, we relied on a sequential testing procedure (Lakens, 2014) with one interim analysis planned at 37 valid data sets or time = .26 for each condition, using a Pocock-type spending function calculated with the GroupSeq package for R (Pahl, 2018). We pre-registered to stop data collection if the observed effects were significant at the interim analysis at $\alpha_1 = .018$. If they were not, we planned to continue data collection until $N = 184$ valid data sets would have been collected and perform the final analysis with $\alpha_2 = .032$. In order to secure the required numbers of valid data sets after applying the pre-registered exclusion criteria, we overpowered both studies by 12.5%, thus planning to collect data from $n = 41$ and $n = 207$ participants per condition for the interim and final analyses, respectively.

3.1.1.2 Participants

The current experiment relies on valid data from a total of $N = 86$ participants (45 female; average age $M = 28.9$ years, $SD = 8.7$, ranging from 18 to 58). The majority of participants indicated being native speakers of German (87%), while 13% indicated to speak German as one of their native languages. Participants were recruited via Prolific (www.prolific.ac) and received monetary compensation of 1.49 GBP (approx. 1.91 USD) for the duration of ten minutes. Participants were randomly assigned to either the trait condition ($n = 45$) or the state condition ($n = 41$). Data of three initial participants who had received erroneous instructions due to a programming error were not included into analyses. Following our pre-registered exclusion criteria, we excluded three further participants because they did not pass the initial attention test and 9 further participants because they aborted the experiment before debriefing and thus did not give informed consent for data analyses.

3.1.1.3 Materials

We used a set of 18 statements that implied, but not explicitly mentioned a trait, as well as 18 statements that implied, but not explicitly mentioned a state in German language. Each statement was presented together with a person's portrait picture (selected from the 10k US adult faces database; Bainbridge, Isola, & Oliva, 2013), and a name, designed such that they appeared like messages from an instant messaging service application (Levordashka & Utz, 2017, see Figure 3.1 for an example, and Tables A.6 - A.9 in the Appendix for the complete list of statements and pre-test results).

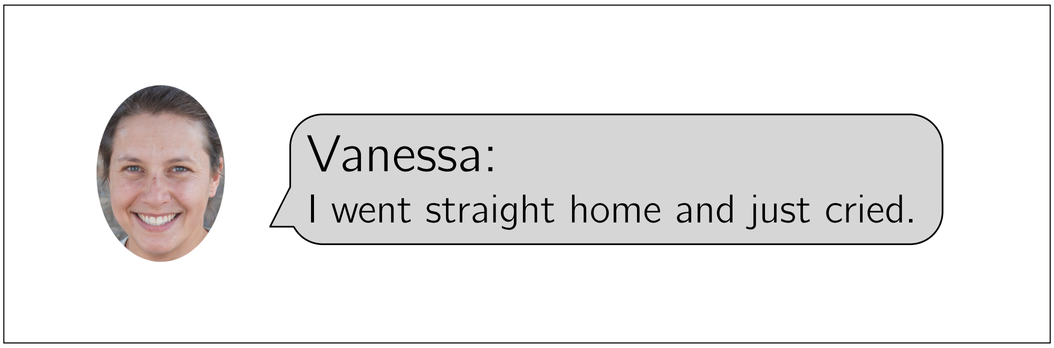

Fig. 3.1.: Example Stimulus (Studies 1 and 2). Original Stimuli were in German language. Original portrait pictures differ. Printed portrait pictures were adapted from Karras et al. (2020), for copyright reasons.

3.1.1.4 Design

We had pre-registered this research prior to data collection as two independent studies, one on trait inferences and one on state inferences. However, because both studies were conducted simultaneously and participants were randomly assigned to either of the studies, we collapsed data and treated both studies as a between-subjects condition in the analyses. Thus, the present experiment followed a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) mixed design with the factor Inference varying between participants and Condition varying within participants. We implemented an additional between-subjects factor based on stimulus-set assignment: Using a counter-balanced Latin square design, stimuli were assigned to separate sets of six to be presented equally often in the implied-, implied-other, and new condition.

3.1.1.5 Procedure

Data were collected online using the platform Qualtrics for online data collection (www.qualtrics.com) and Prolific for participant recruitment (www.prolific.ac). The experiment started with a welcome page that contained an initial attention check requiring participants to click on a logo instead of the continue button (Oppenheimer et al., 2009). If participants failed the attention check, they received a notice and were asked to re-read the instructions on the welcome page. If they failed the attention check a second time, they were excluded from participation and directed back to Prolific.

Participants then started the learning phase of the false recognition paradigm (Todorov & Uleman, 2002), in which they were presented with one of two sets of stimuli (trait vs. state) containing 18 targets and nine fillers, presented in individual random order. Each individual stimulus was displayed for five seconds. Participants were instructed to read the stimuli carefully in preparation of a memory test.

Directly following the learning phase, participants completed the recognition task. In each recognition trial, an actor's name and portrait from the learning phase were presented together with a single probe adjective. Participants were instructed to click on a "yes" button when they recognized this probe word to have been presented earlier in the statement of the same actor and to click on a "no" button when they did not recognize the probe word to have been presented in the statement. The 18 portraits from the target trials of the learning phase were split into three recognition conditions, such that within each participant, six portraits each were presented with (a) the specific trait/state adjective previously implied by the respective actor's statement (implied trait/state probe), (b) with a trait/state adjective implied by a *different* actor's statement (implied-other trait/state probe), and (c) with a new trait/state adjective, which had not been implied in any of the statements (new trait/state probe). We further balanced the valence of implied-other and new control probes, so that half were of opposite valence, and half were of matching valence with regard to the implied trait or state (e.g., Schneid, Carlston, et al., 2015; Schneid, Crawford, et al., 2015)[1]. All target probes required a "no" response. To avoid response biases, the recognition test included nine filler trials in which we presented the actors' portraits paired with the trait or state adjectives that were explicitly mentioned in the filler statements during the learning phase, thus requiring "yes" responses. We recorded responses and response latencies. After the test phase, participants provided demographic information (age, gender, education, profession). Participants also self-reported language proficiency using a 7-point scale (1: "German is my only native language" to 7: "My German is not good enough to understand this question"). Finally, participants were fully debriefed about the purpose of the experiment and once again asked for consent for data storage and analyses. We had pre-registered an exploratory measure of Implicit Personality Theories (Dweck et al., 1995), which, however, was not included in this experiment because of a programming oversight. This measure was included in Experiment 2 instead.

3.1.2 Results

We conducted all data analyses using R 3.6.3 (R Core Team, 2020). As customary with the false recognition paradigm, we focused analyses on false response rates (Todorov & Uleman,

[1]Exploratory analyses demonstrated that valence (in)congruence of implied-other and control traits did not significantly qualify the reported results. We therefore collapsed analyses across this factor.

2003). Analyses of response latencies for correct rejections and further exploratory analyses are reported in Table A.12 in the Appendix. The pre-registered main analyses use ANOVA to analyze differences between participants' scores averaged per condition, aggregated across statements. Data were submitted to a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) ANOVA with Inference varying between participants and Condition varying within participants. Results show a significant main effect for Condition, $F(2, 168) = 28.56$, $p < .001$, $\eta^2_G = .13$, 90% CI [.06, .21], and a significant main effect for inference, $F(1, 84) = 11.45$, $p = .001$, $\eta^2_G = .06$, 90% CI [.01, .17], but no significant interaction, $F(2, 168) = 0.46$, $p = .634$, $\eta^2_G < .01$, 90% CI [0, .03], see upper left panel of Figure 3.2. Separate analyses confirmed expected effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied trait condition ($M = .30$, $SD = .27$) than in the implied-other trait condition ($M = .21$, $SD = .22$), $t(44) = 2.246$, $p = .03$ (one-tailed), $d_z = 0.376$, 95% CI [0.03, 0.72] and in the new trait condition ($M = .11$, $SD = .15$), $t(44) = 4.818$, $p < .001$ (one-tailed), $d_z = 0.868$, 95% CI [0.45, 1.29]. More importantly, in the state inference condition, participants also showed higher false recognition rates in the implied state condition ($M = .45$, $SD = .27$), than in the implied-other state condition ($M = .33$, $SD = .22$), $t(40) = 2.508$, $p = .016$ (one-tailed), $d_z = 0.474$, 95% CI [0.07, 0.87] and in the new state condition ($M = .20$, $SD = .24$), $t(40) = 5.124$, $p < .001$ (one-tailed), $d_z = 0.97$, 95% CI [0.51, 1.43].

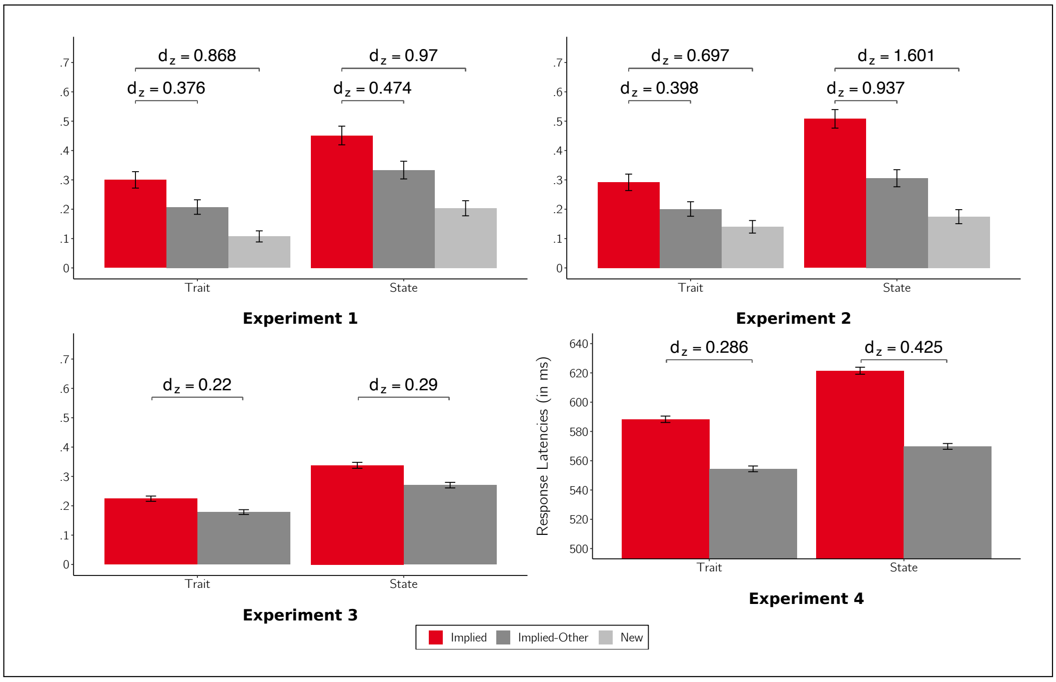

Fig. 3.2.: Mean proportion of false recognition responses (Experiments 1-3) and mean response latencies for correct responses (Experiment 4) as a function of type of inference and experimental condition. Error bars represent *SE*.

3.1.3 Discussion

The results of this first experiment provided initial evidence that people *can* draw both trait and state inferences when processing behavioral information about strangers: In our adaption of the false recognition paradigm, participants were more likely to falsely recognize state and trait adjectives that were implied by behavioral descriptions as compared to non-implied state or trait words. Furthermore, effect sizes were in a comparable range for both the state inference effects as well as the trait inference effects. Note that our stimuli were created such that they neither enhanced trait nor state inferences, by describing relatively mundane behaviors, carefully avoiding temporal markers and being scarce with situational descriptions.

The interpretation of our results is, however, limited by the following caveat: When conducting the interim analyses according to the sequential testing procedure, we had erroneously accepted the difference between implied and implied-other conditions in the trait-inference condition as significant albeit the p-value of $p = .03$ did not fulfill the preregistered significance criterion of $p \leq .018$. We thus prematurely stopped data collection. However, given that the pattern of results replicates typical STI findings, we opted against resuming data collection when noticing our mistake, and decided instead to invest our resources into a replication experiment conducted in the laboratory.

3.2 Experiment 2

Experiment 2 served as close replication of Experiment 1 with the only difference being that data collection was conducted in the laboratory instead of online.

3.2.1 Method

3.2.1.1 Sample size determination

Based on the effect sizes resulting from Experiment 1, we aimed at providing enough statistical power ($1 - \beta = .80$) to detect effect sizes of $d_z = 0.30$ with $\alpha = .05$ (one-tailed) for both, the state and trait conditions of this experiment. This would require a sample size of $n = 142$ per condition, calculated using G*Power 3.1 (Faul et al., 2007). Again, we relied on a sequential testing procedure (Lakens, 2014) with one interim analysis planned at 68 valid data sets or time = .38. We pre-registered to stop data collection if the observed effects were significant at the interim analysis at $\alpha_1 = .025$. If they were not, we planned to continue data collection until 180 valid data sets would have been collected and perform the final analysis with $\alpha_2 = .025$. In order to obtain the required numbers of valid data sets after applying the pre-registered exclusion criteria, we overpowered both experiments by 12.5%, resulting in $n = 71$ and $n = 189$ for the interim and final analyses, respectively, for each condition. During data collection it became apparent that an unexpectedly high number of participants appeared to be non-native speakers of German - whose data would eventually need to be excluded from analyses (see sample description). Additionally, the randomized assignment of participants lead to a high imbalance of participants in the trait and state condition (with only $n = 28$ in the state and $n = 40$ in the trait condition at interim analyses). We therefore deviated from the pre-registered sample size and collected data from 117 participants in order to achieve a more balanced assignment of participants with sufficient language proficiency to both conditions. Note that interim analyses at the pre-registered $n = 68$ already fulfilled the aforementioned decision criteria to warrant applying the stopping rule (see Appendix, p. 156).

3.2.1.2 Participants

The current experiment relies on valid data from a total of $N = 91$ participants (34 female, average age $M = 29.9$, $SD = 11.2$, ranging from 18 to 71 years). The majority of participants indicated being native speakers of German (65%), 21% indicated to speak German as one of their native languages, and 14% indicated that they spoke German very well, albeit it was not their native language. Participants were mainly students from various faculties of Hamburg University, recruited via a university online job platform and were compensated 2.50 EUR (approx. 3.13 USD) for the duration of 15 minutes. Our experiment was the first to be conducted in an one-hour lab session followed by an unrelated experiment on face recognition. Participants were randomly assigned to either the trait condition (n = 42) or the state condition (n = 44). Following our pre-registered exclusion criteria, we excluded data from further participants: Six because they did not pass the initial attention test, twelve because they aborted the experiment before debriefing and thus did not give informed consent for data analysis, and 18 because they self-reported insufficient language proficiency.

3.2.1.3 Procedure

Experiment 2 used the same materials and followed the same procedure for the false recognition task as Experiment 1, with the exception that it was conducted in the laboratory. After completion of the false recognition paradigm, participants additionally completed a Navon task (Navon, 1977) and a measure of implicit personality theories (Implicit Personality Theories Questionnaire [8 items], translated into German; Dweck et al., 1995). Description and results of the exploratory analyses are reported in the Appendix (p. 157).

3.2.2 Results

Individual false recognition rates were submitted to a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) ANOVA with the factor Inference varying between participants and Condition varying within participants. Results show significant main effects for Condition, $F(2, 168) = 40.75$, $p < .001$, $\eta^2_G = .18$, 90% CI [0.10, 0.27] and for Inference, $F(1, 84) = 12.56$, $p < .001$, $\eta^2_G = .07$, 90% CI [.01, .18]. We also observed a significant interaction between Inference and Condition, $F(2, 168) = 5.84$, $p = .004$, $\eta^2_G = .03$, 90% CI [.00, .08], see upper right panel of Figure 3.2. Separate analyses confirmed expected effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied condition ($M = .29$, $SD = .24$) than in the implied-other trait condition ($M = .20$, $SD = .21$), $t(43) = 2.779$, $p = .008$ (one-tailed), $d_z = 0.398$, 95% CI [0.1, 0.69], and in the new trait condition

(M = .14, SD = .18), $t(43) = 4.224$, $p < .001$ (one-tailed), $d_z = 0.697$, 95% CI [0.33, 1.06]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition (M = .51, SD = .22) than in the implied-other state condition (M = .31, SD = .21), $t(41) = 4.336$, $p < .001$ (one-tailed), $d_z = 0.937$, 95% CI [0.42, 1.45], and in the new state condition (M = .18, SD = .20), $t(41) = 8.061$, $p < .001$ (one-tailed), $d_z = 1.601$, 95% CI [1, 2.2].

Further analyses exploring the significant interaction effect of Inference type by Condition indicated that the mean difference of implied state and implied-other state (M = .202, SD = .303) was more than twice as large as the mean difference of implied trait and implied-other trait, albeit not significantly (M = .091, SD = .217), $t(84) = 1.97$, $p = .052$, $d_s = 0.425$, 95% CI [-0.01, 0.86].

3.2.3 Discussion

Experiment 2 replicated the results of Experiment 1, thus strengthening our previous conclusion that participants can spontaneously draw state inferences as well as trait inferences. What is more, state inference effects were numerically larger than trait inference effects.

While Experiments 1 and 2 serve as important proofs of concept and starting point for our research, there are, however, two characteristics of these first two experiments that limit the interpretability of the results. On the one hand, we had employed only behavioral descriptions that uniquely and unambiguously implied *either* trait *or* state inferences. Using a Devil's advocate argumentation, one may assume that via extensive pretesting we succeeded in creating behavioral statements tailored to actively inhibit trait inferences, thereby prompting state inferences. For example, if someone states that "*There is nothing going on this weekend, I am wasting my time channel-surfing*", this person is most obviously bored by the specific situation which is hardly attributable to a personal disposition without further information. Furthermore, we had implemented the state and trait condition as a between-participants factor, which may have affected participants' general mode of information processing during the learning phase, increasing the general likelihood of inferring states from behavior as compared to a potentially dominant focus on trait inferences in spontaneous impression formation.

In Experiments 3 and 4, we employed stimulus material with higher ecologic validity that – in principle – allow for the simultaneous occurrence of state and trait inferences. For example, in our introductory example of Marvin laughing at the joke, one may assume that he is a *jolly* person who generally laughs a lot or that he is so *amused* by this specific joke that he bursts out laughing. We use a false recognition paradigm (Experiment 3) and a probe recognition paradigm (Experiment 4). While Experiments 1 and 2 were conducted in German language (with native German speakers as participants), Experiments 3 and 4 were conducted in English language (with native English speakers as participants).

3.3 Experiment 3

With Experiment 3, we aimed at providing an extended replication of our previous results by using novel behavioral descriptions during the learning phase of the false recognition paradigm that allowed for both, trait and state inferences to occur. Furthermore, all participants were probed for both, trait and state inferences during the recognition phase, thus avoiding potential systematic differences in processing mode that may have affected results based on the between-participants design of Experiments 1 and 2.

3.3.1 Method

3.3.1.1 Sample Size Determination

Our primary interest in this experiment was the main effect of Condition (implied vs. implied-other) for both, trait and state inferences. In order to provide sufficient statistical power ($1 - \beta = .80$) to detect a main effect size of $\eta_p^2 = .091$ (estimation based on the smallest effect size observed in our previous experiments) with $\alpha = .05$ in a 2 x 2 repeated-measures ANOVA, valid data of $N = 82$ participants were required. Because we aimed to additionally test for a potential interaction effect – thus directly comparing state and trait inference effects – we included it into our power analyses. For this interaction, we considered a small effect size of $\eta_p^2 = .022$ as the smallest effect size of interest. In order to provide sufficient statistical power ($1 - \beta = .80$) with $\alpha = .05$ for this effect size in a 2 x 2 interaction in a repeated-measures ANOVA, we planned to collect 352 valid data sets. Given an estimated exclusion rate of 5% with online data collection, we collected data from $N = 376$ participants. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

3.3.1.2 Participants

The current experiment relies on valid data from a total of $N = 365$ participants (109 female, average age $M = 35.6$, $SD = 12.3$, ranging from 18 to 75 years). The majority of participants indicated being native speakers of English (96%), 3% indicated to speak English as one of their native languages, and 1 participant indicated that they spoke English very well, albeit it was not their native language. Participants were recruited using Prolific (www.prolific.ac) and received monetary compensation of 1.50 GBP (approx. 1.88 USD) for the average experiment duration of nine minutes. Following our pre-registered exclusion criteria, we excluded data of six further participants because they had failed an attention test and two further participants because they self-reported insufficient language proficiency. Data of 90 further participants were excluded because they aborted the experiment before debriefing and thus did not give informed consent for data analyses.

3.3.1.3 Materials

We developed and tested a new set of dual-implication behavioral statements that were ambiguous in the sense that they implied both a trait and a state. We selected 24 statements as target stimuli along with 12 filler statements of a similar structure to the target stimuli, with the exception that they explicitly mentioned a trait or state word (see Tables A.10 and A.11 in the Appendix for the complete list of statements and pre-test results). Contrary to the previous experiments, stimuli were not designed to look like messages from an instant messaging service application but formulated in third person and paired with a portrait and a name (see Figure 3.3 for an example).

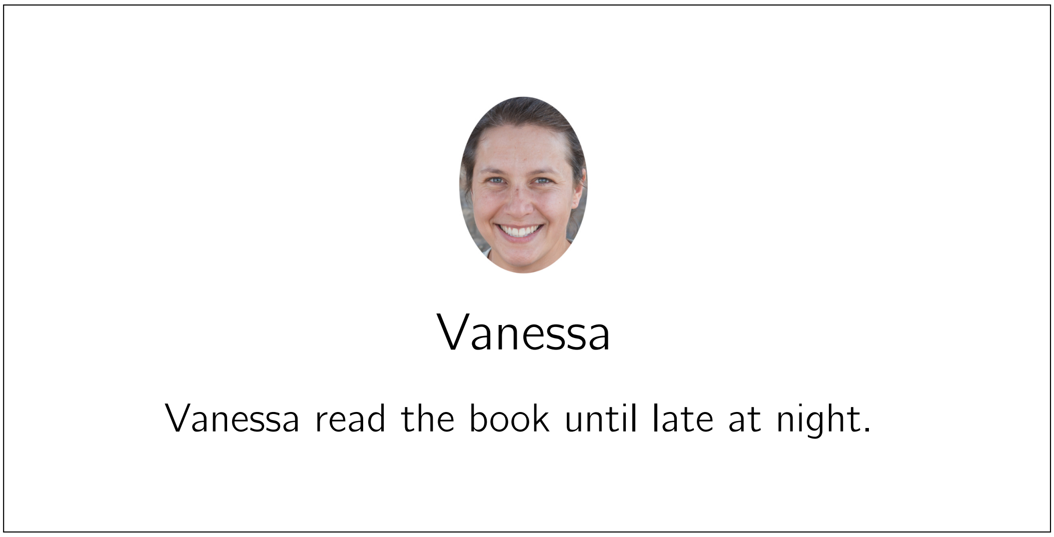

Fig. 3.3.: Example Stimulus (Experiment 3). Original portrait pictures differ. Printed portrait pictures were adapted from Karras et al. (2020), for copyright reasons.

3.3.1.4 Design

We employed a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject design with rate of false recognitions (answers 'yes' in test phase) as dependent variable. In Experiment 3, we omitted the new trait and state conditions, because the implied-other condition represents the more conservative test for the occurrence of trait and state inferences. We implemented an additional between-subjects factor based on stimulus-set assignment: Stimuli were randomly assigned to separate sets of six to be presented equally often in the implied and implied-other condition across participants (using a Latin square design), in individual random order.

3.3.1.5 Procedure

Experiment 3 followed the same procedure as Experiment 1, with the exception that participants self-reported language proficiency using a 6-point scale (1: "English is my first (native) language" to 6: "It is very hard for me to speak and understand English").

3.3.2 Results

False recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject ANOVA. Results show significant main effects for Inference, $F(1, 364) = 135.79$, $p < .001$, $\eta^2_G = .05$, 90% CI [.02, .09], and for Condition, $F(1, 364) = 43.89$, $p < .001$, $\eta^2_G = .02$, 90% CI [0, .04], with no significant interaction, $F(1, 364) = 1.64$, $p = .202$, $\eta^G_p < .01$, 90% CI [0, .01], see lower left panel of Figure 3.2. Separate analyses confirmed spontaneous person inference effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied trait condition ($M = .22$, $SD = .21$) than in the implied-other trait condition ($M = .18$, $SD = .20$), $t(364) = 4.009$, $p < .001$ (one-tailed), $d_z = 0.22$, 95% CI [0.11, 0.33]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition ($M = .34$, $SD = .24$) than in the implied-other state condition ($M = .27$, $SD = .22$), $t(364) = 5.292$, $p < .001$ (one-tailed), $d_z = 0.29$, 95% CI [0.18, 0.4]. Albeit effect sizes were somewhat larger in the state condition as compared to the trait condition, the inference effects (calculated as difference scores of implied – implied-other) did not differ significantly between conditions, $t(738) = 1.28$, $p = .202$, $d = 0.094$, 95% CI [-0.05, 0.24].

3.3.3 Discussion

Results of Experiment 3 replicate and extend the findings of Experiments 1 and 2 by providing first evidence that participants can draw both trait and state inferences from one and the same ambiguous behavior description. These results can, however, not yet be interpreted as indicators that trait and state inferences are *simultaneously* drawn when processing behavioral information (similar to the assumption of simultaneous trait and situation inferences in Krull & Ericson's model; Krull & Ericson, 1995). Indeed, given our specific implementation of the false recognition paradigm, the same results would be observed if some behavioral descriptions exclusively or dominantly triggered trait inferences and other stimuli exclusively or dominantly triggered state inferences. Alternatively, some participants may have systematically only drawn trait inferences and no state inferences and other participants may have systematically drawn state inferences and no trait inferences. Albeit our auxiliary analyses do not support these alternative accounts, it is desirable to provide a more direct empirical test of the simultaneity of trait and state inferences.

The current results also do not fully refute the assumption that trait inferences may be the dominant inference from behavior. For example, one may assume that participants spontaneously only draw trait inferences at encoding of the behavioral information but later (re)consider state inferences once probed with a fitting state word during the recognition test. However, note that the same argument applies to the opposing assumption that participants spontaneously only draw state inferences at encoding and later (re)consider trait inferences when probed with a fitting trait word (but see Todorov & Uleman, 2002). We conducted auxiliary analyses of response times to explore whether response times of false recognition responses differed between the implied state and trait conditions, presuming that false recognition responses based on inferences drawn at encoding may be faster as compared to responses after retrospect reconsiderations. However, there were no significant differences in false recognition response latencies between implied-state and implied-trait conditions (see Table A.12 in the Appendix). In Experiment 4, we used a more direct approach to investigate if and to what extent trait and state inferences are drawn *simultaneously* when participants process behavioral information about actors.

3.4 Experiment 4

In Experiment 4, we implemented a probe recognition paradigm (McKoon & Ratcliff, 1986) – another indirect paradigm frequently used in research on spontaneous trait and situation inferences (e.g., Ham & Vonk, 2003; Newman, 1991; Ramos et al., 2012). In this paradigm, participants read individual behavioral statements, each *immediately* followed by probe words. Participants are instructed to decide as quickly and accurately as possible whether the probes were part of the statement or not. In typical studies on STIs, these probe words are trait adjectives implied by the behavior or non-implied control probes, thus both requiring a negative response. If encoding of the behavioral statement automatically triggers a trait inference, participants should have more difficulty rejecting the implied trait probe and thus demonstrate slower response latencies and/or higher error rates as compared to non-implied control trait probes. One advantage of this paradigm is that each behavioral statement can be followed by several probes, thus allowing for testing multiple inferences referring to the same statement (e.g., Todd et al., 2011). In our adaptation of the paradigm, behavioral statements were followed by both, implied state and implied trait probes, which allows investigating the simultaneous occurrence of both types of inferences.

3.4.1 Method

3.4.1.1 Sample Size Determination

We followed the same pre-registered sample size rationale as in Experiment 3, thus planning to collect valid data from 352 participants.

3.4.1.2 Participants

The current experiment relies on valid data from a total of N = 341 participants (211 female, four other, one unspecified, average age M = 33.2, SD = 11.5, ranging from 18 to 71 years). The majority of participants indicated being native speakers of English (96%), 3% indicated to speak English as one of their native languages, and one participant indicated that they spoke English very well, albeit it was not their native language. Following our pre-registered exclusion criteria, we excluded the data of six further participants because they aborted the experiment before debriefing and thus did not give informed consent for data analysis, five because they self-reported insufficient language proficiency, and two because they responded accurately in less than 60% of trials. Participants were recruited via Prolific (www.prolific.ac) and received 2.20 GBP (approximately 2.75 USD) for the duration of 15 minutes.

3.4.1.3 Materials

We used a different subset of 24 behavioral statements developed for Experiment 3 as target statements and 24 filler statements explicitly mentioning a state or a trait word (see Tables A.10 and A.11 in the Appendix for the complete list of statements and pre-test results). Given the framing of the task as measuring automatic text comprehension, behavioral statements were presented on screen without images. Actor names were presented as part of the statements.

3.4.1.4 Design

We employed a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject design with response latencies of correct probe rejections of target trials as dependent variable.

3.4.1.5 Procedure

Participants completed the same introduction procedure and attention check as in the previous studies. The probe recognition paradigm was then introduced to participants as an experiment on language comprehension. The experiment was designed in PsychoPy 3.1.2 (Peirce et al., 2019) and conducted online on Pavlovia (www.pavlovia.org). Participants were presented with 48 behavioral statements (24 targets, 24 filler) in individual random order and were instructed to read them carefully. Each statement was presented for 3s. Immediately after each statement, participants completed eight probe recognition trials, in which they indicated for each probe word whether it had been part of the previous statement or not. Each probe was preceded by a blank screen (250ms) and a fixation cross (500ms) and remained on screen until a response was recorded. Participants were instructed to indicate via key press whether the word had appeared (press [D]) or not (press [K]) in the statement. Participants were instructed to respond as accurately and as fast as possible. Erroneous responses were signaled to the participants by a red cross displayed for 1000ms. After completion of all eight probe trials for each statement, the next statement was presented after an inter-trial interval of 500ms.

Target probe trials always consisted of the trait and the state adjectives implied by the behavioral statement of the same person (implied condition), as well as a trait and a state adjective implied by the behavioral statement of another person (implied-other condition), thus all requiring a "no" response. In order to balance the ratio of correct "yes" and "no" responses for each trial, we additionally presented four filler probes consisting of words that had actually appeared in the statement (i.e., names, objects, verbs, and prepositions). In order to avoid that participants recognized that any type of probe consistently required a "yes" or "no" response, probes for the filler statements were chosen such that correct responses were "yes" for adjectives and "no" for names, objects, verbs, and prepositions. For each trial, the order of the eight probes was individually randomized, with the restriction that the first probe was never a target probe (Stewart et al., 2004). Responses and response latencies were recorded. After completion of the probe recognition task, participants provided demographic information (age, gender, language proficiency, education, profession). Additionally, participants were asked to indicate how seriously they complied with task instructions using a 10-point scale (0 = not at all and 10 = very much, M = 9.5, SD = 0.8) and asked to speculate about the hypothesis of the experiment. At the end of the experiment, participants were fully debriefed about the purpose of the experiment and once again asked for consent for data storage and analyses.

3.4.2 Results

We had pre-registered response latencies of correct rejections as main dependent variable for our analyses. Analyses of response latencies usually require corrections of outlying slow responses (Ratcliff, 1993). To our knowledge, there is no convention how to correct for outliers in the probe recognition paradigm. Therefore, we applied different trimming criteria for slow responses (2500, 2000, 1500ms, individual $M \pm 2.5 * SD$), and transformations (log- and inverse transformation) and compared their impact on analyses. Across the different trimming criteria and transformations, effect sizes differed by small to medium amounts (trait inference effect: d_z = 0.580 - 0.806; state inference effect: d_z = 0.790 - 1.117; Inference x Condition interaction: η_p^2 = .03 - .07; see Table A.14 in the Appendix). Analyses reported in-text are based on log-transformed data with a cut-off of 1500 ms. For ease of understanding, descriptive statistics are based on untransformed but trimmed response latencies.

The pre-registered 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) repeated-measures ANOVA of response latencies for correct rejections of target probes documented significant main effects of Inference, $F(1, 340) = 224.09$, $p < .001$, $\eta^2_G = .01$, 90% CI [0, .04], and Condition, $F(1, 340) = 592.05$, $p < .001$, $\eta^2_G = .03$, 90% CI [.01, .07], that were qualified by a significant interaction, $F(1, 340) = 23.55$, $p < .001$, $\eta^2_G = .001$, 90% CI [0, .02], see lower right panel of Figure 3.2. Separate analyses confirmed expected effects in both inference conditions: Participants were slower to reject the implied trait probes ($M = 592$ ms, $SD = 122$) compared to the implied-other trait probes ($M = 557$ ms, $SD = 105$), $t(340) = 13.928$, $p < .001$ (one-tailed), $d_z = 0.286$, 95% CI [0.25, 0.33]. Similarly, participants were slower to reject the implied state probes ($M = 625$ ms, $SD = 115$) compared to the implied-other state probes ($M = 573$ ms, $SD = 115$), $t(340) = 19.502$, $p < .001$ (one-tailed), $d_z = 0.425$, 95% CI [0.38, 0.47]. The mean difference of implied state and implied-other state probes ($M = 50$ ms, $SD = 46$) was significantly larger than the mean difference of implied trait and implied-other trait probes ($M = 34$ ms, $SD = 43$), $t(680) = 4.469$, $p < .001$, $d_z = 0.342$, 95% CI [0.19, 0.49].

3.4.3 Discussion

The results of Experiment 4 replicate and complement the results of our previous studies in a different experimental paradigm. Again, we observed significant state and trait inference effects, with state effects being significantly larger than trait effects. The use of the probe recognition paradigm in Experiment 4 allows the conclusion that observers drew state and trait inferences both spontaneously and simultaneously when encoding the behavioral information.

Our research thus far strongly supports the notion that people draw spontaneous trait and state inferences when forming impressions from others' behaviors. However, the interpretability of these effects remains limited because the differentiation of trait and states in these studies entirely relies on our pilot experiment. In this experiment, we had asked participants to deliberately judge lists of adjectives with regard to their perceived temporal and situational stability after explicitly instructing them about our theoretical conceptualization of states and traits. We cannot be entirely sure that participants draw the same conceptual distinction between traits and states when they spontaneously infer them from behavioral statements. Albeit our pilot experiment clearly indicated that the employed adjectives were understood as either states or traits, several of these adjectives may still appear as possibly referring to both, states and traits - depending on the contexts in which they are used. For example, we may use the state adjective *sad* when considering a person to *feel sad* (state) or to be a *sad person* (trait). Because German and English have no clear linguistic markers that conceptually distinguish between traits and states (as in Spanish or Portuguese, for example), it is conceivable that the state inferences we observed in our studies are nothing but trait inferences in disguise. So far, we cannot conclude from our data that spontaneous state inferences actually differ functionally from trait inferences in the eyes of perceivers. We thus conducted a fifth experiment in order to investigate whether participants actually represent the pre-defined state and trait inferences as functionally different from each other.

3.5 Experiment 5

From a theoretical standpoint, traits and states can be distinguished from one another quite precisely (e.g., Hamaker et al., 2007; but see Fleeson and Jayawickreme, 2015). The most obvious difference is their relative stability: traits describe dispositions that are relatively stable over time and across situations, whereas states describe dispositions that are considerably less stable over time and across situations. This stability advantage of traits over states renders trait attributions appealing for impression formation because they may reduce uncertainty to a higher degree than state attributions - by increasing the predictability of other people's future behavior. Trait inferences should thus influence future behavior predictions, and have been demonstrated to do so (McCarthy & Skowronski, 2011b; Nussbaum et al., 2003). State inferences, on the other hand, describe variable, more fleeting properties of a person and should therefore have less influence on future behavior predictions. If people mentally represent this functional differentiation of trait and state inferences, they should rely on them differently when using observed behavior to make future behavioral predictions about others. In Experiment 5, we investigated whether traits and states do indeed hold differential value for predicting behavior. We therefore provided participants with the same behavioral statements from Experiments 3 and 4, either paired with the implied state inference or the implied trait inference, and asked them to judge the likelihood that the actors would show such behavior again in the future. We expected to observe trait inferences leading to a higher perceived probability of repeated behavior, as compared to state inferences.

3.5.1 Method

3.5.1.1 Sample Size Determination

Our primary interest in this experiment was the effect of Adjective (trait vs. state) in a t-test with prediction rating as the dependent variable. We aimed to provide sufficient statistical power ($1 - \beta = .80$) to detect the minimal effect size of interest of $d_s = .50$ for the between-subjects comparison of only the first response (see Analysis Plan). Thus, valid data of $N = 102$ participants was required. Power analyses were conducted using G*Power 3.1 (Faul et al., 2007).

3.5.1.2 Participants

The current experiment relies on valid data from a total of $N = 97$ participants (64 female, one other, average age $M = 33.7$, $SD = 11.9$, ranging from 18 to 66 years). Following our pre-registered exclusion criteria, we excluded the data of five further participants because they aborted the experiment before completion due to a programming error. Participants were recruited via Prolific (www.prolific.ac) and received 1.56 GBP (approximately 2 USD) for the duration of 10 minutes.

3.5.1.3 Materials

We used all 41 behavioral statements used in Experiments 3 and 4 as target statements, paired with a statement describing the actor with the implied state or trait adjectives, respectively (e.g., Marvin laughed at the joke. Marvin is jolly *vs.* Marvin is amused.). Behavioral statements were presented on screen without portrait pictures.

3.5.1.4 Procedure

Participants were introduced to an experiment on behavior prediction and memory performance. The experiment was designed and conducted in Qualtrics software. Participants were instructed to read the behavioral statements and judge how likely each actor would perform behavior such as the one described again in the future using a slider (0 = not at all likely to 100 = very likely). Additionally they were instructed to memorize actor, behavior, and adjective for a later memory test. We implemented the additional memory test in order to ensure that participants actually process the trait and state adjectives and do not base their judgements on the relative frequency of the described behaviors only. Participants completed the same attention check procedure with the first instruction page as in the previous experiments. They were then presented with 82 trials including each behavioral statement once paired with the implied state, once paired with the implied trait. Participants submitted their slider responses by pressing the space bar, which allowed us to record responses and response times.

In order to verify that participants had actually processed behaviors and state and trait adjectives, they completed ten recognition trials consisting of an actor's name and one of the adjectives from the rating phase. In five of these recognition trials, the adjectives had been presented with that same actor, and in the other five, the adjectives had been presented with a different actor. Participants were asked to judge whether actor and adjective had been presented together before (correct decisions: M = .67, SD = .18).

Finally, participants provided demographic information (age, gender, language proficiency, education, profession, ethnicity). Additionally, participants were asked to indicate how seriously they complied with task instructions using a 10-point scale (0 = not at all and 10 = very much, M = 8.9, SD = 1.4).

3.5.1.5 Analysis Plan

We had planned two ways of analysing the results of this Experiment. First, we planned a simple within-participants t-test in order to inspect whether participants attributed higher likelihood of behavior repetition when the behavioral statements were paired with the implied trait, as compared to the implied state adjective. However, because this within-participants design asks participants to estimate the repetition likelihood twice for the same behaviors, differences between implied state and trait adjectives may be under-estimated: Regardless of randomized sequence, participants may be inclined to base their second estimation on the response given to the same behavior earlier. To account for this possibility and rule out any form of cross-contamination, we planned a second analysis using only participants' first responses to each statement, discarding their second judgement.

3.5.2 Results

The within-participants analyses show that participants rated the probability for behavior repetition higher when the behavioral statement was paired with a trait adjective (M = 73.46 , SD = 11.4) than with a state adjective (M = 69.06 , SD = 10.69), $t(96) = 4.95$, $p < .001$ (one-tailed), $d_z = 0.398$, 95% CI [0.233, 0.563]. Additional analyses of the first judgements on each behavioral statement revealed a similar effect with a slightly larger effect size: Participants who saw the statement paired with a trait (M = 73.21, SD = 11.68) rated the probability of behavior repetition significantly higher than participants who saw the statements paired with the respective states (M = 68.44, SD = 10.48), $t(96) = 4.958$, $p < .001$ (one-tailed), $d_z = 0.427$, 95% CI [0.25, 0.605]. We observed the same pattern of results for the second presentation of statements with a slightly reduced effect size. Participants who saw the statement paired with a trait (M = 73.68, SD = 12.11) rated the probability of behavior repetition significantly higher than participants who saw the statements paired with the respective states (M = 69.64, SD = 12.17), $t(96) = 3.696$,

$p < .001$ (one-tailed), $d_z = 0.332$, 95% CI [0.15, 0.514]. We further conducted exploratory analyses in order to explore whether the relative strength of state and trait inferences in Experiments 3 and 4 were related to the aggregated repetition predictions in Experiment 5. We computed relative inference scores for each statement aggregated across participants as a simple difference between implied and implied-other responses. Because both experiments used different dependent variables, we z-transformed these inference scores and averaged the z-values for those statements that were used in Experiments 3 and 4. Simple bi-variate correlations indicate that the size of spontaneous trait inferences in Experiments 3 and 4 was related to the perceived likelihood of behavior repetition in Experiment 5, $r = .439$, $t(46) = 3.316$, $p = .002$, 95% CI [.177, .643], whereas the size of the state inference effects in Experiments 3 and 4 was not significantly related to the repetition predictions in Experiment 5, $r = .078$, $t(46) = 0.533$, $p = .596$, 95% CI [-.210, .355].

3.5.3 Discussion

Results of Experiment 5 support the assumption that participants processed trait and state inferences as indicating differential predictive value. Participants judged the likelihood that actors would repeat their behaviors in the future as higher when the behaviors were paired with the implied trait inference than the implied state inference. Albeit limited by their relatively low power, the exploratory by-item analyses further support this interpretation. The stronger the trait inference from a behavior, the more likely it seems that this behavior may be repeated in the future. No such relation was observed for state inferences drawn from the same behaviors. We are thus confident that the trait and state inferences measured in Experiments 3 and 4 do indeed represent functionally distinct inferences.

3.6 General Discussion

Previous research on impression formation from behavior has provided a vast amount of evidence that when observing others' behaviors, people spontaneously draw trait inferences (e.g., Moskowitz & Olcaysoy Okten, 2016; Uleman et al., 2012). In a series of five pre-registered experiments we provide consistent evidence that people spontaneously and simultaneously infer both traits and states from behavior.

In Experiments 1 and 2, we employed behavioral statements with single-implications of either a trait or a state in a false recognition paradigm. We observed significant inference effects for both trait and state adjectives, in that participants more frequently falsely recognized implied trait and state probes as having been previously mentioned in an actor's behavioral statement, as compared to a trait or state implied by a different actor's behavior, or a new trait or state adjective. We thus replicated the established spontaneous trait inference effect and, more importantly, provided first evidence that people can spontaneously infer states as well. In Experiments 3 and 4, we used dual-implication stimuli, that is, behavioral statements that allowed for both trait and state inferences. Again, we observed significant inference effects for both trait and state adjectives in a false recognition paradigm (Experiment 3) and in a probe recognition paradigm (Experiment 4). Results of Experiment 5 support the assumption that trait and state inferences are represented as functionally different: Participants rated actors as more likely to show similar behavior again in the future when the behavior was paired with the implied trait adjective, as compared to the implied state adjective.

Our experiments show that person inferences from others' behaviors include both, considerations of the actors' enduring traits as well as considerations of their more transient states. The most important contribution of our findings is that they show that state and trait inferences occur *simultaneously* at encoding of the behavioral information and are *mutually non-exclusive*. Specifically, in Experiments 3-5, we employed dual-implication stimuli that were carefully pretested to support both trait and state inferences in deliberate impression formation - and lead to both trait and state inference effects in both the false recognition and the probe recognition paradigm assessing effects of spontaneous impression in Experiments 3 and 4. To better understand the results of the false recognition effects in Experiment 3, we had conducted auxiliary correlational analyses both on the by-participant and the by-item level (see Appendix, p. A.6). Observing negative correlations on either level would have implied that trait and state inferences may be mutually exclusive; on the participant level, if participants who show strong trait inference effects show weak state inference effects or vice versa; on the item-level, if statements leading to strong trait-inference effects (aggregated across participants) lead to weak state inference effects or vice versa. The results indicate, however, that state and trait inferences were not drawn at the expense of each other: There was no indication that participants within our samples had individual response tendencies towards either spontaneous trait or spontaneous state inferences, nor that individual behavioral statements would only prompt either spontaneous state or spontaneous trait inferences. The strongest evidence for the mutual non-exclusiveness of simultaneously and spontaneously drawn state and trait inferences, however, stems from the probe recognition paradigm used in Experiment 4, in which we prompted participants with both implied trait and state adjectives immediately after encoding of the behavioral statements. Again, we observed both inferences effects and again, we did not find any negative correlations between state and trait inferences neither on the participant nor on the item level. Thus, we feel confident to conclude that spontaneous person inferences simultaneously include considerations about stable person dispositions as well as transient mental states and that these are mutually non-exclusive at the stage of spontaneous information processing.

Note that the semantic and pragmatic rules of the employed experimental paradigms and their instructions actually call for noninferences: Participants were never asked to form impressions of others – neither state nor trait impressions – but to merely process and memorize statements and images. Deliberate impression formation would have been a rather distracting activity during information encoding and also hindering during task performance, given that these inferences may increase false recognition rates and slow down responding. Thus, the indirectness of the inference effects – the increase of false recognition responses to implied states and traits in the false recognition paradigm and the slowing of correct rejection responses in the probe recognition paradigm – support the assumption that both state and trait inference effects are *spontaneous* both in the sense of *unintentional* as well as in the sense of *uncontrolled* (Uleman, Hon, et al., 1996). Whether state inferences are drawn with a comparable degree of automaticity as trait inferences is an important empirical distinction, which will be addressed further in Chapter 5.

Using actors' portraits for the assessment of inference effects in the recognition test of the false recognition paradigm further allows for the conclusion that the observed trait and state inference effects should not be interpreted as mere side-effects of text-comprehension, but as actor-specific person inferences. Furthermore, we always compared responses to the state and trait adjectives implied by the actors' behavior to trait and state adjectives implied by behaviors of other actors during the same learning phase. It is thus not mere familiarity or traces of prior activation of text-based associations of behaviors with trait or state words that increase false recognition rates (e.g., reading the verb *laughing* activating associations like *jolly* or *amused*), but inferences directly tied to the specific actors of that behavior (e.g., the laughing person *is* jolly and/or amused; see also Orghian et al., 2018). We further address this question in Chapter 4, in which we investigate to which degree spontaneous trait and state inferences represent actual inferential processes, as compared to mere associative processes.

4 Inferential and Associative Processes in Spontaneous Social Inferences

4.1 Inferential and Associative Processes in STIs

In Chapter 3, we have demonstrated that participants commit more errors (or take more time) when rejecting a probe word that was previously implied, but not mentioned in the actor's statement in a false or probe recognition task. We interpreted these results as evidence for the occurrence of both state and trait inferences. However, the results allow for an alternative interpretation. Bassili (1989) argued that STI-effects might merely reflect categorizations of behavior, instead of attributional inferences concerning the actor. Trait adjectives may serve as categorizations of both, persons and behavior, and categorizing a behavior does not necessarily imply that the person is categorized by that trait as well. If Marvin laughs at a joke, and *jolly* came to your mind, you might have inferred that the *Marvin is jolly*, but you might as well simply have identified laughing as a *jolly behavior*, without necessarily inferring that Marvin is jolly.

Participants' false recognitions of the probe word "jolly" are thus not necessarily evidence for the occurrence of spontaneous trait inferences about the actor. They might simply reflect an association between the actor and the word "jolly", which was activated by the presented behavior. Associations reflect generic, unlabeled linkages in memory as a result of incidental spatial and temporal contiguity of activated constructs (Carlston et al., 1995), for instance when a portrait is presented next to a trait-implying statement, like in evaluative conditioning research or attribute conditioning effects (e.g., Gawronski & Bodenhausen, 2018; Unkelbach & Förderer, 2018). Associations are shallow in the sense that they do not describe any relationship between the two contingent constructs, but merely a bi-directional link representing prior co-occurrence (cf. Gawronski & Bodenhausen, 2011). Inference processes do not result from simple associative links, as they involve three links: person-trait, behavior-trait, and person-behavior. These links are not symmetrical. People infer traits from behaviors more easily than the other way around (Maass et al., 2001), and the semantic links between traits and behaviors are causal (Kressel & Uleman, 2010, for a thorough discussion, see Orghian, Garcia-Marques, Uleman, & Heinke, 2015).

Bluntly said, if STIs merely reflect associative processes, then trait "inferences" should occur for pretty much anything you pair with a trait. Brown and Bassili (2002) illustrated this point with evidence for associations between traits and objects: the case of the *superstitious banana*. They presented participants with pairs of objects and traits and found effects similar to those observed for spontaneous trait inferences. These results can hardly reflect inferential processes, but suggest that associative processes might play a role for STI effects.

Similarly, Carlston et al. (1995) observed that participants would spontaneously infer traits about persons paired with statements in which they ostensibly describe not themselves, but other people - and thus no inference about the person would be warranted. Yet participants readily associated the trait implied by the description with the communicator in a savings-in-relearning paradigm, thus spontaneously transferring the trait to them. This phenomenon, spontaneous trait transference (STT), has since been thouroughly investigated (Carlston & Skowronski, 2005; Crawford, Skowronski, & Stiff, 2007; Crawford, Skowronski, Stiff, & Scherer, 2007; Goren & Todorov, 2009; McCarthy & Skowronski, 2011b; Orghian et al., 2015; Skowronski et al., 1998; Todorov & Uleman, 2004), and typically yields about half the effect size of STIs. Carlston & Skowronski (1995) suggested that associative processes account for some of the variance of STI. The fact that STI effect sizes are typically larger than the ones observed for STT suggests that they reflect more than mere associations, typically interpreted as inferential processes (but see Orghian et al., 2015). Skowronski et al. (1998) suggested that different processes underlie STI and STT. In the case of STI, when perceivers are presented with trait-implying behavior of an actor, they activate the trait implied by the behavior and spontaneously generate a trait inference which is linked to the mental representation of the actor. When probed later with the implied trait, they simply access the inference. In the case of STT, perceivers also activate the trait when presented with an informant's behavior. But the activated trait is not linked inferentially to the mental representation of the actor. Instead, both concepts are represented separately in memory. When participants are then later probed with the implied trait, the spatial and temporal contiguity of trait and informant may cause similar effects as in the case of STI, without the occurrence of an actual inference.

The evidence we have provided so far does not prove spontaneous state inferences to be actual social inferences. We have merely shown that participants associated state adjectives with actors to a similar degree as trait adjectives, and interpreted this pattern of results as indicating that inferential processes had occurred. However, trait and state inferences might differ with respect to the relative contributions of inferential and associative processes to their emergence. As the evidence regarding spontaneous trait transferences as discussed above suggests, spontaneous trait inference effects are likely comprised of both associative and inferential components: Actors are associated with the trait implied by their statement, and that association is also inferentially linked to them. We observed spontaneous state inferences with about the same effect sizes as trait inferences, but yet it is conceivable that spontaneous state inference effects reflect only associative processes. If we assume that states are more readily associated with actors and their behaviors, possibly because perceivers are more used to recognizing others' mental states than inferring their personality traits, then associative processes might account for a larger portion of the observed effects thus far interpreted as evidence for spontaneous state inference, as compared to spontaneous trait inferences. To investigate this important distinction, we conducted two experiments to assess both inferences and transferences, for both traits and states, using the false recognition paradigm. In both experiments, we observed larger effect sizes for inference effects as compared to transference effects, indicating that both, spontaneous trait and state inferences, at least partially reflect results of inferential processes.

4.2 Experiment 6

In Experiment 6, we investigated whether participants would infer traits and states implied by behavioral descriptions with portraits, even if they were instructed that portraits and behaviors are randomly paired (transference condition), and compared effect sizes to inference effects. Participants were presented with a person's name and their portrait, each paired with a dual-implication statement (Goren & Todorov, 2009, Experiment 1; see Fig. 4.1 for an example stimulus), and asked to memorize the information. The statements were printed in either green or red font, indicating whether they were relevant to that person (actor condition) or not (communicator condition). Later, participants were presented with the implied states and traits, paired with either the actor or the communicator. If participants inferred the trait or state implied by the statement relevant to that actor, pairing the actor with the implied trait or state should lead to an increased false recognition rate. If a person merely communicated a trait- or state-implying behavior that another person had performed (and which thus should not be relevant to the communicator), participants may still associate the communicator with the implied characteristic. This association, in turn, should lead to an increased false recognition rate. If inferences reflect more than mere associations, the effect size for false recognition rates in the communicator condition should be reduced, as compared to the actor condition.

4.2.1 Method

4.2.1.1 Sample Size Determination

Our primary interest in this experiment were the 2 x 2 within-subjects interaction effects between Information (actor vs. communicator) and Condition (implied vs. implied-other) in a repeated-measures ANOVA for both between-participant conditions of Inference (trait vs. state). In order to provide sufficient statistical power ($1 - \beta = .80$) with $\alpha = .05$ (one-tailed) to detect a minimal effect size of interest of $\eta_p^2 = .022$ for the 2 x 2 x 2 interaction between Information (relevant vs. irrelevant) x Condition (implied vs. implied-other) x Inference (trait vs. state). Thus, valid data of $N = 352$ participants were required. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

4.2.1.2 Participants

The current experiment relies on valid data from a total of $N = 354$ participants (120 female, average age $M = 33$, $SD = 12.4$, ranging from 18 to 76 years). Participants were recruited via Prolific (www.prolific.ac) and received monetary compensation of 1.56 GBP (approx. 1.98 USD) for the duration of 12 minutes. Participants were randomly assigned to either the trait condition ($n = 177$) or the state condition ($n = 177$). Following our pre-registered exclusion criteria, we excluded data from further participants: Eighteen because they did not pass an initial attention test, two because they revoked their consent for data analysis after being debriefed, and seven because they self-reported insufficient language proficiency.

4.2.1.3 Materials

We selected 24 target stimuli consisting of a behavioral statement, a portrait (selected from the 10k US adult portraits database; Bainbridge, Isola, & Oliva, 2013), and a given name, see Figure 4.1 for an example. All behavioral statements were dual-implication: they implied, but not explicitly mentioned both a state and a trait. Additionally, we used 12 filler stimuli equal to the target stimuli, with the exception that they explicitly mentioned a trait or

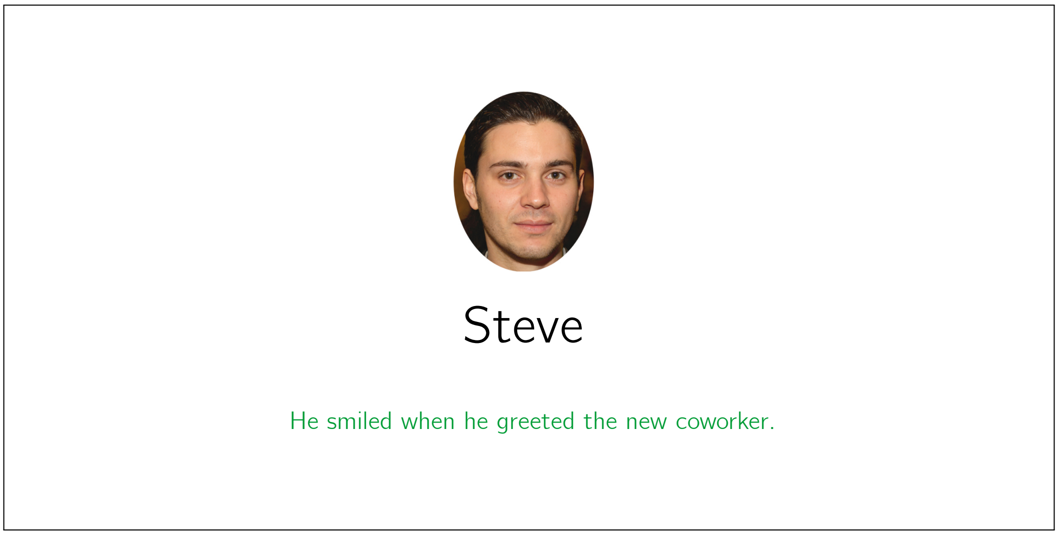

Fig. 4.1.: Example stimulus used in Experiment 6. The green font indicates that the statement is about the person depicted in the portrait. Original portrait pictures differ. Printed portrait pictures were adapted from Karras et al. (2020), for copyright reasons.

state word (see Table A.15 in the Appendix for the complete list of statements and pre-test results).

4.2.1.4 Procedure

Participants took part in a variant of the false recognition paradigm (Goren & Todorov, 2009), which was introduced as an experiment on memory. The experiment was designed and conducted in Qualtrics software. During the initial learning phase, participants were presented with the 36 stimuli. Participants were presented with 36 trials in which single stimuli appeared for six seconds each. Twenty-four of these trials were target stimuli and 12 trials were filler stimuli. Half of the target statements were printed in green font, the other half in red font. Participants were informed that statements in green font were relevant to the person depicted (i.e., represent their behavior; actor condition), whereas statements in red font were randomly paired with people and thus irrelevant (i.e., represent someone else's behavior; communicator condition; Goren & Torodov, 2009, Experiment 1). Half of the irrelevant statements matched the respective portrait's gender, the other half were of opposite gender[1]. The instructions were further clarified in two practice trials. Participants were instructed to read the stimuli carefully in preparation of a memory test later in the experiment.

[1]Exploratory analyses demonstrated that gender (mis)match did not significantly qualify the reported results. We therefore collapsed analyses across this factor.

Immediately after the learning phase, participants completed one practice recognition trial plus 36 experimental recognition trials, in which the portraits and names from the learning phase were presented together with an adjective probe. Participants were instructed to indicate via key press whether the probe word had appeared (press [A]) or not (press [K]) in the behavioral statement previously paired with the same person. Responses and response latencies were recorded. All target trials required a *no* response because the used probes had never been presented during the learning phase. We implemented a between-participants manipulation of probe-type for the target trials in that participants were presented with either only trait adjective probes or only state adjective probes. In six of the target trials, the portraits were paired with the adjectives (state or trait) that were implied by the behavioral statements printed in green font corresponding to those same portraits during the experiment phase (implied - relevant trials, used to measure trait or state inferences) – in this subset, a *yes* response indicates false recognition of a corresponding implied trait or state. In another six target trials, portraits were paired with the adjectives (state or trait) implied by the behavioral statements printed in red font corresponding to those same portraits (implied - irrelevant trials, used to measure trait or state transferences) – in this subset, a *yes* response indicates false recognition of a non-corresponding implied trait or state. In six of the control trials, portraits were paired with adjectives (trait or state) implied by behavioral statements printed in green font corresponding to other experiment trials (implied other - relevant trials, used as a control condition for trait or state inferences), and in another six control trials, portraits were paired with traits implied by behavioral statements printed in red font corresponding to other experiment trials (implied other - irrelevant trials, used as a control condition for trait or state transferences). Finally, in the 12 filler trials, portraits were presented with the adjectives that had previously appeared in the behavioral statement, thus requiring a "yes" response. After completion of the recognition task, participants were asked to provide demographic information (age, gender, language proficiency, education, profession). Additionally, participants were asked to write what they think the experiment was about and to indicate how seriously they complied with task instructions using a ten-point scale (0 = not at all and 10 = very much). At the end of the experiment, participants were fully debriefed about the purpose of the experiment and once again asked for consent for data storage and analyses.

4.2.2 Results

Individual false recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) x 2 (Information: actor vs. communicator) ANOVA with the factor Inference varying between participants and Information and Condition varying within participants. Results showed significant main effects for Condition, $F(1{,}350) = 17.471$, $p < .001$, $\eta^2_G = 0.024$, 90% CI [.005, .058] and for Inference, $F(1{,}350) = 30.607$, $p < .001$, $\eta^2_G = 0.02$, 90% CI [.003, .051]. The interaction between Condition and Information, $F(1{,}350) = 5.348$, $p = .021$, $\eta^2_G = 0.004$, 90% CI [.000, .022], and between Inference and Condition were significant, $F(1{,}350) = 14.884$, $p < .001$, $\eta^2_G = 0.011$, 90% CI [.000, .036], while the interaction between Inference and Information, $F(1{,}350) = 0.116$, $p = .734$, $\eta^2_G = 0$, 90% CI [.000, .006], as well as the triple-interaction between Inference, Condition, and Information $F(1{,}350) = 2.204$, $p = .139$, $\eta^2_G = 0.002$, 90% CI [.000, .016] were non-significant.

False recognition rates for both inference conditions were submitted to separate 2 (Condition: implied vs. implied-other) x 2 (Information: actor vs. communicator) ANOVAs. In the trait condition, we observed significant main effects for Condition, $F(1{,}170) = 27.088$, $p < .001$, $\eta^2_G = 0.021$, 90% CI [.000, .071], and Information, $F(1{,}170) = 5.899$, $p = .016$, $\eta^2_G = 0.006$, 90% CI [.000, .041]. However, the interaction effect was non-significant, $F(1{,}170) = 0.298$, $p = .586$, $\eta^2_G = 0$, 90% CI [.000, .015]. Separate analyses revealed that in the trait inference condition, participants showed higher false recognition rates in the implied condition ($M = .33$, $SD = .24$) than in the implied-other trait condition ($M = .26$, $SD = .21$), $t(170) = 3.767$, $p < .001$ (one-tailed), $d_z = 0.318$, 95% CI [0.148, 0.489]. In the transference condition, participants also showed higher false recognition rates in the implied condition ($M = .29$, $SD = .23$) than in the implied-other trait condition ($M = .23$, $SD = .20$), $t(170) = 3.089$, $p = .001$ (one-tailed), $d_z = 0.267$, 95% CI [0.094, 0.441]. However, contrary to prior research (e.g., Crawford, Skowronski, Stiff, & Scherer, 2007; Skowronski et al., 1998), the transference effect did not differ significantly from the inference effect, $t(170) = 0.546$, $p = .293$ (one-tailed), $d_z = 0.063$, 95% CI [-0.164, 0.289].

In the state condition, we observed significant main effects for Condition, $F(1{,}180) = 7.992$, $p = .005$, $\eta^2_G = 0.007$, 90% CI [.000, .042], and Information, $F(1{,}180) = 9.138$, $p = .003$, $\eta^2_G = 0.008$, 90% CI [.000, .044], which were qualified by a significant interaction, $F(1{,}180) = 7.292$, $p = .008$, $\eta^2_G = 0.007$, 90% CI [.000, .040]. Separate analyses confirmed expected effects: Participants showed higher false recognition rates in the implied

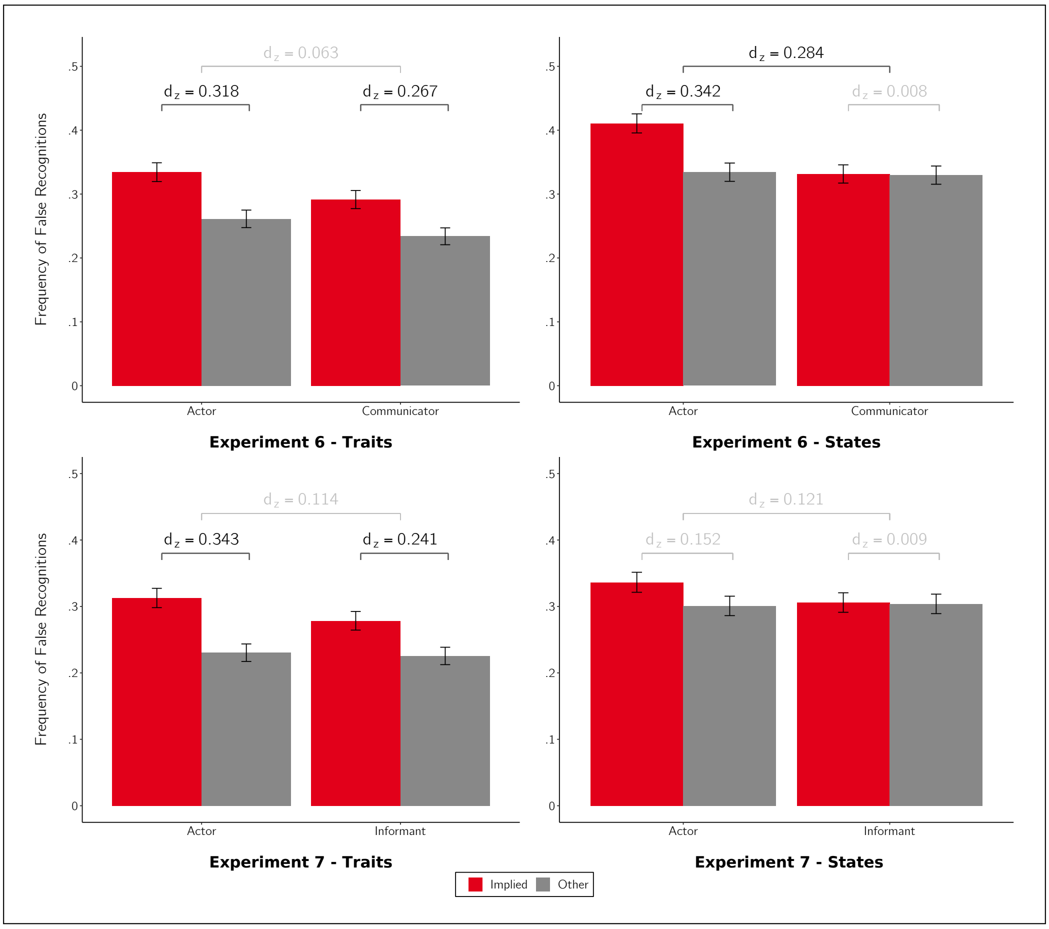

Fig. 4.2.: Mean proportion of false recognition responses for Experiments 6 ($N = 354$) and 7 ($N = 334$) as a function of type of inference and experimental condition with effect sizes Cohen's d_z from pairwise comparisons. Effect sizes printed in grey were non-significant. Error bars represent *SE*.

state condition ($M = .41$, $SD = .21$) than in the implied-other state condition ($M = .33$, $SD = .23$), $t(180) = 3.881$, $p < .001$ (one-tailed), $d_z = 0.342$, 95% CI [0.164, 0.521]. In the transference condition, participants' false recognition rates did not differ between implied state condition ($M = .33$, $SD = .22$) and implied-other state condition ($M = .33$, $SD = .24$), $t(180) = 0.095$, $p = .462$ (one-tailed), $d_z = 0.008$, 95% CI [-0.155, 0.171]. The transference effect differed significantly from the inference effect, $t(180) = 2.7$, $p = .004$ (one-tailed), $d_z = 0.284$, 95% CI [0.073, 0.494].

4.2.3 Discussion

Results of Experiment 6 replicate the pattern of results of Experiments 3 and 4: Participants spontaneously inferred both traits and states from dual-implication stimulus materials. More importantly, however, in the transference condition, in which there was no logical basis for drawing inferences, we observed reduced effect sizes for both traits and states.

For trait inferences, we observed significant inference and transference effects. Effect sizes in the transference condition were reduced, as compared to the inference condition. Previous research has found that effect sizes for STT are usually reduced by half, as compared to STI effects (e.g., Carlston & Skowronski, 2005; Skowronski et al., 1998). We did not observe this pattern in Experiment 6: effect sizes for trait transferences ($d_z = 0.267$) were almost as large as effect sizes for trait inferences ($d_z = 0.318$), and more importantly, did not differ significantly. This pattern of results indicates that a substantial proportion of spontaneous trait inferences may actually reflect associative processes, but this conclusion cannot be drawn from a single experiment. If this result replicates, it warrants further exploration - see Experiment 7.

For state inferences, we observed a significant inference effect and a null-effect for transferences. While the experiment is underpowered to detect a null-effect with sufficient certainty (note that the 95 %CI for the state transference effect was as large as [-0.155, 0.171]), this pattern of results still suggests that state inference effects reflect actual inferential processes to a significantly larger degree than mere associations.

While the pattern of results observed for state inferences and transferences confirmed our hypotheses, the results observed for trait inferences warrant further investigation. We did not replicate the results typically observed in previous research, namely a significantly reduced transference effect, as compared to the inference effect. We thus decided to conduct another experiment to investigate the phenomenon, using a manipulation that should lead to a stronger reduction in trait transference effects.

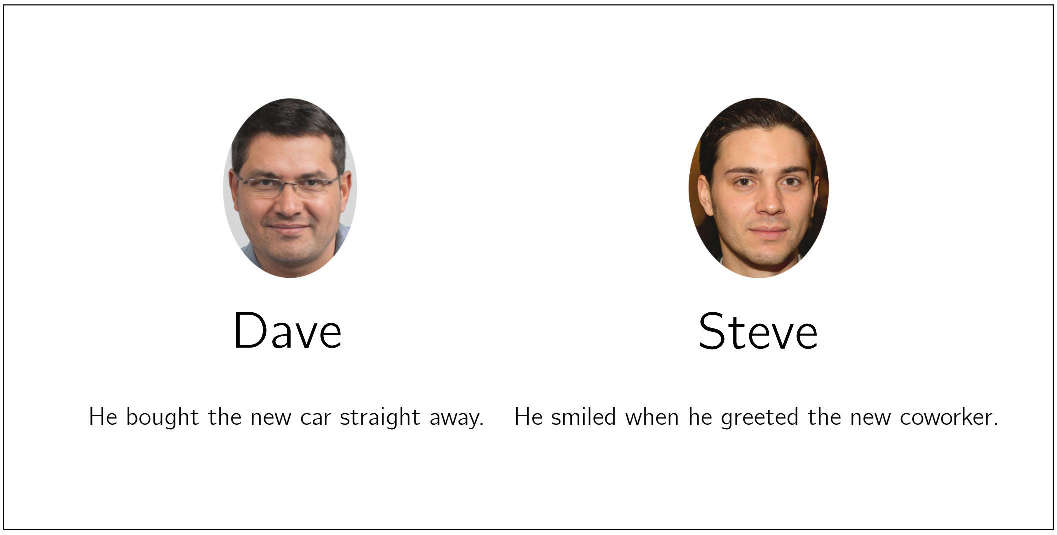

Fig. 4.3.: Example stimulus used in Experiment 7. Dave's statement refers to Steve, and Steve's statement refers to Dave. Original portrait pictures differ. Printed portrait pictures were adapted from Karras et al. (2020), for copyright reasons.

4.3 Experiment 7

In Experiment 7, we employed a manipulation that has been shown to eliminate spontaneous trait transference effects altogether (Goren & Todorov, 2009, Experiment 4). In a false recognition paradigm, participants were presented with two actors on the same screen, each paired with a statement about the respective other person (see Fig.4.3 for an example stimulus), and asked to memorize the information. Later, participants were presented with the implied states and traits, paired with either the actor or the communicator. If participants inferred the trait or state implied by the statement about that actor, pairing the actor with the implied trait or state should lead to an increased false recognition rate. If a person merely communicated a trait- or state-implying behavior that another person had performed (and which thus should not be relevant to the communicator), participants may still associate the communicator with the implied characteristic. This association, in turn, should lead to an increased false recognition rate. If inferences reflect more than mere associations, the effect size for false recognition rates in the communicator condition should be reduced, as compared to the actor condition.

4.3.1 Method

4.3.1.1 Sample Size Determination

Our primary interest in this experiment were the 2 x 2 within-subjects interaction effects between Information (Actor vs. Communicator) and Condition (Implied vs. Implied-Other) in a repeated-measures ANOVA for both between conditions of Inference (Trait vs. State). Goren and Todorov (2009, Experiment 4) reported a significant interaction effect size of η_p^2 = .094 for traits. To safeguard against imprecise power estimates (Perugini et al., 2014), we aimed to provide sufficient statistical power (1 - β = .80) with α = .05 to detect the lower-bound estimate η_p^2 = .046 of a 60% CI of the effect size. Thus, valid data of n = 168 participants were required for both the state- and the trait-condition, respectively. Assuming an exclusion rate of 5%, we planned to collect data from 352 participants. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

4.3.1.2 Participants

The current experiment relies on valid data from a total of N = 334 participants (123 female, average age M = 36.4, SD = 12.9, ranging from 18 to 71 years). Participants were recruited using Prolific (www.prolific.ac) and received monetary compensation of 1.56 GBP (approx. 1.88 USD) for the average experiment duration of 12 minutes. Following our pre-registered exclusion criteria, we excluded data of 12 participants because they had failed an attention test and eight further participants because they self-reported insufficient language proficiency. Data of four further participants were excluded because they retracted their consent for data analysis after debriefing.

4.3.1.3 Materials

We selected 24 target and 12 filler dual-implication stimuli from the materials generated for Experiments 3 and 4 (See Table A.15 in the Appendix for the complete list of statements and pre-test results).

4.3.1.4 Design

We employed a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) x 2 (Information: actor vs. communicator) mixed design, with the factors information and condition varying within participants and the factor inference varying between participants, with rate of false recognitions (answers *yes* in test phase) as dependent variable. We implemented an additional between participants factor based on stimulus-set assignment: Stimuli were randomly assigned to separate sets of six to be presented equally often in the implied and implied-other condition across participants (using a Latin square design).

4.3.1.5 Procedure

Experiment 7 followed the same procedure as Experiment 6, with the following differences: During the initial learning phase, participants were presented with the 18 trials in which two stimuli appeared simultaneously, horizontally aligned, for twelve seconds. Twelve of these trials contained two target stimuli, and six trials contained two filler stimuli, respectively. Participants were informed that within each pair of presented individuals, both were communicators regarding the other's behavior (Goren & Todorov, 2009, Experiment 4). This instruction was further clarified in two practice trials. In six of the target trials, the portraits were paired with the trait or state adjectives that had been implied by the behaviors corresponding to those same portraits (i.e., actor) during the experiment phase (implied-same actor trials, used to measure trait or state inferences) – in this subset, a *yes* response indicates false recognition of a corresponding implied trait or state. In another six target trials, portraits were paired with the trait or state adjective implied by the behavior corresponding to the other portraits (i.e., communicator) that had been shown simultaneously on screen during the experiment phase (implied-other actor trials, used to measure trait or state transferences) – in this subset, a *yes* response indicates false recognition of a non-corresponding implied trait or state. In six of the control trials, portraits were paired with traits or states corresponding to portraits that had been presented in the same position on screen in other experiment trials (implied other - same actor trials, used as a control condition for trait or state inferences), and in another six control trials, portraits were paired with traits corresponding to other portraits that had been presented in the opposite position on screen in other experiment trials (implied other - other actor trials, used as a control condition for trait or state transferences). Finally, in the 12 filler trials, portraits were presented with the adjectives that had previously appeared in the behavioral statement, thus requiring a *yes* response.

4.3.2 Results

Individual false recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) x 2 (Information: actor vs. communicator) ANOVA with the factor Inference varying between participants and Information and Condition varying within participants. Results showed significant main effects for Condition, $F(1,332) = 18.9$, $p < .001$, $\eta^2_G = 0.013$, 90% CI [.000, .041], and for Inference, $F(1,332) = 7.397$, $p = .007$, $\eta^2_G = 0.011$, 90% CI [.000, .037]. The interaction between Inference and Condition was significant, $F(1,332) = 5.912$, $p = .016$, $\eta^2_G = 0.004$, 90% CI [.000, .023], while the interactions between Condition and Information $F(1,332) = 2.198$, $p = .139$, $\eta^2_G = 0.002$, 90% CI [.000, .017], between Inference and Information, $F(1,332) = 0.097$, $p = .756$, $\eta^2_G = 0$, 90% CI [.000, .005], as well as the triple-interaction between Inference, Condition, and Information $F(1,332) = 0.009$, $p = .923$, $\eta^2_G = 0$, 90% CI [.000, .000], were non-significant.

False recognition rates for both inference conditions were submitted to separate 2 (Condition: implied vs. implied-other) x 2 (Information: actor vs. communicator) ANOVA. In the trait condition, we observed significant main effects for Condition, $F(1,169) = 25.176$, $p < .001$, $\eta^2_G = 0.021$, 90% CI [.000, .071]. However, the main effect for Information was non-significant, $F(1,169) = 2.13$, $p = .146$, $\eta^2_G = 0.002$, 90% CI [.000, .027], neither was the interaction effect, $F(1,169) = 1.042$, $p = .309$, $\eta^2_G = 0.001$, 90% CI [.000, .023]. Separate analyses revealed that in the trait inference condition, participants showed higher false recognition rates in the implied condition ($M = .31$, $SD = .24$) than in the implied-other trait condition ($M = .23$, $SD = .24$), $t(169) = 4.065$, $p < .001$ (one-tailed), $d_z = 0.343$, 95% CI [0.172, 0.514]. In the transference condition, participants also showed higher false recognition rates in the implied condition ($M = .28$, $SD = .23$) than in the implied-other trait condition ($M = .22$, $SD = .20$), $t(169) = 2.759$, $p = .003$ (one-tailed), $d_z = 0.241$, 95% CI [0.067, 0.416]. Importantly, the transference effect did not differ significantly from the inference effect, $t(163) = 1.075$, $p = .284$, $d_z = 0.121$, 95% CI [-0.101, 0.344], see lower left panel of Fig. 4.2.

In the state condition, neither the main effect for Condition, $F(1,163) = 1.585$, $p = .21$, $\eta^2_G = 0.002$, 90% CI [.000, .027], nor for Information, $F(1,163) = 1.062$, $p = .304$, $\eta^2_G = 0.001$, 90% CI [.000, .022], nor their interaction were significant, $F(1,163) = 1.156$, $p = .284$, $\eta^2_G = 0.001$, 90% CI [.000, .025]. Participants did not show higher false recognition rates in the implied state condition ($M = .34$, $SD = .23$) than in the implied-other state condition (M = .30, SD = .24), $t(163) = 1.635$, $p = .052$ (one-tailed), $d_z = 0.152$, 95% CI [-0.032, 0.336]. In the transference condition, participants' false recognition rates also did not differ between implied state condition (M = .31, SD = .22) and implied-other state condition (M = .30, SD = .24), $t(163) = 0.095$, $p = .462$ (one-tailed), $d_z = 0.009$, 95% CI [-0.173, 0.19], see lower right panel of Fig. 4.2.

4.3.3 Discussion

Results of Experiment 7 only replicate some of the results of Experiment 6: Participants spontaneously inferred traits from dual-implication stimulus materials, but the state inference effect did not reach statistical significance. More importantly, however, in the transference condition, in which there was no logical basis for drawing inferences, we observed reduces effects for traits and a null-effect for states.

For trait inferences, we observed significant inference and transference effects. Effect sizes in the transference condition were reduced, as compared to the inference condition. Previous research has found that effect sizes for STT are usually reduced by half, as compared to STI effects (e.g., Carlston & Skowronski, 2005; Skowronski et al., 1998). We did not observe this pattern in Experiment 7: Effect sizes for transferences ($d_z = 0.241$) were almost as large as effect sizes for inferences ($d_z = 0.343$), and more importantly, did not differ significantly. This pattern of results indicates that a substantial proportion of spontaneous trait inferences may actually reflect associative processes.

For state inferences, the inference effect did not reach statistical significance. In the state transference condition, we again observed a null-effect, with the same limitation of insufficient statistical power to detect an actual null-effect (the 95% CI for the state transference condition was as large as [-0.173, 0.19]).

While we closely followed the methodology described in previous research (Goren & Todorov, 2009, Experiment 4), there are some differences. Most importantly, the dual-implication stimulus materials we employed are more complex than the frequently used single-implication stimuli. As our materials imply not only a trait, but at the same time a state, cognitive demand on perceivers may be increased, as compared to materials that unambiguously imply only a trait (see Chapter 2.3, for a detailed discussion). Additionally, participants had to process multiple pieces of information in each trial: Two portraits, two names, two behavioral statements, each implying both a trait and a state, and finally, which information referred to which actor. The high degree of complexity may in part explain why effect sizes were reduced, and why the differences between inference and transference conditions were non-significant.

4.4 General Discussion

The two experiments we conducted provided some indication that spontaneous state inferences are most likely inferences, and not mere associations. In fact, the evidence for state inferences was quite strong, as we observed a null-effect for state transferences in both experiments. While the experiments were underpowered to observe an actual null-effect, the evidence does not contradict the hypothesis that state inferences reflect inferential processes. We are thus confident in assuming that spontaneous state inferences are the product of inferential, as compared to associative processes. It even seems that participants did not associate states with communicators at all. Given the readiness that perceivers show in associating traits with communicators, bystanders, and even objects (Bassili, 1989), it is particularly remarkable that participants did not appear to form such associations for states. It seems that spontaneous state inferences are bound specifically to the actor about whom the inference has been drawn, without being associated with communicators - notably less so than spontaneous trait inferences.

Interestingly, the evidence for trait inferences is less conclusive: We observed trait inference effect sizes that were only slightly and non-significantly larger than trait transference effect sizes. The effect size difference was consistently smaller than in previous research across the two experiments we conducted. Trait transferences were also not eliminated in Experiment 7, as in previous research (Goren & Todorov, 2009, Experiment 4). Nevertheless, we did observe a tendency for trait inference effects to be larger than transference effects, albeit non-significantly. Notable differences between our experiments and previous research, that could account for the deviations are, of course, the use of dual-implication materials that implied both a trait and a state simultaneously. Our materials are more complex than the single-implication trait-implying materials used in previous research, which might affect participants' processing of the behavioral descriptions. One may speculate that multiple-implication materials require more cognitive effort than single-implication materials. When participants process multiple possible implications implied by the descriptions, this might bind some of their resources, which in turn might reduce their capacity to draw inferences. This might explain the non-significant difference between the trait inference and trait transference effects, as the inference effect may be reduced due to the complex materials. This interpretation is consequential for the assumed automaticity of spontaneous trait inferences: If increased working memory load reduces spontaneous trait inferences, they might not be as efficient as previously assumed.

It is noteworthy that this seemed not to be true for states, for which we observed a significant difference between inference and transference effects in Experiment 6. The state inference effect was, however, reduced in Experiment 7, in which the presentation of two actors informing about each other at a time might have induced higher load on participants' working memory. We did not systematically manipulate working memory load in the two experiments reported in this chapter, but the need to assess its influence has nonetheless become apparent. In Chapter 5, we report a series of experiments in which we manipulated working memory load and its effects on spontaneous trait and state inferences, to further assess the influence of working memory load on spontaneous trait and state inferences.

4.4.1 Directions for Future Research

The two experiments reported in this chapter are a first step at investigating the inferential vs. associative nature of spontaneous state inferences, and cannot yet provide definite answers to this question. In order to strengthen the evidence and further our understanding of the processes underlying the formation of both spontaneous trait and state inferences, we have devised a series of experiments.

In an experiment that we are currently conducting, we investigate if or to what extent the pattern of results observed in Experiments 6 and 7 can be replicated with single-implication trait and state materials in the false recognition paradigm with the transference manipulation employed in Experiment 6 (red and green statements). This approach allows for eliminating the potentially confounding factor of working memory load induced by the dual-implication materials.

Another experiment in preparation aims at investigating transference effects using a different experimental paradigm. In order to be able to draw reliable conclusions, it is important to rule out possible methodological artefacts. Converging evidence from a multi-method approach would increase our confidence in the results. We opted to adapt a lie-detection manipulation in combination with the savings-in-relearning paradigm (Crawford, Skowronski, Stiff, & Scherer, 2007). In this approach, only self-informant behavioral statements are presented. Participants' processing goals during behavior presentation are manipulated: In one condition, participants are instructed to familiarize themselves with the trait- and state-implying materials. In this condition, spontaneous inferences should occur. In another condition, participants are instructed to judge whether the actors are lying about their behavior or not. This is assumed to trigger a subjective validation vs. invalidation of behavioral statements. The rationale behind this manipulation is that the subjective truth-value should moderate inferential processing by providing the alternative inference *liar*, interfering with inferences implied by the behavior. Associative effects should remain largely unaffected by subjective truth-value (e.g., De Houwer et al., 2020; Gawronski & Bodenhausen, 2018).

Both a successful replication and a failure to replicate our previous experiments are of high theoretical interest: If the results replicate with single-implication materials, the evidence for differences in the processes underlying spontaneous state and trait inferences are strengthened. Pending further replication and confirmation using different experimental paradigms and materials, of course, the combined evidence would then suggest that the proportion of inferential processes involved in state inferences is actually larger than for trait inferences. This, in turn, could suggest that there is a hierarchy or sequential processing involved in spontaneous social inferences.

Another prospect for further research is warranted in a more general way. Thus far, we have interpreted the increased effect sizes in the inference conditions, as compared to transference conditions (in which there was no logical basis for spontaneous social inferences) as evidence for an inferential process that goes beyond mere associations. However, Orghian et al. (2015) conducted simulations using a simple model based solely on associative learning (model of associative trait inferences and transferences; MATIT). They demonstrated that MATIT was able to produce the results commonly observed in experiments on spontaneous social transferences using only associative learning, suggesting that the evidence available so far is not sufficient to support a dual-process view of spontaneous social inferences. The MATIT highlights the importance of further investigation of the underlying processes of spontaneous social inferences and transferences, and whether or to what extent they rely on inferential and/or associative processes. For the purposes of this work, we did not address these important concerns, but limited ourselves to investigating whether the processes underlying spontaneous state inferences are at least as inferential (or associative) as the ones underlying spontaneous trait inferences. And this is precisely what we observed, furthering the evidence that spontaneous state inferences are comparable to spontaneous trait inferences.

To summarize, in this chapter we have provided evidence that spontaneous state inferences are comparable to spontaneous trait inferences when both are contrasted to conditions in which there is no logical basis for inferences. If anything, the evidence for the inferential nature of state inferences was stronger, as we observed no significant differences between trait inferences and transferences. While the results need further exploration, as we have outlined above, we can safely consider spontaneous state inferences an alternative social inference to the established spontaneous trait inferences, as they are not mere associations activated by the perception of behavior, but inferences linked to the actors.

5 Efficiency of Spontaneous Social Inferences

5.1 On the Automaticity of Social Inferences

Winter and colleagues suggested that trait inferences seemed to be "largely, but not entirely, automatic" (Winter et al., 1985, p. 904), as the evidence the authors provided indicated that they occurred unintentionally and without perceivers' awareness. Most research on spontaneous trait inferences has focused on the aspect of spontaneity, that is, the unintentionality of trait inferences. And with good reason: When Winter and Uleman published their first studies on spontaneous trait inferences, their research surprised many scholars, as in the spirit of the cognitive miser's pursuit of resource preservation, traits were deemed one of the least likely concepts to be inferred spontaneously. The concept of automaticity however, while itself subject to extensive debate in the scientific community, encompasses more than unintentionality and awareness. Bargh (1994) figuratively named the four *horsemen of automaticity*: Awareness, efficiency, intention, and control. Our understanding of automaticity has since evolved (for a review, see Moors & De Houwer, 2006), and robust evidence has only been provided for the unintentionality and efficiency of spontaneous trait inferences. I will discuss the evidence available for the different aspects of automatic processes involved in STIs below, but note that the evidence for (un)consciousness and (un)controllability stems from few studies only, and should thus not be over-interpreted.

While often not explicitly defined, the therm *spontaneous* in STI is used mainly in reference to *unintentional*. I understand unintentional as goal-independent (Moors & De Houwer, 2006), that is, trait inferences occur unintentionally when behavior is processed in the absence of an explicit goal to form an impression. Usually, in STI research, unintentionality is investigated by manipulating the instructions given to participants when they are presented with trait-implying statements. For example, Carlston and Skowronski (1994) presented participants with trait-implying behavioral statements in a savings-in-relearning paradigm and asked them to either generate a specific trait word that could describe the respective actors' personality (specific impression condition), form a general impression about the actor (general impression condition), or to merely familiarize themselves with stimulus materials that would be used later in the experiment. After a filler task, participants were asked to learn pairs of portraits and trait words. In the critical trials, the portraits had been presented earlier, and the trait word was the one implied by the behavioral description. In control trials, neither portrait nor trait had been presented before. After another filler task, participants were once again presented with the portraits and asked to recall the corresponding trait. Carlston and Skowronski (1994) observed that participants recalled more traits for critical trials, as compared to the control condition, and interpret this pattern of results as evidence that the critical trials effectively constituted a re-learning, because participants had inferred the traits paired with the portraits in these trials earlier, when the portrait was presented with a behavioral statement implying that trait. Most importantly, Carlston and Skowronski observed this pattern of results for all instructions: Participants inferred traits from behavioral statements, even when no conscious intention to form an impression was elicited by an impression formation instruction. Further research robustly provided evidence that trait inferences occur unintentionally, when participants were instructed to merely familiarize themselves, memorize, or ignore the contents of the materials (Bassili & Smith, 1986; Carlston et al., 1995; Costabile, 2016; Ferreira et al., 2012; Ham & Vonk, 2003; McCarthy & Skowronski, 2011a; Na & Kitayama, 2011; Todd et al., 2011; Uleman & Moskowitz, 1994; Whitney et al., 1992).

Further evidence regarding the unintentionality of spontaneous trait inferences is their observation in tasks such as the probe recognition paradigm and the false recognition paradigm. In these paradigms, task performance is actually hindered by the occurrence of STIs: If participants infer a trait, responding accurately or quickly becomes more difficult, so participants clearly should suppress any inherent intentions to form STIs.

Second, spontaneous trait inferences are assumed to occur *unconsciously*. I understand unconscious inferences such that they occur outside of perceivers' subjective awareness: Perceivers are potentially able to access the inferences they draw unconsciously, but would need to allocate attention to the inference process (Moors & De Houwer, 2006). In the experimental paradigms used in STI research, participants are unlikely to attend to the inference process (Ferreira et al., 2012), as task demands are usually pitted against impression formation (for instance, experiments are often presented as memory experiments, without explicit mention of impression formation). And indeed, most perceivers appear to be unaware that they formed spontaneous impressions: When queried, participants in some studies self-reported not having formed impressions at all (Lupfer et al., 1990;

Moskowitz, 1993; Winter & Uleman, 1984). However, Uleman and Moskowitz (1994) found that between 6 and 14% of participants in one of their studies reported some awareness of trait inferences, and that this awareness correlated with the STI measure. If participants become aware of their inferences, this awareness appears to facilitate recall of trait-implying behavior. However, it is complicated to investigate non-conscious processing of stimuli or stimulus features (Wentura et al., 2017), and the zero awareness criterion is often not perfectly met.

Third, STIs also appear to be difficult to control or even *uncontrollable*: They occur in a savings-in-relearning paradigm when participants are instructed to actively avoid forming inferences (Shimizu, 2017) and participants do not seem to be able to suppress the effects of STIs on further information processing, such as predicting future behavior (McCarthy & Skowronski, 2011a). While STIs appear to possess some characteristics of uncontrollable processes, they cannot be considered entirely outside of perceivers' control. More modern accounts of automatic and controlled processes state that phenomena are rarely process-pure (Conrey et al., 2005; Ferreira et al., 2012; Jacoby, 1991), but an interplay of both automatic and controlled processes with varying degrees. The process-dissociation procedure (PDP; Jacoby, 1991) allows for identifying contributions of automatic and controlled processes to a given phenomenon. McCarthy and Skowronski (2011b) conducted such analyses for both explicit and spontaneous trait inferences. They found that the contribution of automatic and controlled processes to the influence of those inferences on subsequent responses was similar for explicit and spontaneous trait inferences, suggesting that STIs are comprised of both, automatic and controlled processes.

Fourth, the evidence on the *efficiency* of STI is mixed: They appear to require at least some amount of cognitive resources. Perceivers must attend to and process the trait-implying behavior, which is more difficult under working memory load (Chun et al., 2002; Wigboldus et al., 2004). Uleman and Moskowitz (1994) asked participants to memorize trait-implying statements while ignoring their meaning, locate specific letters in the statements, or attend to the sound of the statements' words to decide whether they rhyme with a given target word. They found that these instructions resulted in reduced trait inferences, yet did not fully eliminate them. Winter et al. (1985) induced working memory load in participants concurrent with the presentation of trait-implying statements. They asked participants to memorize single vs. multiple digits, and presented trait-implying statements as ostensible distractors. They found that participants in both digit conditions inferred traits, albeit somewhat reduced in the multiple digit condition. This finding has been replicated repeatedly

(Lupfer et al., 1990; Todorov & Uleman, 2003; Uleman et al., 1992; Wells et al., 2011). Wells et al. (2011) found that STI effects were correlated with individual differences in working memory, suggesting that STIs at least partially rely on working memory capacity. Other manipulations of working memory load during encoding of trait-implying stimuli yielded similar results, such as a concurrent lie-detection task (Crawford, Skowronski, Stiff, & Scherer, 2007), and the presentation of trait-implying statements in rapid succession (Todorov & Uleman, 2003), both of which reduced spontaneous trait inference effects.

Even though not all features of automatic processing have been investigated systematically and conclusively, STIs are still presumed to possess a high degree of automaticity. The presumed automaticity of STIs, in turn, has been cited as a possible explanation for correspondence biases (Ferreira et al., 2012; Moskowitz, 2005). While there is ample evidence suggesting that spontaneous trait inferences occur with some degree of efficiency, no such evidence has been provided for state inferences (or other social inferences, for that matter). The paradigms employed in Experiments 1-7 do not explicitly instruct participants to form impressions about the target actors, indicating that the observed inferences did not require conscious intention. However, it is still conceivable that spontaneous trait and state inferences differ qualitatively with regards to their efficiency. State inferences may require more (or less) cognitive resources than trait inferences, which could indicate a hierarchy or sequential processing of spontaneous social inferences. Such processing would be compatible with the assumption that correspondence biases occur because traits represent the dominant category in impression formation, and are drawn routinely and effortlessly. To our knowledge, so far no empirical evidence for this assertion has been provided.

In the current chapter, we present research investigating if and to what extent spontaneous state inferences possess similar characteristics of automaticity as STIs. As investigating all aspects of automaticity is beyond the scope of this dissertation, the experiments reported in this chapter focus on the relative efficiency of spontaneous trait and state inferences. In Experiment 8, we manipulated working memory load during encoding (Todorov & Uleman, 2003, Exp. 3) in a false recognition paradigm using single-implication behavioral statements. In Experiment 9, working memory load was manipulated during encoding of dual-implication behavioral statements. We observed the occurrence of spontaneous state and trait inferences under load for single-implication materials, but no inferences under load for dual-implication materials.

5.2 Experiment 8: Efficient Social Inferences in Single-Implication Behaviors

To investigate the degree of efficiency of spontaneous state inferences, we manipulated working memory load during encoding of single-implication behavioral statements. In the high load condition, participants were asked to memorize a sequence consisting of six random digits during encoding of each behavioral statement and presented with a digit recognition task after each statement. In the low load condition, the sequences contained of six identical digits. In the low load condition, trait and state inference effects should be observed, with significantly higher false recognition rates for implied state and trait adjectives as compared to implied-other adjectives. In the high load condition, we expected the same pattern of results for traits, possibly with reduced effect size (Wells et al., 2011). For state inferences in the high load condition, two outcomes were plausible: If state inferences are comparable to trait inferences in their relative resource-independence, we expected to observe a similar pattern of results as for trait inferences. If, however, state inferences are more resource-dependent than trait inferences, the difference between false recognition rates for implied and implied-other state adjectives should be diminished or even disappear in the high load condition.

5.2.1 Method

5.2.1.1 Sample Size Determination.

Our primary interest in this experiment are the 2 x 3 interaction effects between Condition (implied vs. implied-other) and Working Memory Load (no vs. low vs. high) in a mixed ANOVA for both conditions of Inference (trait vs. state). We aimed to provide sufficient statistical power ($1 - \beta = .80$) to detect the minimal effect size of interest of $\eta_p^2 = .022$ for the 2 x 2 x 3 interaction between Condition (implied vs. implied-other) x Inference (trait vs. state) x Working Memory Load (low vs. high). Thus, valid data of $N = 352$ participants were required. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

5.2.1.2 Participants

The current experiment relied on valid data from a total of $N = 358$ participants (183 female; average age $M = 31.1$ years, $SD = 10.3$, ranging from 18 to 69). Participants were recruited via Prolific (www.prolific.ac) and received monetary compensation of 2.05 GBP (approx. 2.84 USD). Following our pre-registered exclusion criteria, we excluded data from ten participants because they did not pass an attention check, and data from 41 participants because they self-reported insufficient language proficiency.

5.2.1.3 Materials

We selected 24 target stimuli of those already employed in Experiments 1 and 2, consisting of a behavioral statement in German language, a portrait (selected from the 10k US adult faces database; Bainbridge, Isola, & Oliva, 2013), and a given name. Twelve behavioral statements unambiguously implied a trait (e.g., "I gave 5 Euros to the homeless person" - generous), and 12 other unambiguously implied a state (e.g., "I yelled at the guy in the bakery" – angry), but these were not mentioned in the statement. Additionally, we used 12 filler stimuli equal to the target stimuli, with the exception that they explicitly mention a trait or state word (see Appendix, Tables A.6 - A.9, for the complete list of statements and pre-test results).

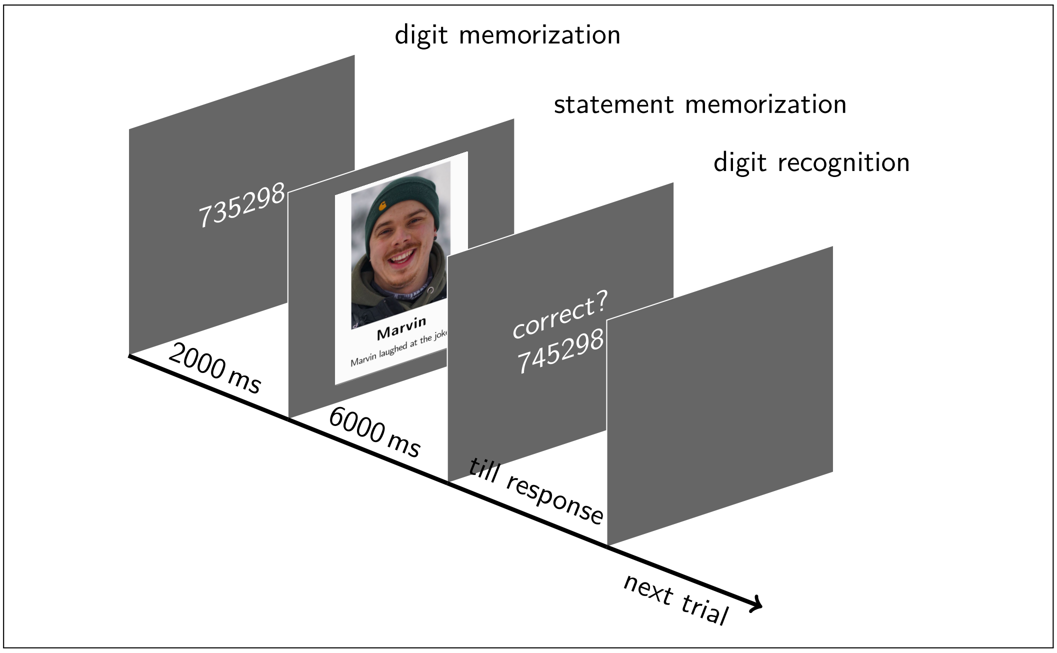

Fig. 5.1.: Example for a learning phase trial in Experiments 8 and 9, high load condition.

5.2.1.4 Procedure

Experiment 8 followed the same procedure as Experiment 3, with the exception that we added a manipulation of working memory load (adapted from Todorov & Uleman, 2003, Experiment 3): During the initial learning phase, participants were presented with 36 trials. A trial consisted of six digits which were presented for two seconds. In the low load condition, the sequence consisted of identical digits (e.g., 111111). In the high load condition, the sequence consisted of random digits (e.g., 735298). Participants were instructed to memorize the digits and to keep them in mind while additionally memorizing the subsequently presented person information, which consisted of a portrait picture, a name and a behavioral statement. The person information was presented for six seconds and followed immediately by a digit recognition screen. Participants were presented with six digits. In the low load condition, the presented digits were either the identical or an entirely altered row of digits (e.g., 111111 vs. 222222). In the high load condition, the presented digits were either the identical row of digits or the identical row in which one randomly selected digit was altered (e.g., 735298 vs. 745298). Participants were asked to decide whether the digits presented were the same they had seen earlier, with the response options yes or no. The next trial began with an interval of two seconds after the response was recorded. In the no load condition, no digits were presented or tested, and there was no interval between trials. Immediately after the learning phase, participants completed one practice recognition trial plus 24 experimental recognition trials as well as 12 filler trials, in which the portraits and names from the learning phase were presented together with an adjective probe. The entire procedure lasted approximately 12 minutes. In Experiment 8, participants were not presented with the implicit personality theories measure (Implicit Personality Theories Questionnaire [8 items], translated into German; Dweck et al., 1995).

5.2.2 Results

We conducted all data analyses using R 3.6.3 (R Core Team, 2020). As a manipulation check, we calculated individual detection parameter A' (ranging from 0 to 1; Pollack, 1970) as a measure of individual working memory performance. A' scores suggested almost perfect performance in the low load condition (M = .975, SD = .026) that significantly differed from .5, $t(116) = 194.319$, $p < .001$ (one-tailed), $d_s = 17.965$, 95% CI [15.61, 20.32]. In the high load condition, participants' performance also differed significantly from .5 (M = .885, SD = .074), $t(120) = 56.858$, $p < .001$ (one-tailed), $d_s = 5.169$, 95% CI [4.419, 5.919]. A' scores in the low load condition were larger compared to the high load condition, $t(150.659) = 12.492$, $p < .001$ (one-tailed), $d_s = 1.599$, 95% CI [1.305, 1.892], indicating that the working memory load manipulation had the intended effect.

As customary with the false recognition paradigm, we focused analyses on response rates (Todorov & Uleman, 2003). Individual false recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) x 3 (Load: no vs. low vs. high) ANOVA with the factor Load varying between participants and Inference and Condition varying within participants. Results showed significant main effects for Inference, $F(1,355) = 128.342$, $p < .001$, $\eta^2_G = 0.092$, 90% CI [.049, .144], Condition, $F(1,355) = 192.319$, $p < .001$, $\eta^2_G = 0.155$, 90% CI [.101, .216], and Load $F(2,355) = 8.137$, $p < .001$, $\eta^2_G = 0.022$, 90% CI [.002, .051], with significant interactions between Condition and Load, $F(2,355) = 3.428$, $p = .034$, $\eta^2_G = 0.007$, 90% CI [.000, .024], between Inference and Condition, $F(1,355) = 14.85$, $p < .001$, $\eta^2_G = 0.011$, 90% CI [.000, .037], and a significant triple-interaction, $F(2,355) = 4.685$, $p = .01$, $\eta^2_G = 0.007$, 90% CI [.000, .025], but no significant interaction between Load and Inference, $F(2,355) = 2.02$, $p = .134$, $\eta^2_G = 0.003$, 90% CI [.000, .016].

5.2.2.1 Trait Inferences

Individual false recognition rates for both inference conditions were submitted to separate 2 (Condition: implied vs. implied-other) x 3 (Load: no vs. low vs. high) ANOVAs. In the trait condition, we observed significant main effects for Load, $F(2,355) = 9.882$, $p < .001$, $\eta^2_G = 0.027$, 90% CI [.004, .058], for Condition, $F(1,355) = 65.834$, $p < .001$, $\eta^2_G = 0.066$, 90% CI [.030, .114], and importantly, a significant interaction between Condition and Load, $F(2,355) = 3.443$, $p = .033$, $\eta^2_G = 0.007$, 90% CI [.000, .025].

Separate analyses confirmed the expected effects in all load conditions: In the no load condition, we observed significant differences in false recognition rates between the implied trait condition ($M = .31$, $SD = .23$) and the implied-other trait condition ($M = .22$, $SD = .22$), $t(119) = 3.227$, $p = .001$ (one-tailed), $d_z = 0.385$, 95% CI [0.141, 0.629], In the low load condition, false recognition rates in the implied trait condition ($M = .38$, $SD = .26$) were significantly larger than in the implied-other trait condition ($M = .20$, $SD = .19$), $t(116) = 7.171$, $p < .001$ (one-tailed), $d_z = 0.762$, 95% CI [0.524, 1]. In the high load condition, we observed significant differences between the implied trait condition ($M = .41$,

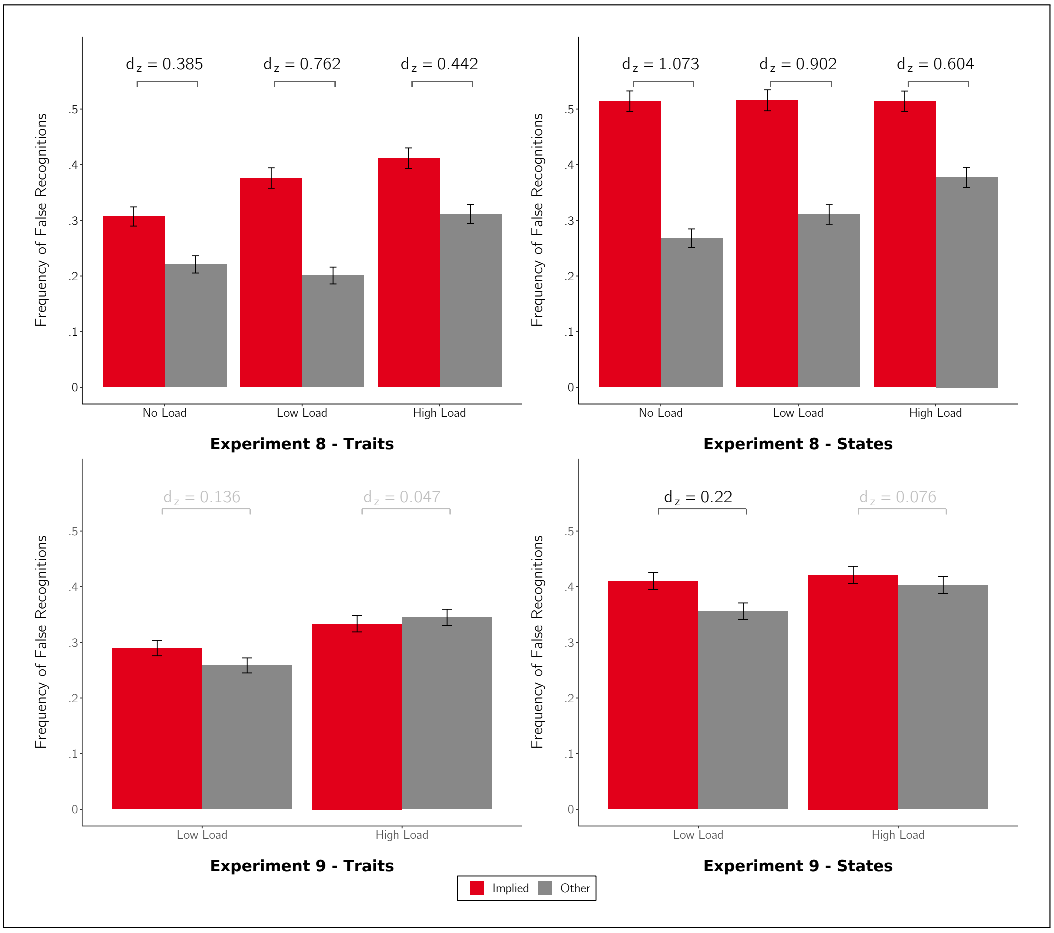

Fig. 5.2.: Mean Proportion of False Recognition Responses (Experiments 8 and 9) as a Function of Type of Inference and Experimental Condition. Error bars represent *SE*. Effect sizes printed in grey font stem from non-significant comparisons.

SD = .22) and the implied-other trait condition (M = .31, SD = .24), $t(120) = 3.914$, $p < .001$ (one-tailed), $d_z = 0.442$, 95% CI [0.209, 0.676].

For trait inferences, the inference effects were largest for the low load condition, as compared to the high load condition $t(235.741) = 2.106$, $p = .036$, $d_s = 0.273$, 95% CI [0.016, 0.529], and the no load condition $t(233.685) = -2.463$, $p = .015$, $d_s = -0.32$, 95% CI [-0.577, -0.062]. We observed no significant difference between no load and high load conditions, $t(238.578) = -0.39$, $p = .697$, $d_s = -0.05$, 95% CI [-0.304, 0.204].

5.2.2.2 State Inferences

For state inferences, we observed a significant main effect for Condition, $F(1{,}355) = 165.224$, $p < .001$, $\eta^2_G = 0.157$, 90% CI [.102, .218], and a significant interaction between Condition and Load, $F(2{,}355) = 4.454$, $p = .012$, $\eta^2_G = 0.01$, 90% CI [.000, .030]. However, the main effect for Load did not reach significance, $F(2,355) = 2.933$, $p = .055$, $\eta^2_G = 0.008$, 90% CI [.000, .027].

Separate analyses confirmed the expected effects in all load conditions: In the no load condition, we observed significant differences in false recognition rates between the implied state condition ($M = .51$, $SD = .23$) and the implied-other state condition ($M = .27$, $SD = .23$), $t(119) = 9.793$, $p < .001$ (one-tailed), $d_z = 1.073$, 95% CI [0.802, 1.344]. In the low load condition, false recognition rates in the implied state condition were significantly larger ($M = .52$, $SD = .22$), as compared to the implied-other state condition ($M = .31$, $SD = .23$), $t(116) = 7.745$, $p < .001$ (one-tailed), $d_z = 0.902$, 95% CI [0.63, 1.174]. In the high load condition, we observed a significant difference between the implied state condition ($M = .51$, $SD = .23$) and the implied-other state condition ($M = .38$, $SD = .22$), $t(120) = 4.98$, $p < .001$ (one-tailed), $d_z = 0.604$, 95% CI [0.344, 0.863].

For state inferences, the inference effects (calculated as difference scores of implied – implied-other) differed significantly between no load and high load conditions, $t(237.379) = 2.947$, $p = .004$, $d_s = 0.379$, 95% CI [0.123, 0.636]. We observed no significant differences between no load and low load conditions $t(233.968) = 1.115$, $p = .266$, $d_s = 0.145$, 95% CI [-0.111, 0.401], nor between low load and high load conditions $t(235.937) = 1.805$, $p = .072$, $d_s = 0.234$, 95% CI [-0.022, 0.49].

The state inference effect was larger as compared to the trait inference effect, $t(713.358) = 3.524$, $p < .001$ (one-tailed), $d_s = 0.263$, 95% CI [0.116, 0.411], albeit this effect was only driven by the no load condition, $t(237.119) = 4.36$, $p < .001$ (one-tailed), $d_s = 0.563$, 95% CI [0.304, 0.822]. Differences between the trait inference effects and state inference effects were non-significant in the low load condition, $t(230.506) = 0.83$, $p = .204$ (one-tailed), $d_s = 0.109$, 95% CI [-0.149, 0.366], and in the high load condition, $t(239.026) = 0.954$, $p = .171$ (one-tailed), $d_s = 0.123$, 95% CI [-0.131, 0.376].

5.2.2.3 Auxiliary Analyses

The interplay of automatic and controlled processes in spontaneous social inferences can be quantified by use of the process dissociation procedure (PDP; Ferreira et al., 2012; Jacoby, 1991; McCarthy & Skowronski, 2011a; Payne, 2005). This simple modeling approach allows approximating the contributions of automatic and controlled processes to an observed outcome, given an appropriate research design. An experiment needs to consist of two conditions: An inclusion condition, in which both automatic and controlled processes work together and an exclusion condition where the two types of processes work in opposition. The PDP model assumes that participants' performance on the inclusion task is due to either controlled processes or, in case controlled processes fail, to automatic processes. In the exclusion task, participants' performance is due to automatic processes and the failure of controlled processes. Combining these assumptions, the contribution of controlled processes to task performance can be calculated as the difference between participants' performance in the inclusion task and participants' performance in the exclusion task. An exemplary path model for the false recognition paradigm is depicted in Figure 5.3.

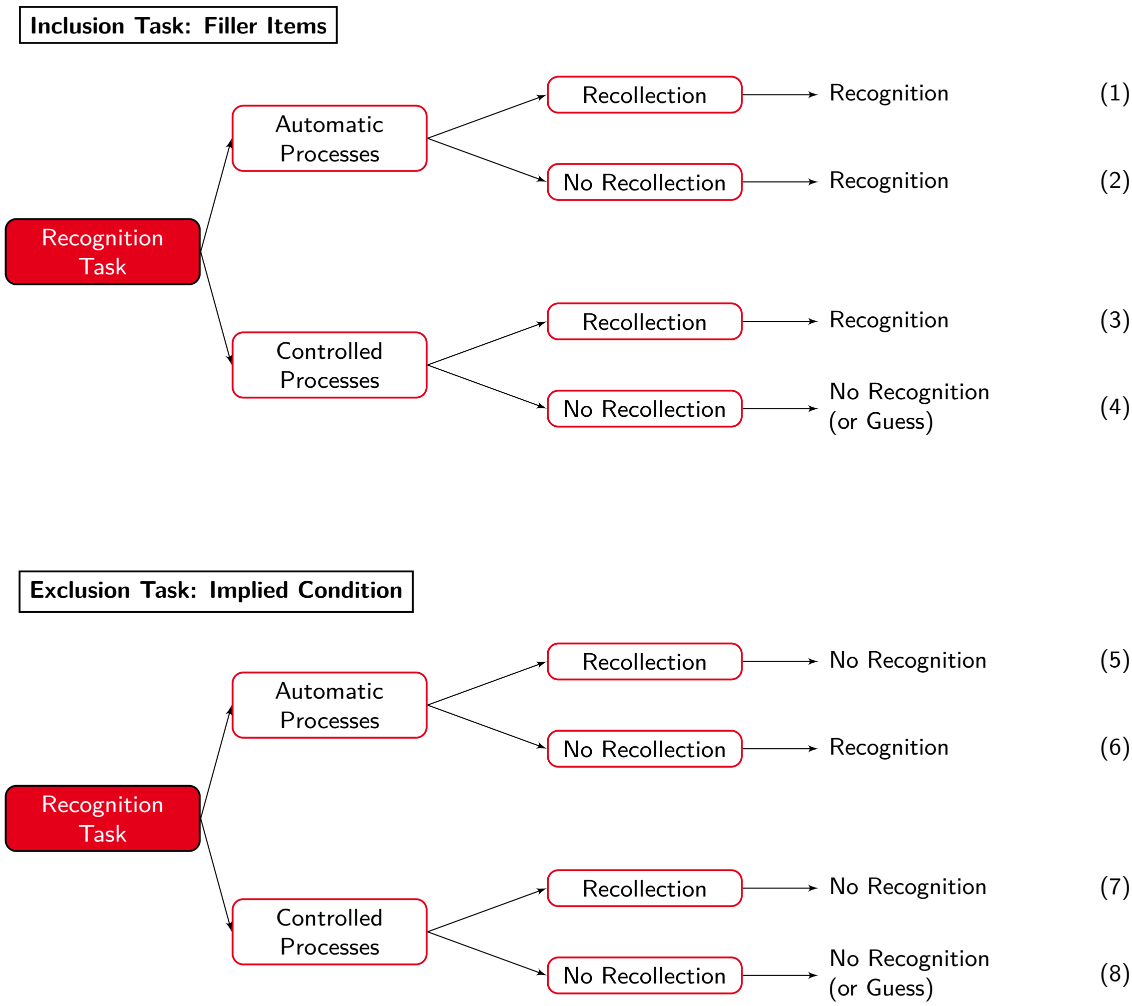

Fig. 5.3.: Path Model for the Process Dissociation Procedure. Adapted from Vaterrodt-Plünnecke et al. (2002).

In the false recognition paradigm, filler trials, in which participants are probed with an adjective that was actually mentioned in the statement presented earlier, can be construed as an inclusion task: both controlled processes and automatic processes favor the same response, with the use of either leading to a correct response, namely recognition of the probe as part of the respective actor's description. Participants may recall the exact behavior description and know that the adjective was part of the description, to then use this information to produce a correct response (Fig. 5.3, Path 3). At the same time, participants may draw a social inference upon reading the behavioral statement and encode the inference as a property of the actor, regardless of the fact that the social inference is also explicitly named in the statement - which is assumed to represent automatic processes. When participants are then asked to recognize the adjective that had been presented in the statement earlier, they may not consciously recall that it had in fact been present in the statement. Instead, the social inference they had drawn may be available, leading participants to produce a correct response (Path 2). If participants have drawn an inference and also recall the statement, they are also likely to produce a correct response (Path 1). In summary, for filler trials, controlled and automatic processes work in concert, and both processes will produce a recognition response.

For target trials in the implied condition, in which participants are probed with an adjective that was implied by, but not actually mentioned in the statement presented earlier, automatic and controlled processes work in opposition to each other. Upon reading the behavioral statement, participants may draw a social inference and encode that social inference as a property of the actor. The effect of this automatic process would hinder the correct rejection of a target adjective in the recognition task, and facilitate false recognition (Path 6). The controlled process, however, which would lead participants to consciously recollect that the target adjective had not been part of the behavioral statement, would favor a correct rejection (Path 7). In other words, false recognitions occur when automatic processes favor a false recognition and controlled processes fail to produce the correct rejection. Automatic and controlled processes thus work in opposition to each other in target trials in the implied condition.

Using PDP, we can compute the relative contributions of automatic and controlled processes via participants' performance in the exclusion and inclusion tasks. The PDP approach allows for the computation of an automatic processing parameter estimate (A parameter) and a controlled processing parameter estimate (C parameter). However, there may be a third process interfering with participants' performance, namely response biases and/or guessing. In the absence of the effect of a previously drawn social inference, and without conscious recollection of the behavioral statement, participants may simply guess yes or no in some trials (Paths 4 & 8). Such guessing tendencies were assessed in our experiments by examining responses on implied-other trials, in which the probe adjective was neither implied nor did it appear in the behavioral statement. In our analyses, we corrected for guessing using a probabilistic correction approach, using the implied-other condition as an estimate for guessing (Buchner et al., 1995; McCarthy & Skowronski, 2011a). As the PDP approach is a relatively simple model, calculated parameter estimates can have negative values for some participants. However, it is theoretically impossible for either automatic processing or controlled processing to have less-than-zero influence on responding. Because of this theoretical impossibility, we re-coded values of participants whose calculated estimates of controlled processing were negative (i.e., a false recognition rate greater than hit rate) to zero. Likewise, values of participants with negative estimates of automatic processing were re-coded to zero (McCarthy & Skowronski, 2011a).

We hypothesized that working memory load during the encoding of behavioral stimuli would affect the social inference process to some degree, which in turn should reduce the automatic processing parameter estimate (A parameter), but leave the controlled processing parameter estimate (C parameter) largely unaffected. Note, however, that the PDP model does not allow specifying process characteristics. The parameters A and C are general terms and we cannot link them directly to any process characteristic (i.e., unintentionality, unconsciousness, and so forth). In order to specify process characteristics we would need to validate the model by systematically manipulating different aspects experimentally and observing changes in the processing parameter estimates.

Individual automatic processing parameter estimates were submitted to a 2 (Inference: trait vs. state) x 3 (Load: no vs. low vs. high) ANOVA with the factor Load varying between participants and Inference varying within participants. Results showed a significant main effect for Inference, $F(1{,}355) = 25.67$, $p < .001$, $\eta^2_G = 0.033$, 90% CI [.009, .069]. The main effect for Load did not reach significance, $F(2{,}355) = 1.506$, $p = .223$, $\eta^2_G = 0.004$, 90% CI [.000, .018]. We observed a significant interaction between Inference and Load, $F(2{,}355) = 4.474$, $p = .012$, $\eta^2_G = 0.012$, 90% CI [.000, .034]. For states, we observed significant automatic processing parameter estimates across all working memory load conditions. We observed a significant difference between estimates in the no load condition ($M = 0.383$, $SD = 0.277$) and the high load condition ($M = 0.289$, $SD = 0.302$), $t(237.564) = 2.538$, $p = .012$, $d_s = 0.327$, 95% CI [0.071, 0.582].

Automatic processing parameter estimates for traits also differed significantly from zero across all working memory load conditions. We observed a significant difference between estimates in the no load condition ($M = 0.205$, $SD = 0.225$) and the low load condition ($M =$

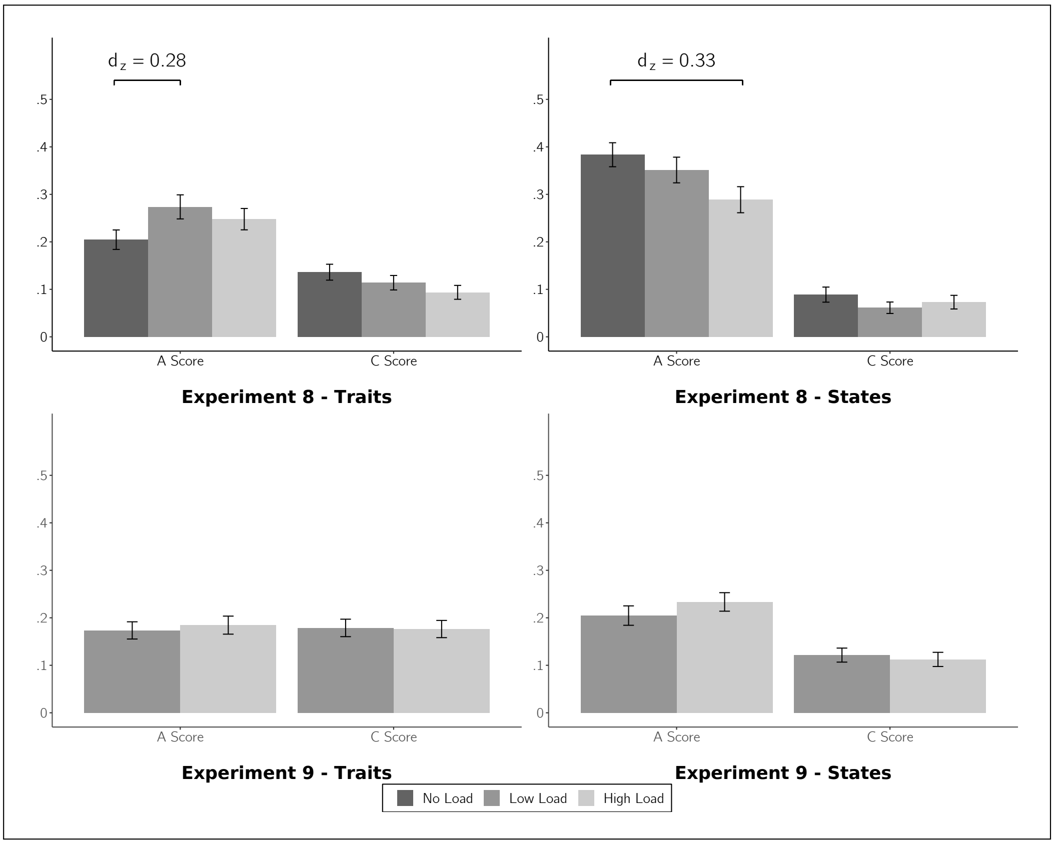

Fig. 5.4.: Mean automatic and controlled processing parameter estimates (Experiments 8 and 9) as a function of type of inference and experimental condition. Error bars represent *SE*. Only significant effect sizes are stated.

0.274, $SD = 0.274$), $t(237.564) = 2.538$, $p = .012$ (one-tailed), $d_s = 0.327$, 95% CI [0.071, 0.582]. Automatic processing parameter estimates were larger for states, as compared to traits, $t(697.396) = 4.875$, $p < .001$, $d_s = 0.364$, 95% CI [0.216, 0.512].

Individual controlled processing parameter estimates (see upper two panels of Fig. 5.4) were submitted to a 2 (Inference: trait vs. state) x 3 (Load: no vs. low vs. high) ANOVA with the factor Load varying between participants and Inference varying within participants. Results showed a significant main effect for Inference, $F(1,355) = 25.67$, $p < .001$, $\eta^2_G = 0.033$, 90% CI [.009, .069]. The main effect for Load did not reach significance, $F(2,355) = 1.506$, $p = .223$, $\eta^2_G = 0.004$, 90% CI [.000, .018]. We observed a significant interaction between Inference and Load, $F(2,355) = 4.474$, $p = .012$, $\eta^2_G = 0.012$, 90% CI [.000, .034]. Controlled processing parameter estimates did not differ significantly across load conditions, but were larger for traits as compared to states, $t(708.084) = 3.285$, $p = .001$, $d_s = 0.246$, 95% CI [0.098, 0.393].

5.2.3 Discussion

Results of Experiment 8 replicate the pattern of results of Experiments 1 and 2: Participants spontaneously inferred both traits and states from single-implication behaviors in all three working memory load conditions, indicating that both spontaneous trait and state inferences occur even under working memory load. Working memory load did, however, qualify the effects. In the trait condition, the inference effect (calculated as difference scores of implied – implied-other) was smaller in the high load condition, as compared to the low load condition. This pattern of results replicates results from previous experiments, indicating that STI may not be fully independent of working memory load (Lupfer et al., 1990; Todorov & Uleman, 2003; Uleman et al., 1992; Wells et al., 2011). We observed the same pattern of results in the state condition: the inference effect was smaller in the high load condition, as compared to the no load condition. These results indicate that the spontaneous state inference process is to some degree impacted by working memory load - but the effect is not eliminated. As in Experiments 1 and 2, effect sizes were larger for state inferences as compared to trait inferences.

We observed inference effects in all three working memory load conditions with effect sizes comparable to our previous experiments for both, states and traits. The trait inference effect was larger in the low load condition, as compared to the high load condition, replicating results from previous studies indicating that STI are dependent on working memory load. More importantly, we observed the same pattern for state inferences, indicating that they depend on working memory load to a similar degree as trait inferences. Exploratory PDP analyses provided converging evidence.

While the pattern of results for most conditions support the hypothesis that working memory load reduces spontaneous social inferences, the trait inference effect was larger in the low load condition, as compared to the no load condition. This increase in effect size for the low load condition does not fit the expected pattern of results. There is no apparent explanation for this deviation. One may speculate that some load could be beneficial in that it increases attention to the stimuli, but as this effect did not occur for the within-participants factor states, this explanation seems unlikely. This result might be a false positive, and warrants replication of the experiments to investigate whether it actually represents a robust finding. This deviation, however, does not contradict our overall findings.

Spontaneous social inferences from behavior appear to be relatively efficient if the behavior offers a clear, single implication. Such behavior, although frequently used in research, is rare in everyday life, as we pointed out earlier (see Chapter 2.3). In Experiment 9, we investigated whether and to what extent social inferences occur in a similarly efficient manner for dual-implication behaviors.

5.3 Experiment 9: Inefficient Social Inferences in Dual-Implication Behaviors

In Experiment 9, we replicated Experiment 8 with dual-implication materials to investigate whether spontaneous social inferences would be drawn efficiently from dual-implication materials, that is, behavioral statements that imply both a trait and a state.

5.3.1 Method

5.3.1.1 Sample Size Determination

Our primary interest in this experiment were the 2 x 2 interaction effects between Load (low vs. high) and Condition (implied vs. implied-other) in a mixed ANOVA for both conditions of Inference (trait vs. state). We aimed to provide sufficient statistical power ($1 - \beta = .80$) to detect the minimal effect size of interest of $\eta_p^2 = .022$ for the 2 x 2 x 2 interaction between Load (low vs. high) x Condition (implied vs. implied-other) x Inference (trait vs. state). Thus, valid data of $N = 352$ participants were required. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

5.3.1.2 Participants

The current experiment relies on valid data from a total of $N = 350$ participants (134 female; average age $M = 35.1$ years, $SD = 13.6$, ranging from 18 to 76). Participants were recruited via Prolific (www.prolific.ac) and received monetary compensation of 1.56 GBP (approx. 1.98 USD). Following our pre-registered exclusion criteria, we excluded data from 13 participants because they did not pass an initial attention test. Due to a programming error, we were unable to assess participants' accuracy in the digit recognition task and were thus unable to exclude participants with area operating characteristics scores (Pollack,

1970) of A' < .50 in this task, as we had initially pre-registered. However, this analysis was merely intended to ensure the efficiency of our working memory load manipulation. As previous research has already demonstrated the efficiency of the exact manipulation used in this experiment (Todorov & Uleman, 2003), we deem it appropriate to assume that the manipulation did have the intended effect.

5.3.1.3 Materials

We selected 24 target stimuli consisting of a behavioral statement, a portrait (selected from the 10k US adult faces database; Bainbridge, Isola, & Oliva, 2013), and a given name, see Figure 3.1 for an example. All behavioral statements were ambiguous in that the behavior could be attributed to a trait or state of the actor (e.g., "Maggie sighed as her unread emails piled up."; unorganized, overwhelmed), but these were not explicitly mentioned in the statement. Additionally, we used 12 filler stimuli equal to the target stimuli, with the exception that they explicitly mentioned a trait or state word (see Table A.15 in the Appendix for the complete list of statements and pre-test results).

5.3.1.4 Procedure

Experiment 9 followed the same procedure as Experiment 8, with the following exceptions: We omitted the no load condition, in which participants were not presented with any digit task. The entire procedure lasted approximately 15 minutes.

5.3.2 Results

We conducted all data analyses using R 3.6.3 (R Core Team, 2020). As customary with the false recognition paradigm, we focused analyses on response rates (Todorov & Uleman, 2003).

Individual false recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) x 2 (Load: low vs. high) ANOVA with the factor Load varying between participants and Inference and Condition varying within participants. Results showed significant main effects for Inference, $F(1{,}348) = 81.851$, $p < .001$, $\eta^2_G = 0.052$, 90% CI [.021, .096], Condition, $F(1{,}348) = 4.318$, $p = .038$, $\eta^2_G = 0.003$, 90% CI [.000, .021], and Load $F(1{,}348) = 6.711$, $p = .01$, $\eta^2_G = 0.01$, 90% CI [.000, .034], with no significant interactions.

Individual false recognition rates for both inference conditions were submitted to separate 2 (Condition: implied vs. implied-other) x 2 (Load: low vs. high) ANOVAs. In the trait condition, we observed a significant main effect for Load, $F(1,348) = 9.532$, $p = .002$, $\eta^2_G = 0.014$, 90% CI [.001, .041]. However, neither the main effect for Condition, $F(1,348) = 0.504$, $p = .478$, $\eta^2_G = 0$, 90% CI [.000, .011], nor the interaction were significant, $F(1,348) = 2.302$, $p = .13$, $\eta^2_G = 0.002$, 90% CI [.000, .018].

In the state condition, we observed a significant main effect for Condition, $F(1,348) = 5.205$, $p = .023$, $\eta^2_G = 0.006$, 90% CI [.000, .026]. However, neither the main effect for Load, $F(1,348) = 2.032$, $p = .155$, $\eta^2_G = 0.003$, 90% CI [.000, .020], nor the interaction were significant, $F(1,348) = 1.272$, $p = .26$, $\eta^2_G = 0.001$, 90% CI [.000, .015].

Separate analyses confirmed expected effects only in the low load condition: For traits, the difference between false recognition rates in the implied trait condition ($M = .29$, $SD = .23$) and the implied-other trait condition ($M = .26$, $SD = .23$) missed significance, $t(175) = 1.635$, $p = .052$, $d_z = 0.171$, 95% CI [0.005, 0.337]. For states, participants showed higher false recognition rates in the implied state condition ($M = .41$, $SD = .25$) than in the implied-other state condition ($M = .36$, $SD = .24$), $t(175) = 2.331$, $p = .01$, $d_z = 0.262$, 95% CI [0.072, 0.453].

In the high load condition, we observed no significant differences between conditions, indicating that under high load, neither traits nor states had been inferred spontaneously. False recognition rates did not differ significantly between the implied trait condition ($M = .33$, $SD = .24$) and the implied-other trait condition ($M = .34$, $SD = .25$), $t(173) = -0.555$, $p = .71$, $d_z = -0.022$, 95% CI [-0.196, 0.153]. We observed no significant difference between the implied state condition ($M = .42$, $SD = .23$) and the implied-other state condition ($M = .40$, $SD = .25$), $t(173) = 0.84$, $p = .201$, $d_z = 0.108$, 95% CI [-0.077, 0.293].

Note that the inference effects (calculated as difference scores of implied – implied-other) did not differ significantly between load conditions, neither for trait inferences $t(345.414) = 1.517$, $p = .065$, $d_s = 0.17$, 95% CI [-0.044, 0.384], nor for state inferences $t(346.733) = 1.128$, $p = .13$, $d_s = 0.13$, 95% CI [-0.084, 0.345].

5.3.2.1 Auxiliary Analyses

We followed the same PDP approach described in Experiment 8. Individual controlled processing parameter estimates (see lower two panels of Fig. 5.4) were submitted to a 2 (Inference: trait vs. state) x 2 (Load: low vs. high) ANOVA with the factor Load varying between participants and Inference varying within participants. Results showed a significant main effect for Inference, $F(1{,}336) = 17.945$, $p < .001$, $\eta^2_G = 0.019$, 90% CI [.003, .051]. The main effect for Load did not reach significance, $F(1{,}336) = 0.094$, $p = .759$, $\eta^2_G = 0$, 90% CI [.000, .007]. We observed no significant interaction between Inference and Load, $F(1{,}336) = 0.055$, $p = .815$, $\eta^2_G = 0$, 90% CI [.000, .005].

Individual automatic processing parameter estimates were submitted to a 2 (Inference: trait vs. state) x 2 (Load: low vs. high) ANOVA with the factor Load varying between participants and Inference varying within participants. Results showed a significant main effect for Inference, $F(1{,}336) = 5.149$, $p = .024$, $\eta^2_G = 0.006$, 90% CI [.000, .028]. The main effect for Load was non-significant, $F(1{,}336) = 0.922$, $p = .338$, $\eta^2_G = 0.001$, 90% CI [.000, .016]. We observed no significant interaction between Inference and Load, $F(1{,}336) = 0.248$, $p = .619$, $\eta^2_G = 0$, 90% CI [.000, .010].

5.3.3 Discussion

Results of Experiment 9 replicate the pattern of results of Experiments 3 and 4: Participants spontaneously inferred both traits and states from dual-implication stimulus materials in the low working memory load condition, albeit the trait inference effect did not reach statistical significance and effect sizes were smaller as compared to single-implication stimulus materials. In the high working memory load condition, we observed no evidence of spontaneous social inferences, failing to replicate previous results on spontaneous trait inferences under working memory load (e.g., Todorov & Uleman, 2003; Wells et al., 2011).

Notable differences between our experiment and previous research is the use of dual-implication materials, that allow for both a trait and a state inference. While dual-implication materials have been employed to demonstrate the occurrence of social inferences, to our knowledge, the efficiency of inferences from such materials has not yet been investigated. Dual-implication material may itself require more cognitive processing capacity, as at least two inferences may be drawn simultaneously, as Experiment 4 indicated. Participants might thus suffer more from reduced working memory capacity when presented with dual-implication behaviors, which may explain why the effect sizes in the low load condition were smaller, as compared to effect sizes in other experiments, in which no load was induced at encoding (Experiments 3, 4, 6, & 7). The cognitive demands of memorizing a simple sequence of equal digits combined with processing dual-implication statements may already reduce social inference effects. Reduced automatic processing parameter estimates from PDP analyses, as compared to Experiment 8, can be interpreted as converging evidence for reduced social inferences in both load conditions – but comparisons across experiments need to be done cautiously, and systematic comparison in a single experiment is needed. Increased cognitive demand may also explain why we did not observe spontaneous social inference effects in the high load condition. Prior research found reduced spontaneous trait inference effects under high load conditions, as compared to low load or no load conditions, using single-implication materials with the same load manipulation in the same paradigm (Todorov & Uleman, 2003).

Future research needs to further investigate the efficiency of spontaneous social inferences using dual-implication materials. While in Experiment 9, we used similar materials as in Experiments 3, 4, 6, and 7, in which no additional working memory load was induced, we did not include a no-load condition, which would have provided insight into whether the particular method and materials we employed in this experiment would generally evoke spontaneous social inferences in the absence of working memory load. Furthermore, we were unable to assess the impact of our manipulation on participants' subjective working memory load due to the programming error described above. Future research would benefit from an objective assessment of the degree of working memory load that participants experience, for instance the individual area operating characteristic A' (Pollack, 1970), which we used in Experiment 8.

5.4 General Discussion

In Experiments 8 and 9, we observed evidence indicating that participants spontaneously inferred both traits and states from behavioral descriptions. Working memory load reduced effect sizes for single-implication behaviors, and strongly reduced or even eliminated the effect for dual-implication behaviors. Complex materials thus did seem to moderate the inference process: We did find spontaneous social inferences under working memory load *only* when we used single-implication materials (Experiment 9). Using single-implication materials, we observed both trait and state inferences in all three load conditions. Smaller effect sizes in the high load conditions indicate that working memory load may reduce spontaneous social inferences. More importantly, however, the use of dual-implication materials eliminated inferences under high load conditions, and resulted in reduced effect sizes even under low load conditions.

5.4.1 Comparably (In)Efficient Trait and State Inferences

The core finding of Experiments 8 and 9 is that effects of both spontaneous trait and state inferences were comparable. It thus seems that both traits and states are inferred efficiently from behavioral descriptions, although working memory load did reduce effect sizes, suggesting that social inference processes do require some available cognitive resources. Most importantly, though, both trait and state inferences seemed to be equally dependent on cognitive resources. Our data thus do not support the often presumed dominance of traits in spontaneous impression formation from behavior: At least one other spontaneous social inference may be drawn at least as efficiently as trait inferences. There is thus no evidence that would suggest a hierarchy in spontaneous social inferences. Neither trait nor state inferences seemed to be processed more efficiently than the other, suggesting that both can occur simultaneously and relatively efficiently. This finding is challenging the assumptions about the emergence of correspondence biases: One hypothesis, as outlined above, assumes that people tend to attribute others' behavior to traits because traits are already available (Ferreira et al., 2012; Moskowitz, 2005). Our results challenge this mechanism: If multiple inferences are drawn simultaneously and with a similar degree of efficiency, how then do perceivers eventually show correspondence biases? To be sure, we do not suggest that perceivers do not exhibit correspondence biases, but want to highlight that we do not yet understand which processes occur when perceivers process multiple spontaneous inferences. It remains unclear which processes are at play in a possible integration of multiple inferences to a consciously accessible impression. Our data, however, contribute to a better understanding of these processes. The relative efficiency of both trait and state inferences we observed in the present experiments do not indicate a hierarchy or sequential processing of spontaneous social inferences, at least not for trait and state inferences. A possible process could have been that perceivers may first infer a state, to then generalize to a trait, or vice versa, similar to hierarchies documented in explicit impression formation (Malle & Holbrook, 2012). Our data regarding the efficiency of spontaneous social inferences suggest that this is not the case. In spite of this new evidence contradicting one possible explanation, however, we are left with the fact that our understanding of the social inference process is still very limited. To further it, it is crucial to investigate the mechanics that govern the interplay of multiple simultaneous social inferences from materials that allow for various inferences.

5.4.2 Inefficient Inferences from Complex Materials

In the present experiments, we observed neither spontaneous trait nor state inference effects under working memory load conditions for dual-implication materials, neither in the low load nor high load conditions. Our first results thus suggest that the complexity of dual-implication materials may require more of perceivers' cognitive resources to begin with, as compared to single-implication materials. The combined effects of complex materials and working memory load manipulation may explain why no spontaneous social inferences seemed to have been drawn, and also suggest that spontaneous social inferences from complex materials are actually more resource-dependent than previously assumed. While preliminary, this finding is incompatible with previous assumptions about the efficiency of spontaneous social inferences. The results of Experiments 8 and 9 indicate that the conclusions about the efficiency of social inferences in general that have been drawn from studies employing only single-implication materials may have been premature. It is possible that these assumptions are valid only for behavioral statements that clearly imply a single characteristic. Such behaviors may be less common in everyday life than behaviors that are attributable to multiple characteristics of both the person and the situation. We need to test whether these results replicate, to investigate whether spontaneous state and trait inference effects are robust when the complexity of materials is manipulated.

5.4.3 Directions for Future Research

In order to further our understanding of the social inference process, we thus need to investigate whether our findings regarding the inefficiency of spontaneous social inferences from dual-implication behaviors replicate. In a future experiment, we plan to replicate Experiment 9 and include a no load control condition, in order to ensure that the materials we employ allow for trait and state inferences in principle, without concurrent working memory load. Further, we will include a measure of individual differences in working memory capacity, the operation span task (OSPAN; Turner & Engle, 1989; Unsworth & Engle, 2007). If the OSPAN correlates with participants' performance in the social inference task at different working memory load levels, we can assume that any effects of the working memory load manipulation occurred because of an impact on working memory capacity, and not because of other possible consequences of the working memory load manipulation (e.g., disruption of attention; Wells et al., 2011).

The generalizability of the present results is limited by the fact that, in this series of experiments, we only provided evidence regarding the *efficiency* of spontaneous state inferences. The use of the false recognition paradigm further warrants the assumption that inferences were drawn unintentionally, as an intention to draw inferences would have hindered task performance. However, we have not yet investigated whether state and trait inferences are also comparable with regards to unconsciousness and uncontrollability. In future research, we plan to investigate the relative controllability of spontaneous social inferences. We will manipulate processing goals at behavior encoding. Here, we will implement a suppression condition, in which participants are instructed to actively avoid forming any impression of the actors upon reading the behavioral statements (Shimizu, 2017). In two further specific suppression conditions, participants are instructed to specifically suppress thinking about the general personality of the actors (trait suppression) or the current state of the actors (state suppression). Inference effects in these three suppression conditions will be compared with a control condition, in which no suppression is instructed. In line with previous research (Carlston et al., 1995; Shimizu, 2017), we expect to observe significant trait inference effects that are unaffected by the suppression instructions. Again, we test alternative hypotheses for state inferences: If state inferences are equally uncontrollable as trait inferences, we should observe comparable state inference effect sizes in all four conditions. However, if the formation of spontaneous state inferences is controllable, state inference effects should be significantly reduced in the impression suppression conditions, as compared to the control condition. Another approach we have planned involves manipulating participants' processing goals during retrieval of formed impressions (Uleman et al., 2005). In a false recognition paradigm, participants first familiarize themselves with a number of statements. In the recognition test, participants in the suppression condition are informed that previously formed person impressions may reduce performance in the recognition test and we thus explicitly ask them to not rely on their impressions of the actors. In the impression support condition, participants are informed that previously formed impressions may increase performance in the recognition test and we thus explicitly ask them to rely on their impressions of the actors. Finally, in the control condition, participants will simply be given a standard instruction to try to respond as accurately as possible. In line with previous research, trait inference effects should occur in all three conditions. This pattern of results should be mirrored for state inferences if they are as difficult to control as trait inferences.

Our dual-implication materials, too, are limited in generalizability, in that they clearly imply only two characteristics, namely traits and states. While the generation of multi-implication materials is challenging, the focus on mostly one implied characteristic in past research on spontaneous impression formation from behavior has, as we have outlined above, severely limited our understanding of the social inference process. In future research, we intend to address these stimulus limitations by investigating the efficiency of different spontaneous social inferences using multiple-implication materials, that may imply traits, states, situations, and/or goals simultaneously, to name but a few. Only by employing behaviors that are more representative of behaviors that perceivers may actually encounter in their everyday lives, we can understand the processes involved in forming spontaneous social inferences.

To sum up, in order to further our understanding of the social inference process and its underlying processes, we need to investigate whether differences in the characteristics of automaticity exist between spontaneous trait, state, and other inferences, and whether the effects generalize beyond single-implication behaviors.

6 Conclusion

6.1 Spontaneous State Inferences

Thirty years of research have provided a substantial amount of evidence supporting the notion that people spontaneously infer traits about actors when confronted with their trait-implying behaviors (e.g., Moskowitz & Olcaysoy Okten, 2016; Uleman et al., 2012). The ubiquity of evidence for spontaneous trait inferences has, however, sometimes been misconceived as evidence for the ubiquity of spontaneous trait inferences (Schneid, Carlston, et al., 2015). In three series of studies, we provided consistent evidence that spontaneous inferences are not limited to traits, but may include various mental states as well.

In Chapter 3, we provided consistent evidence across four experiments and two experimental paradigms that participants inferred both traits and states from behavioral descriptions. Effect sizes for state inferences were at least as large as those for trait inferences. Participants inferred states both from behavioral descriptions that implied only a state (single-implication), and from those that implied both a trait and a state (dual-implication). That is, even if materials afforded a trait inference, participants still showed indications that they also inferred a state. In Experiment 5, we observed that participants represented states and traits as functionally different: They rated the likelihood that actors would repeat a behavior as higher if the behavior was paired with a trait, as compared to a state implied by the behavior, respectively.

In Chapter 4, we investigated the relative contributions of associative and inferential processes in spontaneous trait and state inferences. We observed first evidence suggesting that state inferences likely result from inferential processes, at least to a comparable degree as trait inferences.

In Chapter 5, we systematically manipulated working memory load during encoding of behavioral stimuli. We observed that for single-implication materials both, trait and state inferences, occurred even under high load conditions. Effect sizes tended to decrease with increasing load. For dual-implication materials, we observed a state inference effect in low load conditions only, and no trait inference effect in either condition. This result may be interpreted as indicating that spontaneous social inferences from dual-implication materials require more cognitive resources, as compared to single-implication materials.

From the three series of experiments presented in this dissertation, we thus conclude that spontaneous social inferences are not limited to traits, but encompass inferences about actors' more transient mental states as well. These inferences can occur simultaneously with trait inferences, and do so with a comparable degree of efficiency. In the following, I will discuss why it makes sense for people to infer states, how the observation of multiple simultaneous spontaneous inferences challenges the available theorizing, and speculate on how these challenges might be addressed. Finally, I will discuss the limitations of the present research.

6.2 Multiple Spontaneous Social Inferences in the Impression Formation Process

The results of our research are in line with the many findings from different domains of person perception showing that people have a strong tendency to spontaneously make sense of their social environment by going beyond the given information, adding self-generated information and thus drawing causal inferences and attributions. Our findings complement empirical evidence that people can draw spontaneous situation inferences from behavior (Ham & Vonk, 2003; Todd et al., 2011). However, while states certainly have a situational component, they differ from situation inferences quite substantially. In the example *Pete tells the waiter the food tastes good* (adapted from Ham & Vonk, 2003), a situation inference would refer to a property of the situation: The food is *delicious*. In contrast, a state inference would refer to the person: Pete is *pleased*. Our results are also compatible with recent documentations of other spontaneous person inferences, such as goals (e.g., Hassin et al., 2005; Olcaysoy Okten & Moskowitz, 2019; Van Overwalle et al., 2012) or motivations (e.g., Reeder, 2009a, 2009b; Reeder et al., 2001). Without question, goals and motivations represent subcategories of mental states. However, our findings indicate that people infer mental states that go beyond these previously documented inferences. In the example above of Pete complementing the waiter on the food, he might have the goal of *wanting to be nice* which may be relatively independent of his state of being *pleased*. Our studies are the first to demonstrate that people spontaneously draw state inferences beyond behavioral goals.

In this respect, our findings are in line with theoretical accounts on social inferences that presume that behavior attribution includes inferences of subcategories of states. For example, in his multiple inference model, Reeder (2009b) suggests that people might draw spontaneous inferences about both motives and traits simultaneously from behavior. Similarly, Malle (2007, 2001) proposed in his folk theory of mind that people deliberately make sense of other people's behavior mainly in terms of beliefs, desires, and intentions. Our research expands the empirical support of these approaches in two ways. On the one hand, and as previously discussed, we have provided empirical evidence regarding the spontaneity, relative resource-independence, and inferential nature of these inferences. Future research will help to further understand the process characteristics of both trait and state inferences. On the other hand, we enlarge the range of inferred states to affective or emotional states (e.g., happy, surprised, anxious, ashamed), cognitive and attentional states (e.g., focused, interested, distracted, bored) and physiological states (e.g., hungry, full, sick, tired), thus supporting the idea that observers' inferences about actors' states are not limited to motives, beliefs, desires, and intentions, but encompass a great variety of mental states.

The finding of simultaneous and spontaneous trait and state inferences questions previous assumptions regarding the underlying mechanisms of the correspondence bias in person perception: Given that we are able to simultaneously and spontaneously draw trait and state inferences, how come that in deliberate impression formation we still appear to favor dispositional inferences above all others? For example, Krull & Erickson's (1995) model of dispositional inference presumes that behavioral observations are initially and automatically processed only in terms of trait inferences, thus causing a dispositional bias in impression formation when controlled processes do not operate to correct these inferences (e.g., under cognitive load or time pressure; Gawronski, 2004). However, this assumption is difficult to uphold given the multiplicity of spontaneous social inferences.

6.3 Implications for Theory Development

The occurrence of spontaneous state inferences - and other inferences - puts the presumed ubiquity of spontaneous trait inferences into question. The prominent theoretical models of the social inference process do not account for the empirical observations of multiple simultaneous inferences: Some models assume traits to be the default inference, which may then be effortfully corrected for situational constraints (Gilbert et al., 1988; Trope, 1986), some others assume that situations may be the default inference, which may then be effortfully corrected for dispositional information. In one of these models, Krull (1993) argued that the processing sequence and thus relative dominance of different inferences is not fixed, but depends on whether the inferential goal of the perceiver focuses on the acting person or the situation as causes of behavior: Whatever is in the perceiver's focal interest is processed spontaneously and effortlessly, thus dominates automatic inferences and is only revised given sufficient motivation and cognitive resources (Krull & Erickson, 1995). Based on this account, we may conclude that participants in our studies drew simultaneous state and trait inferences because both equally satisfy a predominant goal of forming an impression about the acting person. The model, however, assumes that only a single inference is drawn initially, and fails to account for the observation of multiple simultaneous inferences. This observation is challenging: How do perceivers integrate multiple inferences into a coherent impression?

If we make the reasonable assumption that spontaneous social inferences form the basis for explicit impression formation, research suggests that perceivers tend to describe actors' behaviors in terms of traits (e.g., Gawronski, 2004). For example, our research suggests that perceivers, upon reading about Marvin (who *laughed at the joke*), infer both a state he is in (*amused*) and a trait that describes his personality (*jolly*). If we ask perceivers to form an impression, they are likely to describe Marvin as *a jolly person*. If we assume that perceivers, when asked to form an impression, also engage in spontaneous impression formation, there must be an intermediate process whose outcome is a trait attribution. How this process works remains unclear - several speculations are possible. I gathered the relevant evidence available for subcomponents of the social inference process that might partially determine the presumed integration processes and will describe them in more detail below.

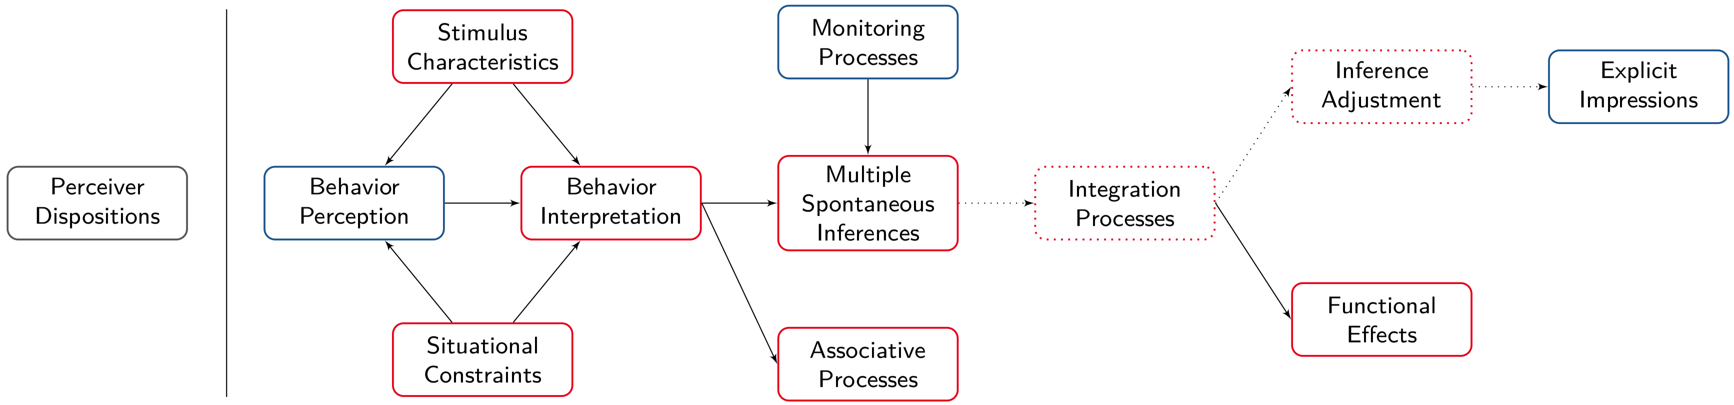

Fig. 6.1.: Proposal of a unifying process framework of social inferences. Red nodes assume spontaneous processes, blue nodes assume deliberate processes. Dashed elements alert to assumptions without empirical support.

6.3.1 A Unifying Process Framework of Social Inferences

In Fig. 6.1, I have gathered the available evidence regarding variables influencing the processes, as well as speculations about processes that might provide insight into how spontaneous social inferences and deliberate impression formation may be intertwined, to supplement the existing social inference models. These ideas are in large part not my own, I merely combine the existing evidence provided elsewhere. The framework assumes the influence of perceiver characteristics on several steps of the social inference process. It further depicts the sequential processes of behavior perception and interpretation already presumed in existing models (Gilbert et al., 1988; Krull & Erickson, 1995). Subsequently, at least three processes have been documented to occur simultaneously: the occurrence of multiple spontaneous inferences, associative processes involved in transference effects, and monitoring processes, which are involved in moderating the influence between spontaneous social inferences and explicit impressions, as I will discuss later (Ferreira et al., 2012). Following the occurrence of multiple spontaneous inferences, the framework posits *integration processes*, which is a placeholder term for the not yet investigated processes involved in managing multiple spontaneous inferences. I will speculate on these processes below. Finally, the framework depicts possible consequences for both explicit and unconscious processing subsequent to spontaneous social inferences, namely explicit impression formation and the umbrella term *functional effects* (such as biases, for instance), which I will define in more detail below.

6.3.1.1 Perceiver Conditions

A multitude of perceiver characteristics have been identified as influential on the social inference process, albeit some of the evidence has so far - and to my knowledge - only been investigated in explicit impression formation. While I often deem it reasonable to assume that these processes also influence spontaneous social inferences, I will point out the lack of empirical data whenever applicable.

Stable person characteristics. Stable person characteristics have been linked to individual tendencies to exhibit increased spontaneous trait inferences. A number of studies have investigated the influence of cultural differences on the tendency to infer traits from behavior, and to attribute traits as causes of behavior, as opposed to inferences or attributions about situational determinants of behavior (for an overview, see Bott et al., 2022). However, the available evidence provided only sample comparisons, not actual investigations of cultural differences. While in our meta-analysis, we could not confirm that these sample comparisons are robust, it does seem plausible that cultural variables impact the inference and/or integration processes presumed by the framework. For example, cultures with a predominantly *independent* person construal emphasize individuals' uniqueness and autonomy in behavioral decisions, whereas cultures with predominantly *interdependent* person construal emphasize effects of the social context on persons and accentuate the individual's social and situational connectedness (e.g., Triandis, 2018). Given these culturally acitivated tendencies to explain behaviors differently, it seems plausible that they influence the inference process in different ways. Past research has focused mostly on traits and/or situations, empirical data regarding cultural differences in drawing and integrating state inferences in the social inference process is lacking.

Another concept relevant to the formation of STI (and thus possibly also for state and other inferences) are implicit personality theories (Chiu et al., 1997; Dweck et al., 1995). The idea behind these theories is that people either believe that attributes (such as intelligence or moral character) are fixed, stable entities (entity theory), or dynamic, malleable, and developable (incremental theory). People with different implicit personality theories have been shown to tend to attribute behaviors differently, either to traits (e.g., "I failed the test because I'm dumb") or to circumstance (e.g., "I failed the test because of my strategy"; Chiu et al., 1997). We measured participants' implicit personality theories in Experiment 2, but did not find any correlation with trait- or state-inference effects. It is plausible that effects of participants' implicit personality theories affect the social inference process at a later stage, in that participants with an entity theory about others' personality show a stronger bias towards trait attributions in integrating multiple inferences, as compared to participants with an incremental theory. This assumption would be testable by asking participants to explicitly form impressions of actors after their spontaneous inferences have been measured.

A related concept is personal need for structure (PNS; Neuberg & Newsom, 1993), a concept which describes a desire for clarity, certainty, and aversion to ambiguity. Participants with high PNS have been shown to exhibit a higher tendency to draw trait inferences, as compared to participants with low PNS (Moskowitz, 1993). In sum, there are a number of stable person characteristics that might differentially influence perceivers' tendencies to spontaneously infer certain social inferences, and/or bias their integration and correction processes. Although our own experiments have not shown correlations between individual differences and trait vs. state inferences, further research needs to be conducted on the matter, as past research on spontaneous trait inferences actually has documented such correlations.

Variable person characteristics. Rim et al. (2009) measured participants' construal level. One aspect of construal level describes individual dispositions to think in abstract vs. concrete terms about events, objects, and importantly, other people (abstract vs. concrete mindset; Trope & Liberman, 2010). Participants who were primed to adopt an abstract mindset showed increased spontaneous trait inferences, as compared to participants who were primed to adopt a concrete mindset. Traits are arguably abstract concepts, while states may represent more concrete representations of a person's dispositions. It thus seems plausible that different mindsets, circumstantially or even chronically primed, might moderate participants' integration of different inferences, possibly by assigning weights. In our research, we assessed participants' chronic construal level mindset (Experiment 2), but did not find correlations between construal level and either trait or state inferences. Further research is needed to investigate. In a future study, we aim at manipulating participants' construal level mindset before presenting them with a social inference task using multiple-implication materials. We would, despite our failure to show effects for chronic construal level mindsets on social inferences, expect participants in an abstract mindset to have a stronger tendency to infer traits, as compared to states or situational properties, for instance. Again, it is also plausible that effects of construal level affect the social inference process at a later stage, in that participants in an abstract mindset show a stronger bias towards trait impressions in integrating multiple inferences. This assumption would be testable by asking participants to explicitly form impressions of actors after their spontaneous inferences have been measured.

Perceivers' current mood is another variable person characteristic that has been shown to influence the social inference process. (Wang et al., 2015) showed that participants who had watched a clip from *Mr. Bean* (positive mood induction), spontaneously inferred traits from behavioral statements in a later probe recognition task. Participants who had watched a clip of an earthquake (negative mood induction), did not. (Shi et al., 2019) showed that participants in negative mood did infer traits from behavioral statements, but more so if they were angry (after watching a video clip about a war trial), as compared to sad participants (who watched earthquake footage). To my knowledge, no empirical evidence has been gathered regarding differential influences of perceivers' mood on the occurrence and processing of spontaneous state inference (or other inferences, for that matter). Given the quite strong effects that Wang et al. found regarding perceivers' mood on STIs, it is plausible that such differences have effects on other inferences and/or perceivers' ability to integrate multiple inferences as well. These assumptions are testable, in future research we aim at investigating the impact of different mood inductions on both the occurrence of spontaneous social inferences and their integration.

A variable person characteristic with consequential effects on the social inference process are perceivers' processing goals. Krull (1993) presented participants with a videotape of an interview with an anxious person and measured participants' ratings of the actor's trait anxiety. Half of the participants were instructed to focus on the situation, namely on how anxiety-provoking the questions were. The other half were instructed to focus on how anxious the actor was. He found that the processing goal altered the degree to which participants rated the actor's trait anxiety: Perceivers who focused on the person rated the actor's trait anxiety as higher than perceivers who focused on the situation. A working memory load manipulation (concurrent rehearsal of digits) moderated the effects: Busy person-focused perceivers' trait ratings increased, while busy situation-focused perceivers' trait ratings decreased. Krull's experiment was conducted with explicit ratings as dependent measure. However, the results demonstrate the importance of processing goals on the inference process: Krull demonstrated that processing goals are able to alter explicit ratings. In the framework, I assume that between behavior perception and explicit impressions (such as trait ratings), spontaneous inferences occur. Do processing goals also alter the occurrence of spontaneous inferences, or do they influence monitoring processes which in turn result in integration processes that favor impressions compatible with the processing goal? Rim et al. (2013) demonstrated that, in fact, processing goals may also alter spontaneous inferencing: They primed participants with an affiliation goal vs. no goal and observed that participants were more likely to falsely recognize positive traits in a subsequent false recognition task. These results suggest that processing goals may also influence the occurrence of spontaneous social inferences. It remains unclear how processing goals might impact multiple social inferences. In future research, we aim at investigating whether or to what extent different processing goals alter effect sizes of simultaneously occurring trait, state, and possibly further spontaneous inferences such as goal or situation inferences. Even more insight into the social inference process might be gained by assessing both spontaneous inferences and subsequent explicit impressions, and comparing the effect of processing goals on effect sizes of multiple inferences at either step of the inference process: If effect sizes for spontaneous inferences correlate with effect sizes for explicit ratings, this would indicate that the integration processes assumed in the framework may be influenced by processing goals as well. If, however, effect sizes differ between spontaneous inferences and explicit ratings, integration and/or correction processes may be quite influential.

6.3.1.2 Behavior Perception & Interpretation

A number of factors determine how perceivers perceive and interpret behavior. In the framework, I have devided these factors into properties of the stimuli, consisting of both characteristics of the actor and the behavior, and properties of the situation, by which I mean possible situational constraints on the perceiver during behavior perception and interpretation, such as concurrent working memory load.

Actor Characteristics. A large number of actor characteristics have been demonstrated to influence impressions of them. For instance, face perception research has shown time and again that facial features result in spontaneous impression formation (e.g., Oosterhof & Todorov, 2008; Zebrowitz, 2017). As behavior is strongly linked to actors who have faces, at least for human actors, impressions from face and behavior perceptions are likely to interact. However, few studies actually investigated this interaction (see Li et al., 2017; Shen et al., 2020, for exceptions). However, the existing data focused entirely on trait ratings. Given the readiness with which perceivers infer emotional states from facial expressions (e.g., Ekman, 1997), a connection between facial expressions and spontaneous state inferences seems obvious. In future studies, we aim at presenting participants with state-implying statements paired with a portrait that depicts an actor in a matching vs. non-matching state. We expect to observe interactions between spontaneous impressions from faces and behaviors: Matching states should lead to an increased effect size for spontaneous state inferences, whereas non-matching states should lead to reduced effect sizes.

Further actor characteristics include prior information that might be available about the actor, such as physical appearance - which has thus far only been demonstrated to influence explicit impression formation processes (Trope, 1986). Prior behavioral information about the actor does not inhibit inferences of other traits from new behaviors, nor are spontaneous trait inferences updated when the initial behavioral information is transformed (Olcaysoy Okten et al., 2019) - suggesting that spontaneous trait inferences are not affected by previously drawn inferences. Various inferences from sequentially perceived behaviors may nonetheless need to be integrated to form explicit impressions.

Further research has documented that perceivers tend to attribute behavior differently depending on how familiar they are with an actor, mediated by how important they perceive the actor to them (Idson & Mischel, 2001): Unfamiliar actors are more likely to be described by trait ascriptions, whereas familiar actors are more likely to be described in terms of their goals. While no research on the impact of familiarity on spontaneous impression formation exists to my knowledge, it is plausible that familiarity only influences the more deliberate inference correction, while spontaneous inference processes remain unaffected - much like is the case for updating inferences. That is, if you knew Marvin (who *laughed at the joke*) as a rather depressed individual, you might still draw the inference *jolly* from his behavior, which you could then contrast with your prior knowledge about him and deliberately correct your inference to arrive at a final conclusion. On the other hand, Rim et al. (2009) observed that STI effects were stronger for spatially distant actors, as compared to spatially near others, which could support the notion that familiarity might also impact spontaneous inferences. One might further hypothesize that familiarity has differential impact on trait vs. state inferences: State inferences should occur more frequently for familiar actors as compared to trait inferences. Note that these are merely speculations, and further research on spontaneous social inferences about familiar vs. unfamiliar actors is needed to substantiate them.

Likewise, research on the linguistic intergroup bias (LIB; Maass et al., 1989) has provided evidence that people represent others' behavior differently depending on its evaluation and whether it is performed by ingroup vs. outgroup members: Whereas desired ingroup behavior and undesired outgroup behavior is represented in more abstract terms (e.g., by using trait adjectives), undesired behavior of ingroup members and desired behavior of outgroup members is represented in less abstract forms (e.g., by using state adjectives). Although we are not aware of any research investigating if the LIB applies to spontaneous impression formation as well, it is conceivable that observers apply state and trait inferences differently to the same behaviors based on the group memberships of actor and perceiver. In a related notion, research showed that spontaneous trait inferences are more likely to be made for stereotype-consistent behaviors, as compared to stereotype-inconsistent behaviors, whereas spontaneous situational inferences tend to follow the opposite pattern (Ramos et al., 2012; Wigboldus et al., 2004).

Behavior Characteristics. Several characteristics of behaviors have been shown to impact the social inference process. Marcelo et al. (2019) showed that variation of specific linguistic features of a behavioral statement impact STI effect sizes. What is more, trait inferences occur for both rather extreme behaviors and more mundane ones (see Chapter 2.3; Levordashka & Utz, 2017). Trait inferences should be more likely to occur for extreme behaviors, as fewer alternative explanations may be available for extreme behaviors - it is hard to attribute *painting a swastika on a synagogue wall* to anything else but bigotry (Winter & Uleman, 1984). This argument is in line with the classic attribution theory devised by Kelley (1967, see Chapter 2), which posits that, when asked to form an impression, perceivers analyze behavior with regard to its consistency, consensus, and distinctiveness. Olcaysoy Okten and Moskowitz (2018) manipulated the consistency (e.g., Marvin *always* laughs vs. *just this once* laughed at the joke) and distinctiveness (e.g., *everyone else* vs. *no one else* laughed at the joke) of behavioral statements. They observed that trait attributions were more likely to occur for behaviors high in consistency or low in distinctiveness, whereas goal attributions were more likely to occur for behaviors high in distinctiveness and low in consistency. It remains to be determined whether perceivers engage in such covariation analyses in spontaneous impression formation as well, or whether they only result from deliberate processes that correct initial inferences. It does seem plausible, however, that additional information provided in behavioral statements may also impact the likelihood of certain inferences. For example, linguistic markers of temporal consistency such as in the example above. Thus, behaviors and behavioral statements themselves may present strong constraints with regard to the inferences that can be drawn from them.

6.3.1.3 Multiple Inferences and Integration Processes

Thus far, we have reviewed the available evidence on factors that might influence the social inference process. The research presented in this paper along with others (Ham & Vonk, 2003; Olcaysoy Okten & Moskowitz, 2019; Todd et al., 2011) documents the occurrence of multiple, simultaneous spontaneous inferences. It remains unclear how perceivers manage this multitude of simultaneously activated concepts. Perceivers may choose among initial inferences, possibly guided by their processing goals, which may lead them to favor a trait, state, situation, goal, or different inference about the actor. This process might be more or less consciously accessible, efficient, and uncontrollable (or not). Perceivers may integrate inferences into a combined impression, possibly weighing different inferences and judging their proportionate contribution to causing the observed behavior. Several theoretical accounts already posit some sort of integration processes. Read and Miller (2005) assume in their social dynamics model that perceivers engage in sequential processing of features of actor, behavior, and situation, which are integrated into an impression of a 'scenario', which identifies actors' goals and motives. These goals and motives, in turn, are then generalized to a trait term. While this model goes beyond the assumption of default trait inferences, the evidence for multiple simultaneous spontaneous inferences provided in this work does not integrate well with this model. Trope and Gaunt (2000) describe an integration process model for explicit impression formation. They posit that, once behavior is identified (Marvin *laughs at the joke*), perceivers consider both dispositional and situative information and engage in evaluating several hypotheses about the actor against each other (is Marvin a *jolly person*? Is it the *joke* that is particularly funny?). According to Trope and Gaunt's integration model, this comparative evaluation of multiple, simultaneously available explanations of behavior is an effortful and deliberate process. Cognitive load may prevent perceivers from fully integrating multiple explanations of behavior. Perceivers should thus overweigh any explanation that happens to be strongly activated, and underweigh all other explanations, when they do not spare the cognitive resources or motivation to engage in the effortful correction process. This failure to integrate multiple explanations could then result in correspondence biases. It seems plausible that an integration process could explain the gap between multiple spontaneous social inferences and explicit impressions. However, Trope and Gaunt tested their hypotheses on single-implication materials, using explicit trait ratings and trait impression instructions only. Empirical support for the applicability of the integration assumptions for spontaneous impression formation is thus still lacking. In future research, we need to investigate if and how perceivers integrate multiple spontaneous inferences, and how the integration processes work.

6.3.1.4 Monitoring & Associative Processes

As discussed at length in Chapter 4, social inferences consist at least in part of associative processes, in which characteristics implied by a behavior may be associated with whomever or whatever may be present during encoding of behavioral information. These processes are assumed to run parallel to inferential processes.

Ferreira et al. (2012) posit that monitoring processes may explain how spontaneous inferences and explicit impressions may be linked. They presented participants with trait-implying behaviors and instructed them to either memorize them or to form an impression. Subsequently, they measured trait inferences and behavior recall. Ferreira et al. observed that participants in the impression formation condition recalled more behaviors and clustered the recalled behavior more according to the implied traits than participants in the memory formation condition did. In other words, participants under memory instructions did produce

STIs, but the inferred traits were not used to organize the behavioral information. Only participants with impression formation goals showed trait clustering at recall. Ferreira et al. suggested that impression formation goals might lead perceivers to engage in deliberate monitoring of spontaneous inferences, which makes use of spontaneously generated inferences in subsequent tasks more likely. They also posit, however, that spontaneous trait inferences may be used as a basis for explicit impression formation (Ferreira et al., 2012, p. 10). If I asked you why Marvin laughed at the joke, the accessibility of the already generated spontaneous inference *jolly* might facilitate generating an explanation (Marvin laughed at the joke *because he is a jolly person*). However, as we have seen, it is likely that multiple spontaneous inferences are accessible simultaneously, and we have yet to investigate how these multiple spontaneous inferences link to explicit impression formation.

6.3.1.5 Inference Adjustment & Explicit Impressions

Most classic models already assume that initial inferences may be corrected if perceivers are motivated, have cognitive resources available and information providing alternative explanations is available to them. If these conditions are met, perceivers may engage in an effortful inference adjustment process. Krull and Erickson (1995) include spontaneous trait and situational inferences in their model. They assume that they are drawn effortlessly, and may be corrected effortfully, if perceivers are motivated and spare cognitive capacity. They do not make assumptions about how this effortful correction process is linked to spontaneous inferences, and this question remains unanswered empirically (but see Ferreira et al., 2012): Do perceivers first need to access unconscious inferences to engage in correction processes? Or may correction also occur unconsciously, that is, more in the sense of integration processes, as posited in my framework? It seems reasonable to assume that any factor that might be considered in a conscious correction of an initial inference is a property of the stimulus (person traits, states, or situational properties) about which a simultaneous spontaneous inference has already been drawn. In the case of Marvin, we could assume that the trait *jolly*, the state *amused*, the situational property *funny*, and possibly even further spontaneous inferences could be activated simultaneously. If multiple spontaneous inferences that offer alternative explanations can occur simultaneously, it is possible that the assumed correction processes (Krull & Erickson, 1995) actually occur spontaneously, without deliberate intent. In light of the evidence provided by Ferreira et al. (2012), we may speculate that processing goals moderate these processes. Perceivers with impression formation goals may engage in inference monitoring, which may allow them to consciously access spontaneous inferences, which in turn may enable them to deliberately correct them.

The interplay of spontaneous social inferences and deliberate impression formation seems very plausible, yet very few empirical data exist on this issue. In order to further our understanding of the social inference process, it is crucial to tackle the procedural details of the link between (multiple) spontaneous inferences and explicit impression formation. Above, I have stated some speculations, which are testable. First of all, we need to investigate whether spontaneous social inferences correspond to explicitly formed impressions from behavior. This is testable by presenting participants with behavioral descriptions implying multiple person and/or situational characteristics, and measuring spontaneous inferences, for instance using the false recognition paradigm. Subsequently, participants are asked to form an impression of the actors, by either asking them to state their impression in a free-text format, or by rating a multitude of characteristics for each actor, including the properties implied by the statement. If spontaneous inference effects correlate with explicit expressions, we can assume that spontaneous inferences form the basis for explicit impression formation. In a second step, we would need to ensure that no deliberate adjustment or correction processes can occur between spontaneous inferencing and explicit impression formation, for instance by introducing a working memory load condition that should occupy cognitive resources and therefore limit participants' ability to engage in effortful processing - for instance by asking them to memorize a sequence of digits while engaging in the explicit impression formation task. In a third step, it would be interesting to investigate the role of monitoring processes for multiple spontaneous inferences and their link to explicit impression formation by introducing an impression formation instruction for some participants, as opposed to memorization instructions. Measuring both spontaneous inferences and explicit impressions is tricky, however. All measures of spontaneous inferences developed to date involve presenting participants with trait- or state-implying behaviors and subsequently probing them with these implied traits or states. If participants are asked to form explicit impressions about the actors, this impression may or may not be influenced by the probe participants were probed with earlier. Nonetheless, the interplay between spontaneous and explicit impression formation merits investigation and the development of suitable research procedures.

6.3.2 Functional Effects

Besides the focus on understanding the psychological mechanisms of inference processes in spontaneous and controlled impression formation, considerations regarding the functional value of spontaneous impressions warrant further thought and development: How do spontaneous inferences that people draw from behavior affect their further judgment, decision making, and interaction behavior? Many scholars have offered a functional perspective of spontaneous impression formation processes, presuming that people draw inferences because they can make use of them for forming expectations and adapting their own (interaction) behavior. This served as an explanation for a frequently presumed dominance of trait inferences - because inferring a stable personal disposition as cause of behavior allows predictions of future behavior and thus adaptation for future interactions (e.g., Newman, 1996). However, such functional consequences of first impression formation have only rarely been studied (but see Costabile & Madon, 2019; McCarthy & Skowronski, 2011a, for exceptions). The results of our research also question this simplified approach: While we show that participants simultaneously draw state and trait inferences from behavior, results of Experiment 5 demonstrate that these same inferences are related to different expectations regarding the future and the likelihood that the same kind of behavior will be repeated. Thus, people may not only draw spontaneous inferences in order to predict future behavior, but also to infer temporal characteristics whose predictive value is limited. Nussbaum et al. (2003) provide evidence that this differentiation applies to the use of explicit inferences as well.

6.4 Boundary Conditions of Spontaneous Social Inferences

Finally, I must describe the limitations of the current research. We interpreted the results observed in the nine experiments reported in this dissertation as evidence for the occurrence of spontaneous state inferences from behavior, while it would technically only be correct to assume that these inferences occur for (a) some, carefully selected behaviors and (b) in text comprehension, as all stimuli employed in this and similar research were written behavioral statements and/or descriptions. I will discuss both limitations in detail below.

6.4.1 Behavior Selection

It is imperative that spontaneous social inferences be measured indirectly to observe actual spontaneous impression formation processes, which may or may not differ from deliberate impression formation, as we discussed before. In indirect measures of spontaneous inferences, participants are typically presented with inference probes and their differential reactions are observed. In all experimental paradigms that have been developed so far, researchers thus need to know beforehand which inference perceivers are likely to draw from any given behavior, to measure reactions to that particular inference probe. This severely limits the choice of stimuli for which spontaneous impression formation can be assessed, namely behaviors for which a high proportion of participants generate the same deliberate inference in a pretest. This eliminates a vast number of behaviors: For example, of the 484 statements we developed, only 60 achieved a consensus greater than 50% in the pretest. This lead us to exclude 88% of the materials we had initially developed, further limiting the selection of behaviors already specifically created to imply a trait and/or state (see also Malle & Holbrook, 2012).

What is more, the complexity of behaviors used as stimuli in spontaneous social inference research is reduced, as compared to everyday behavior. Chapters 4 and 5 revealed notable differences between single- and dual-implication behaviors: The relatively complex dual-implication behaviors appeared to require more cognitive resources as compared to single-implication behaviors. Everyday behaviors are often even more complex, in that they may imply various characteristics to differing degrees. The spontaneous impression formation process for such behaviors might be substantially different from the one for behaviors with relatively clear implications typically used in research. Perceivers might require even more cognitive resources to disambiguate the behavior, or, to the contrary, identify the behavior as non-diagnostic with regards to their predominant goal of forming an impression and interrupt the spontaneous impression formation process entirely. These speculations reveal how limited the scope of research concerning spontaneous impression formation from behavior is to date, and what challenges still remain unaddressed.

6.4.2 Text Comprehension

Another strong limitation to the generalizability of the present research and research on spontaneous social inferences from behavior in general is the use of written behavior descriptions. While text has certainly gained relevance in our lives in times of digital communication such as instant messaging or social network posts, Uleman, Newman, and Moskowitz (1996, p.146) once described the conditions under which spontaneous impressions are supposed to occur in daily life as "uneventful people watching". The written stimuli typically employed in research, however, differ substantially from this idea, as others have noted before (Bassili, 1993; Fiedler & Schenck, 2001; Fiedler et al., 2005; Malle & Holbrook, 2012). Most of the time, we use language and specifically written text to inform others about things we deem important or informative, only rarely to convey information accurately and completely (Grice, 1975). Following this idea of the cooperative nature of language, presenting behavior descriptions to participants suggests that these descriptions are meaningful with regard to the actor, which might itself trigger impression formation processes in participants. The choice of words also carries additional meaning and presupposes specific attributions. For example, research on the linguistic-category model (Fiedler & Semin, 1988) suggests that different verb classes practically determine the attributions from statements that contain these verbs.

Thus, text is already loaded with interpretations. On the other hand, text is quite simple, and the amount of information conveyed by it is quite limited, as compared to real behavior observation. The behavior descriptions frequently used in research typically consist of short sentences with easy syntax, and contain the behavioral information but not much else. During behavior observation in real life, however, perceivers may observe the same behavior as described in a written description, but need to devote attention specifically to the behavior, while a vast amount of other information is also available, such as gestures, facial expressions, clothing, the surroundings, other people, and much more. There are few papers investigating spontaneous social inferences in more ecologically valid settings, such as in pictures (Fiedler & Schenck, 2001) and video clips (Fiedler et al., 2005), as well as the interplay of multiple explicit social inferences from video clips (Malle & Holbrook, 2012). However, most of the evidence on spontaneous social inferences available to date, including the research presented in this dissertation, is based on studies relying on written stimuli, and it remains unclear whether the processes investigated actually generalize to spontaneous impression formation that occurs in everyday perception of behavior. First results of an experiment we conducted on spontaneous trait and state inferences from video clips, however, do suggest that both traits and states may be inferred from short video clips.

6.5 Final Conclusion

The process of person perception from behavior is very complex, and it is likely that a multitude of factors influence which inferences a specific perceiver draws from a specific behavior shown by a specific actor. Caution is advised to not over-generalize the current findings to person perception in general. Further research is needed to investigate which processes contribute to social inferences from behavior and to better understand the systematic behind the aforementioned stimulus constraints and how target and perceiver effects may differentially impact spontaneous state and trait inferences in deliberate and spontaneous impression formation from behavior.

In light of the multitude of spontaneous social inferences, we need to further develop our theorizing of impression formation from behavior to understand why and how we tend to overly attribute observed behaviors to peoples' dispositions. We are confident that the presented research constitutes an important first step in this direction.

Bibliography

Ames, D. R. (2004). Inside the mind reader's tool kit: Projection and stereotyping in mental state inference. *Journal of Personality and Social Psychology*, *87*(3), 340–353. https://doi.org/10.1037/0022-3514.87.3.340

Apperly, I. A., Riggs, K. J., Simpson, A., Chiavarino, C., & Samson, D. (2006). Is belief reasoning automatic? *Psychological Science*, *17*(10), 841–844. https://doi.org/10.1111/j.1467-9280.2006.01791.x

Aronson, E., Wilson, T., & Sommers, S. (2019). *Social psychology*. Pearson.

Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258–290.

Bainbridge, W. A., Isola, P., & Oliva, A. (2013). The intrinsic memorability of face photographs. *Journal of Experimental Psychology: General*, *142*(4), 1323–1334. https://doi.org/10.1037/a0033872

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, efficiency, intention, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 1. basic processes* (2nd ed., pp. 1–40). Erlbaum. https://doi.org/10.1007/s00572-005-0022-9

Bassili, J. N. (1989). Traits as action categories versus traits as person attributes in social cognition. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 61–89). Lawrence Erlbaum Associates, Inc.

Bassili, J. N. (1993). Procedural efficiency and the spontaneity of trait inference. *Personality and Social Psychology Bulletin*, *19*(2), 200–205. https://doi.org/10.1177/0146167293192009

Bassili, J. N., & Smith, M. C. (1986). On the spontaneity of trait attribution. Converging evidence for the role of cognitive strategy. *Journal of Personality and Social Psychology*, *50*(2), 239–245. https://doi.org/10.1037/0022-3514.50.2.239

Bauman, C. W., & Skitka, L. J. (2010). Making attributions for behaviors: The prevalence of correspondence bias in the general population. *Basic and Applied Social Psychology*, *32*(3), 269–277. https://doi.org/10.1080/01973533.2010.495654

Biner, P. M., Angle, S. T., Park, J. H., Mellinger, A. E., & Barber, B. C. (1995). Need state and the illusion of control. *Personality and Social Psychology Bulletin*, *21*(9), 899–907. https://doi.org/10.1177/0146167295219004

Bott, A., Brockmann, L., Denneberg, I., Henken, E., Kuper, N., Kruse, F., & Degner, J. (2022). *Spontaneous trait inferences from behavior: A systematic meta-analysis.* Manuscript submitted for publication.

Brown, R. D., & Bassili, J. N. (2002). Spontaneous trait associations and the case of the superstitious banana. *Journal of Experimental Social Psychology*, *38*(1), 87–92. https://doi.org/10.1006/jesp.2001.1486

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, *64*(2), 123–152.

Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward unbiased measurement of conscious and unconscious memory processes within the process dissociation framework. *Journal of Experimental Psychology: General*, *124*(2), 137–160. https://doi.org/10.1037/0096-3445.124.2.137

Campbell, J. I. D., & Thompson, V. A. (2012). MorePower 6.0 for ANOVA with relational confidence intervals and Bayesian analysis. *Behavior Research Methods*, *44*(4), 1255–1265. https://doi.org/10.3758/s13428-012-0186-0

Carlston, D. E., Skowronski, J. J., & Sparks, C. (1995). Savings in relearning: II. On the formation of behavior-based trait associations and inferences. *Journal of Personality and Social Psychology*, *69*(3), 420–35. https://doi.org/10.1037/0022-3514.69.3.429

Carlston, D. E., & Skowronski, J. J. (1994). Savings in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology*, *66*(5), 840–856. https://doi.org/10.1037/0022-3514.66.5.840

Carlston, D. E., & Skowronski, J. J. (2005). Linking versus thinking: Evidence for the different associative and attributional bases of spontaneous trait transference and spontaneous trait inference. *Journal of Personality and Social Psychology*, *89*(6), 884–98. https://doi.org/10.1037/0022-3514.89.6.884

Chiu, C.-y., Hong, Y.-y., & Dweck, C. S. (1997). Lay dispositionism and implicit theories of personality. *Journal of Personality and Social Psychology*, *73*, 19–30. https://doi.org/10.1037/0022-3514.73.1.19

Chun, W. Y., Spiegel, S., & Kruglanski, A. W. (2002). Assimilative behavior identification can also be resource dependent: The unimodel perspective on personal-attribution phases. *Journal of Personality and Social Psychology*, *83*(3), 542–555. https://doi.org/10.1037/0022-3514.83.3.542

Conrey, F. R., Gawronski, B., Sherman, J. W., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, *89*(4), 469–487. https://doi.org/10.1037/0022-3514.89.4.469

Cook, A. E., & O'Brien, E. J. (2017). Fundamentals of inferencing during reading. *Language and Linguistics Compass*, *11*(7), 1–16. https://doi.org/10.1111/lnc3.12246

Costabile, K. A. (2016). Narrative construction, social perceptions, and the situation model. *Personality and Social Psychology Bulletin*, *42*(5), 589–602. https://doi.org/10.1177/0146167216636627

Costabile, K. A., & Madon, S. (2019). Downstream effects of dispositional inferences on confirmation biases. *Personality and Social Psychology Bulletin*, *45*(4), 557–570. https://doi.org/10.1177/0146167218789624

Crawford, M. T., Skowronski, J. J., & Stiff, C. (2007). Limiting the spread of spontaneous trait transference. *Journal of Experimental Social Psychology*, *43*(3), 466–472. https://doi.org/10.1016/j.jesp.2006.04.003

Crawford, M. T., Skowronski, J. J., Stiff, C., & Scherer, C. R. (2007). Interfering with inferential, but not associative, processes underlying spontaneous trait inference. *Personality and Social Psychology Bulletin*, *33*(5), 677–690. https://doi.org/10.1177/0146167206298567

De Houwer, J., Van Dessel, P., & Moran, T. (2020). Chapter three - attitudes beyond associations: On the role of propositional representations in stimulus evaluation. In B. Gawronski (Ed.), *Advances in experimental social psychology* (pp. 127–183). Academic Press. https://doi.org/https://doi.org/10.1016/bs.aesp.2019.09.004

Diergarten, A. K., & Nieding, G. (2016). Online emotional inferences in written and auditory texts: a study with children and adults. *Reading and Writing*, *29*(7), 1383–1407. https://doi.org/10.1007/s11145-016-9642-z

Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*(4), 568–581.

Dweck, C. S., Chiu, C., & Hong, Y. (1995). Implicit theories and their role in judgments and reactions : A word from two perspectives. *Psychological Inquiry*, *6*(4), 267–285. https://doi.org/10.1207/s15327965pli0604

Ekman, P. (1997). Should we call it expression or communication? *Innovation: The European Journal of Social Science Research*, *10*(4), 333–344. https://doi.org/10.1080/13511610.1997.9968538

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

Ferreira, M. B., Garcia-Marques, L., Hamilton, D., Ramos, T., Uleman, J. S., & Jerónimo, R. (2012). On the relation between spontaneous trait inferences and intentional inferences: An inference monitoring hypothesis. *Journal of Experimental Social Psychology*, *48*(1), 1–12. https://doi.org/10.1016/j.jesp.2011.06.013

Fiedler, K., & Schenck, W. (2001). Spontaneous inferences from pictorially presented behaviors. *Personality and Social Psychology Bulletin*, *27*(11), 1533–1546. https://doi.org/10.1177/01461672012711013

Fiedler, K., & Semin, G. R. (1988). On the causal information conveyed by different interpersonal verbs: The role of implicit sentence context. *Social Cognition*, *6*(1), 21–39. https://doi.org/10.1521/soco.1988.6.1.21

Fiedler, K., Watling, M., Menges, J. I., & Schenck, W. (2005). Priming trait inferences through pictures and moving pictures: The impact of open and closed mindsets. *Journal of Personality and Social Psychology*, *88*(2), 229–244. https://doi.org/10.1037/0022-3514.88.2.229

Fleeson, W., & Jayawickreme, E. (2015). Whole trait theory. *Journal of Research in Personality*, *56*, 82–92. https://doi.org/10.1016/j.jrp.2014.10.009

Frank, M. G., & Ekman, P. (1996). Physiological effects of the smile. *Directions in Psychiatry*, *16*(25), 1–8.

Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European Review of Social Psychology*, *15*(1), 183–217. https://doi.org/10.1080/10463280440000026

Gawronski, B., & Bodenhausen, G. V. (2011). The associative–propositional evaluation model: Theory, evidence, and open questions. In J. M. Olson & M. P. Zanna (Eds.). Academic Press. https://doi.org/https://doi.org/10.1016/B978-0-12-385522-0.00002-0

Gawronski, B., & Bodenhausen, G. V. (2018). Evaluative conditioning from the perspective of the associative-propositional evaluation model. *Social Psychological Bulletin*, *13*(3), 1–33. https://doi.org/10.5964/spb.v13i3.28024

Gernsbacher, M. A., Goldsmith, H. H., & Robertson, R. R. (1992). Do readers mentally represent characters' emotional states? *Cognition and Emotion*, *6*(2), 89–111. https://doi.org/10.1080/02699939208411061

Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, *54*(5), 733–740. https://doi.org/10.1037/0022-3514.54.5.733

Goren, A., & Todorov, A. (2009). Two faces are better than one: Eliminating false trait associations with faces. *Social Cognition*, *27*(2), 222–248. https://doi.org/10.1521/soco.2009.27.2.222

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, *101*(3), 371–395. https://doi.org/10.1037/0033-295x.101.3.371

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Speech acts* (pp. 41–58). Academic.

Ham, J., & Vonk, R. (2003). Smart and easy: Co-occurring activation of spontaneous trait inferences and spontaneous situational inferences. *Journal of Experimental Social Psychology*, *39*(5), 434–447. https://doi.org/10.1016/S0022-1031(03)00033-7

Hamaker, E. L., Nesselroade, J. R., & Molenaar, P. C. M. (2007). The integrated trait-state model. *Journal of Research in Personality*, *41*(2), 295–315. https://doi.org/10.1016/j.jrp.2006.04.003

Hamlin, J. K., Ullman, T., Tenenbaum, J., Goodman, N., & Baker, C. (2013). The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Developmental Science*, *16*(2), 209–226. https://doi.org/10.1111/desc.12017

Hassin, R. R., Aarts, H., & Ferguson, M. J. (2005). Automatic goal inferences. *Journal of Experimental Social Psychology*, *41*(2), 129–140. https://doi.org/10.1016/j.jesp.2004.06.008

Heider, F. (1944). Social perception and phenomenal causality. *Psychological Review*, *51*(6), 358–374. https://doi.org/10.1037/h0055425

Heider, F. (1958). *The psychology of interpersonal relations*. John Wiley & Sons.

Hewstone, M., Stroebe, W., & Jonas, K. (2016). *An introduction to social psychology*. John Wiley & Sons.

Hoffman, C., Mischel, W., & Mazze, K. (1981). The role of purpose in the organization of information about behavior: Trait-based versus goal-based categories in person cognition. *Journal of Personality and Social Psychology*, *40*(2), 211–225. https://doi.org/10.1037/0022-3514.40.2.211

Hooper, N., Erdogan, A., Keen, G., Lawton, K., & McHugh, L. (2015). Perspective taking reduces the fundamental attribution error. *Journal of Contextual Behavioral Science*, *4*(2), 69–72. https://doi.org/10.1016/j.jcbs.2015.02.002

Hopthrow, T., Hooper, N., Mahmood, L., Meier, B. P., & Weger, U. (2017). Mindfulness reduces the correspondence bias. *Quarterly Journal of Experimental Psychology*, *70*(3), 351–360. https://doi.org/10.1080/17470218.2016.1149498

Idson, L. C., & Mischel, W. (2001). The personality of familiar and significant people: The lay perceiver as a social - cognitive theorist. *Journal of Personality and Social Psychology*, *80*, 585–596. https://doi.org/10.1037/0022-3514.80.4.585

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*(5), 513–541. https://doi.org/10.1016/0749-596X(91)90025-F

Jayawickreme, E., Zachry, C. E., & Fleeson, W. (2019). Whole trait theory: An integrative approach to examining personality structure and process. *Personality and Individual Differences*, *136*, 2–11. https://doi.org/10.1016/j.paid.2018.06.045

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*(1), 3–28. https://doi.org/10.1037/0033-2909.114.1.3

Jones, E. E. (1979). The rocky road from acts to dispositions. *American Psychologist*, *34*(2), 107–117. https://doi.org/10.1037/0003-066X.34.2.107

Jones, E. E. (1990). *Interpersonal perception*. W.H. Freeman.

Jones, E. E., & Davis, K. (1965). From acts to dispositions: The attribution process in social psychology. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 219–266). Academic Press. https://doi.org/10.1016/S0065-2601(08)60107-0

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, *3*(1), 1–24. https://doi.org/10.1016/0022-1031(67)90034-0

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of stylegan. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.

Kelley, H. H. (1967). Attribution theory in social psychology. *Nebraska symposium on motivation*, *15*, 192–238.

Kelley, H. H. (1971). *Attribution in social interaction*. General Learning Press.

Kelley, H. H. (1972). *Causal schemata and the attribution process*. General Learning Press.

Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, *28*(2), 107–128. https://doi.org/10.1037/h0034225

Kelley, H. H., & Michela, J. L. (1980). Attribution theory and research. *Annual Review of Psychology*, *31*, 457–501. https://doi.org/10.1146/annurev.ps.31.020180.002325

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge University Press.

Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., . . . Nosek, B. A. (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, *1*(4), 443–490. https://doi.org/10.1177/2515245918810225

Korman, J., & Malle, B. F. (2016). Grasping for traits or reasons? How people grapple with puzzling social behaviors. *Personality and Social Psychology Bulletin*, *42*(11), 1451–1465. https://doi.org/10.1177/0146167216663704

Kressel, L. M., & Uleman, J. S. (2010). Personality traits function as causal concepts. *Journal of Experimental Social Psychology*, *46*(1), 213–216. https://doi.org/10.1016/j.jesp.2009.08.018

Krull, D. S. (1993). Does the grist change the mill? The effect of the perceiver's inferential goal on the process of social inference. *Personality and Social Psychology Bulletin*, *19*(3), 340–348.

Krull, D. S., & Erickson, D. J. (1995). Inferential hopscotch: How people draw social inferences from behavior. *Current Directions in Psychological Science*, *4*(2), 35–38. https://doi.org/10.1111/1467-8721.ep10770986

Kruse, F., & Degner, J. (2021). Spontaneous state inferences. *Journal of Personality and Social Psychology*, *121*(4), 771–791. https://doi.org/10.1037/pspa0000232

Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, *44*(7), 701–710. https://doi.org/10.1002/ejsp.2023

Lalljee, M., Lamb, R., Furnham, A., & Jaspars, J. (1984). Explanations and information search: Inductive and hypothesis-testing approaches to arriving at an explanation. *British Journal of Social Psychology*, *23*(3), 201–212. https://doi.org/10.1111/j.2044-8309.1984.tb00631.x

Levordashka, A., & Utz, S. (2017). Spontaneous trait inferences on social media. *Social Psychological and Personality Science*, *8*(1), 93–101. https://doi.org/10.1177/1948550616663803

Li, T., Liu, X., Pan, J., & Zhou, G. (2017). The interactive effect of facial appearance and behavior statement on trust belief and trust behavior. *Personality and Individual Differences*, *117*, 60–65. https://doi.org/10.1016/j.paid.2017.05.038

Lillard, A. S., & Skibbe, L. (2006). Theory of mind: Conscious attribution and spontaneous trait inference. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 277–305). Oxford University Press. https://doi.org/10.1093/acprof

Lupfer, M. B., Clark, L. F., & Hutcherson, H. W. (1990). Impact of context on spontaneous trait and situational attributions. *Journal of Personality and Social Psychology*, *58*(2), 239–249. https://doi.org/10.1037/0022-3514.58.2.239

Maass, A., Colombo, A., Colombo, A., & Sherman, S. J. (2001). Inferring traits from behaviors versus behaviors from traits: The induction-deduction asymmetry. *Journal of Personality and Social Psychology*, *81*(3), 391–404. https://doi.org/10.1037/0022-3514.81.3.391

Maass, A., Salvi, D., Arcuri, L., & Semin, G. (1989). Language use in intergroup contexts: the linguistic intergroup bias. *Journal of Personality and Social Psychology*, *57*(6), 981–993. https://doi.org/10.1037/0022-3514.57.6.981

Malle, B. F. (2004). *How the mind explains behavior*. MIT Press.

Malle, B. F. (2005). Folk theory of mind: Conceptual foundations of human social cognition. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 225–255). Oxford University Press.

Malle, B. F. (2008). Fritz Heider's legacy: Celebrated insights, many of them misunderstood. *Social Psychology*, *39*(3), 163–173. https://doi.org/10.1027/1864-9335.39.3.163

Malle, B. F., & Holbrook, J. (2012). Is there a hierarchy of social inferences? The likelihood and speed of inferring intentionality, mind, and personality. *Journal of Personality and Social Psychology*, *102*(4), 661–684. https://doi.org/10.1037/a0026790

Malle, B. F., Knobe, J. M., & Nelson, S. E. (2007). Actor-observer asymmetries in explanations of behavior: New answers to an old question. *Journal of Personality and Social Psychology*, *93*(4), 491–514. https://doi.org/10.1037/0022-3514.93.4.491

Malle, B. F., Moses, L. J., & Baldwin, D. A. (Eds.). (2001). *Intentions and intentionality: Foundations of social cognition*. MIT Press.

Marcelo, D., Garcia-Marques, L., & Duarte, I. (2019). Language as a 'game changer' for spontaneous trait inference. *Cognitive Linguistic Studies*, *6*, 185–209. https://doi.org/10.1075/cogls.00035.mar

Maro, P. V. (29 BC). *Georgicon*.

McArthur, L. A. (1972). The how and what of why: Some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology*, *22*(2), 171–193. https://doi.org/10.1037/h0032602

McCarthy, R. J., & Skowronski, J. J. (2011a). The interplay of controlled and automatic processing in the expression of spontaneously inferred traits: A PDP analysis. *Journal of Personality and Social Psychology*, *100*(2), 229–240. https://doi.org/10.1037/a0021991

McCarthy, R. J., & Skowronski, J. J. (2011b). What will Phil do next? Spontaneously inferred traits influence predictions of behavior. *Journal of Experimental Social Psychology*, *47*(2), 321–332. https://doi.org/10.1016/j.jesp.2010.10.015

McClung, J. S., & Reicher, S. D. (2018). Representing other minds: Mental state reference is moderated by group membership. *Journal of Experimental Social Psychology*, *76*, 385–392. https://doi.org/10.1016/j.jesp.2017.11.013

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(1), 82–91. https://doi.org/10.1037/0278-7393.12.1.82

Mischel, W. (1968). *Personality and assessment*. John Wiley & Sons, Inc.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, *132*(2), 297–326. https://doi.org/10.1037/0033-2909.132.2.297

Moskowitz, G. B. (1993). Individual differences in social categorization: The influence of personal need for structure on spontaneous trait inferences. *Journal of Personality and Social Psychology*, *65*(1), 132–142. https://doi.org/10.1111/j.1471-6402.2009.01539.x

Moskowitz, G. B. (2005). Correspondence bias and spontaneous trait inferences. In G. B. Moskowitz (Ed.), *Social cognition: Understanding self and others* (pp. 267–309). Guilford Press.

Moskowitz, G. B., & Olcaysoy Okten, I. (2016). Spontaneous goal inference (SGI). *Social and Personality Psychology Compass*, *10*(1), 64–80. https://doi.org/10.1111/spc3.12232

Myers, D., & Twenge, J. (2018). *Social psychology*. McGraw-Hill Education.

Na, J., & Kitayama, S. (2011). Spontaneous trait inference is culture-specific: Behavioral and neural evidence. *Psychological Science*, *22*(8), 1025–1032. https://doi.org/10.1177/0956797611414727

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9*(3), 353–383. https://doi.org/10.1016/0010-0285(77)90012-3

Neuberg, S. L., & Newsom, J. T. (1993). Personal need for structure: Individual differences in the desire for simple structure. *Journal of Personality and Social Psychology*, *65*, 113–131.

Newman, L. S. (1991). Why are traits inferred spontaneously? A developmental approach. *Social Cognition*, *9*(3), 221–253. https://doi.org/10.1521/soco.1991.9.3.221

Newman, L. S. (1996). Trait impressions as heuristics for predicting future behavior. *Personality and Social Psychology Bulletin*, *22*(4), 395–411. https://doi.org/10.1177/0146167296224006

Nussbaum, S., Trope, Y., & Liberman, N. (2003). Creeping dispositionism: The temporal dynamics of behavior prediction. *Journal of Personality and Social Psychology*, *84*(3), 485–497. https://doi.org/10.1037/0022-3514.84.3.485

Olcaysoy Okten, I., & Moskowitz, G. B. (2018). Goal versus trait explanations: Causal attributions beyond the trait-situation dichotomy. *Journal of Personality and Social Psychology*, *114*(2), 53608. https://doi.org/10.1037/pspa0000104

Olcaysoy Okten, I., & Moskowitz, G. B. (2019). Spontaneous goal versus spontaneous trait inferences and explanations: How ideology shapes attributions. *European Journal of Social Psychology*, *117*(1), 1–12. https://doi.org/10.1002/ejsp.2611

Olcaysoy Okten, I., Schneid, E., & Moskowitz, G. B. (2019). On the updating of spontaneous impressions. *Journal of Personality and Social Psychology*, *117*, 1–25. https://doi.org/10.1037/pspa0000156.supp

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Science*, *105*, 11078–11092.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872. https://doi.org/10.1016/j.jesp.2009.03.009

Orghian, D., Garcia-Marques, L., Uleman, J. S., & Heinke, D. (2015). A connectionist model of spontaneous trait inference and spontaneous trait transference: Do they have the same underlying processes? *Social Cognition*, *33*(1), 20–66. https://doi.org/10.1521/soco.2015.33.1.20

Orghian, D., Ramos, T., & Garcia-Marques, L. (2018). You are cruel even if he did It: Behavior and face processing in spontaneous trait inference and transference. *Basic and Applied Social Psychology*, *40*(2), 104–114. https://doi.org/10.1080/01973533.2018.1436056

Orghian, D., Ramos, T., Reis, J., & Garcia-Marques, L. (2019). Activation is not always inference: Word-based priming in spontaneous trait inferences. *Social Cognition*, *37*(2), 145–173. https://doi.org/10.2139/ssrn.2919048

Pahl, R. (2018). *Groupseq: A gui-based program to compute probabilities regarding group sequential designs* [R package version 1.3.5].

Payne, B. K. (2005). Conceptualizing control in social cognition: How executive functioning modulates the expression of automatic stereotyping. *Journal of Personality and Social Psychology*, *89*(4), 488–503. https://doi.org/10.1037/0022-3514.89.4.488

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Perugini, M., Gallucci, M., & Costantini, G. (2014). Safeguard power as a protection against imprecise power estimates. *Perspectives on Psychological Science*, *9*(3), 319–332. https://doi.org/10.1177/1745691614528519

Pollack, I. (1970). A nonparametric procedure for evaluation of true and false positives. *Behavior Research Methods & Instrumentation*, *2*(4), 155–156. https://doi.org/10.3758/BF03209289

R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.

Ramos, T., Garcia-Marques, L., Hamilton, D. L., Ferreira, M., & Van Acker, K. (2012). What I infer depends on who you are: The influence of stereotypes on trait and situational spontaneous inferences. *Journal of Experimental Social Psychology*, *48*(6), 1247–1256. https://doi.org/10.1016/j.jesp.2012.05.009

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, *114*(3), 510–532. https://doi.org/10.1037/0033-2909.114.3.510

Read, S. J., & Miller, L. C. (2005). Explanatory coherence and goal-based knowledge structures in making dispositional inferences. In B. F. Malle & S. D. Hodges (Eds.), *Other minds: How humans bridge the divide between self and others* (pp. 124–139). Guilford Press.

Reeder, G. D. (2009a). Mindreading and dispositional inference: MIM revised and extended. *Psychological Inquiry*, *20*(1), 73–83. https://doi.org/10.1080/10478400902868565

Reeder, G. D. (2009b). Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychological Inquiry*, *20*(1), 1–18. https://doi.org/10.1080/10478400802615744

Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review*, *86*(1), 61–79. https://doi.org/10.1037/0033-295X.86.1.61

Reeder, G. D., Hesson-McInnis, M., Krohse, J. O., & Scialabba, E. A. (2001). Inferences about effort and ability. *Personality and Social Psychology Bulletin*, *27*(9), 1225–1235. https://doi.org/10.1177/0146167201279014

Rim, S. Y., Min, K. E., Uleman, J. S., Chartrand, T. L., & Carlston, D. E. (2013). Seeing others through rose-colored glasses: An affiliation goal and positivity bias in implicit trait impressions. *Journal of Experimental Social Psychology*, *49*, 1204–1209. https://doi.org/10.1016/j.jesp.2013.05.007

Rim, S., Uleman, J. S., & Trope, Y. (2009). Spontaneous trait inference and construal level theory: Psychological distance increases nonconscious trait thinking. *Journal of Experimental Social Psychology*, *45*, 1088–1097. https://doi.org/10.1016/j.jesp.2009.06.015

Ross, L. D. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 173–220). Academic Press. https://doi.org/10.1016/S0065-2601(08)60357-3

Ross, L. D., Amabile, T. M., & Steinmetz, J. L. (1977). Social roles, social control, and biases in social-perception processes. *Journal of Personality and Social Psychology*, *35*(7), 485–494. https://doi.org/10.1037/0022-3514.35.7.485

Ruffman, T. (2014). To belief or not belief: Children's theory of mind. *Developmental Review*, *34*(3), 265–293. https://doi.org/10.1016/j.dr.2014.04.001

Schneid, E. D., Carlston, D. E., & Skowronski, J. J. (2015). Spontaneous evaluative inferences and their relationship to spontaneous trait inferences. *Journal of Personality and Social Psychology*, *108*(5), 681–696. https://doi.org/10.1037/a0039118

Schneid, E. D., Crawford, M. T., Skowronski, J. J., Irwin, L. M., & Carlston, D. E. (2015). Thinking about other people: Spontaneous trait inferences and spontaneous evaluations. *Social Psychology*, *46*(1), 24–35. https://doi.org/10.1027/1864-9335/a000218

Scott, R. M., & Baillargeon, R. (2017). Early false-belief understanding. *Trends in Cognitive Sciences*, *21*(4), 237–249. https://doi.org/10.1016/j.tics.2017.01.012

Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, *54*(4), 558–568. https://doi.org/10.1037/0022-3514.54.4.558

Shen, X., Mann, T. C., & Ferguson, M. J. (2020). Beware a dishonest face?: Updating face-based implicit impressions using diagnostic behavioral information. *Journal of Experimental Social Psychology*, *86*. https://doi.org/10.1016/j.jesp.2019.103888

Shi, W., Yang, F., Xia, J., & Wang, M. (2019). Differentiating the effects of anger and sadness: A perspective of spontaneous trait inference. *International Journal of Psychology*, *54*, 487–494. https://doi.org/10.1002/ijop.12491

Shimizu, Y. (2017). Why are negative behaviours likely to be immediately invoked traits? The effects of valence and frequency on spontaneous trait inferences. *Asian Journal of Social Psychology*, *20*(3-4), 201–210. https://doi.org/10.1111/ajsp.12183

Skowronski, J. J., Carlston, D. E., Mae, L., & Crawford, M. T. (1998). Spontaneous trait transference: Communicators take on the qualities they describe in others. *Journal of Personality and Social Psychology*, *74*(4), 837–848. https://doi.org/10.1037/0022-3514.74.4.837

Smith, E., Mackie, D., & Claypool, H. (2014). *Social psychology: Fourth edition*. Psychology Press.

Stewart, T. L., Weeks, M., & Lupfer, M. B. (2004). Spontaneous stereotyping: A matter of prejudice? *Social Cognition*, *21*(4), 263–298. https://doi.org/10.1521/soco.21.4.263.27003

Taylor, S. E. (1981). The interface of cognitive and social psychology. In J. Harvey (Ed.), *Cognition, social behavior, and the environment* (pp. 189–211). Lawrence Erlbaum.

Taylor, S. E., & Fiske, S. T. (1975). Point of view and perceptions of causality. *Journal of Personality and Social Psychology*, *32*(3), 439–445. https://doi.org/10.1037/h0077095

Thornton, M. A., & Tamir, D. I. (2020). People represent mental states in terms of rationality, social impact, and valence: Validating the 3d Mind Model. *Cortex*, *125*, 44–59. https://doi.org/10.1016/j.cortex.2019.12.012

Todd, A. R., Molden, D. C., Ham, J., & Vonk, R. (2011). The automatic and co-occurring activation of multiple social inferences. *Journal of Experimental Social Psychology*, *47*(1), 37–49. https://doi.org/10.1016/j.jesp.2010.08.006

Todorov, A., & Uleman, J. S. (2002). Spontaneous trait inferences are bound to actors' faces: Evidence from a false recognition paradigm. *Journal of Personality and Social Psychology*, *83*(5), 1051–1065. https://doi.org/10.1037/0022-3514.83.5.1051

Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology*, *39*(6), 549–562. https://doi.org/10.1016/S0022-1031(03)00059-3

Todorov, A., & Uleman, J. S. (2004). The person reference process in spontaneous trait inferences. *Journal of Personality and Social Psychology*, *87*(4), 482–493. https://doi.org/10.1037/0022-3514.87.4.482

Triandis, H. (2018). *Individualism and collectivism*. Taylor & Francis.

Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, *93*(3), 239–257. https://doi.org/10.1037/0033-295X.93.3.239

Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology*, *79*(3), 344–354. https://doi.org/10.1037//0022-3514.79.3.344

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, *117*(2), 440–463. https://doi.org/10.1037/a0018963

Turner, M. L., & Engle, R. W. (1989). Is working memory capacity task dependent? *Journal of Memory and Language*, *28*(2), 127–154. https://doi.org/10.1016/0749-596X(89)90040-5

Uleman, J. S. (2005). Introduction: Becoming aware of the new unconscious. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 3–18). Oxford University Press.

Uleman, J. S., Blader, S. L., & Todorov, A. (2005). Implicit impressions. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 362–392). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195307696.003.0015

Uleman, J. S., Hon, A., Roman, R. J., & Moskowitz, G. B. (1996). On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin*, *22*(4), 377–394. https://doi.org/10.1177/0146167296224005

Uleman, J. S., & Moskowitz, G. B. (1994). Unintended effects of goals on unintended inferences. *Journal of Personality and Social Psychology*, *66*(3), 490–501. https://doi.org/10.1037/0022-3514.66.3.490

Uleman, J. S., Newman, L. S., & Moskowitz, G. B. (1996). People as flexible interpreters: Evidence and issues from spontaneous trait inference. *Advances in Experimental Social Psychology*, *28*(100), 211–279. https://doi.org/10.1016/S0065-2601(08)60239-7

Uleman, J. S., Newman, L. S., & Winter, L. (1992). Can personality traits be inferred automatically? Spontaneous inferences require cognitive capacity at encoding. *Consciousness and Cognition*, *90*(1), 77–90. https://doi.org/10.1016/1053-8100(92)90049-G

Uleman, J. S., Rim, S., Saribay, A. S., & Kressel, L. M. (2012). Controversies, questions, and prospects for spontaneous social inferences. *Social and Personality Psychology Compass*, *6*(9), 657–673. https://doi.org/10.1111/j.1751-9004.2012.00452.x

Uleman, J. S., Saribay, S. A., & Gonzalez, C. M. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology*, *59*, 329–60. https://doi.org/10.1146/annurev.psych.59.103006.093707

Unkelbach, C., & Förderer, S. (2018). A model of attribute conditioning. *Social Psychological Bulletin*, *13*(3), 1–23. https://doi.org/10.5964/spb.v13i3.28568

Unsworth, N., & Engle, R. W. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, *114*(1), 104–132. https://doi.org/10.1037/0033-295X.114.1.104

Van Overwalle, F., Van Duynslaeger, M., Coomans, D., & Timmermans, B. (2012). Spontaneous goal inferences are often inferred faster than spontaneous trait inferences. *Journal of Experimental Social Psychology*, *48*(1), 13–18. https://doi.org/10.1016/j.jesp.2011.06.016

Vaterrodt-Plünnecke, B., Krüger, T., & Bredenkamp, J. (2002). Process-dissociation procedure: A testable model for considering assumptions about the stochastic relation between consciously controlled and automatic processes. *Experimental Psychology*, *49*(1), 3–26. https://doi.org/10.1027/1618-3169.49.1.3

Wang, M., Xia, J., & Yang, F. (2015). Flexibility of spontaneous trait inferences: The interactive effects of mood and gender stereotypes. *Social Cognition*, *33*, 345–358.

Wells, B. M., Skowronski, J. J., Crawford, M. T., Scherer, C. R., & Carlston, D. E. (2011). Inference making and linking both require thinking: Spontaneous trait inference and spontaneous trait transference both rely on working memory capacity. *Journal of Experimental Social Psychology*, *47*(6), 1116–1126. https://doi.org/10.1016/j.jesp.2011.05.013

Wentura, D., Rohr, M., & Degner, J. (2017). Masked emotional priming: A double dissociation between direct and indirect effects reveals non-conscious processing of emotional information beyond valence. *Consciousness and Cognition*, *49*, 203–214. https://doi.org/10.1016/j.concog.2017.01.016

Whitney, P., Waring, D. A., & Zingmark, B. (1992). Task effects on the spontaneous activation of trait concepts. *Social Cognition*, *10*(4), 377–396. https://doi.org/10.1521/soco.1992.10.4.377

Wigboldus, D. H., Sherman, J. W., Franzese, H. L., & Knippenberg, A. V. (2004). Capacity and comprehension: Spontaneous stereotyping under cognitive load. *Social Cognition*, *22*, 292–309. https://doi.org/10.1521/soco.22.3.292.35967

Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of inferences. *Journal of Personality and Social Psychology*, *47*(2), 237–252. https://doi.org/10.1037/0022-3514.47.2.237

Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology*, *49*(4), 904–917. https://doi.org/10.1037/0022-3514.49.4.904

Zebrowitz, L. A. (2017). First impressions from faces. *Current Directions in Psychological Science*, *26*(3), 237–242. https://doi.org/10.1177/0963721416683996

Appendix A

A

A.1 Pilot Study: Perceived Variability of Person-Describing Adjectives

A.1.1 Participants

Analyses rely on valid data from a total of $N = 55$ participants (41 female; average age $M = 2.8$ years, $SD = 6.7$, ranging from 18 to 55). Participants were recruited on campus and received partial course credit. We excluded the data of six participants because they self-reported insufficient language proficiency.

A.1.2 Materials

We compiled a list of 323 person-describing adjectives.

A.1.3 Procedure

Data collection was conducted online using the platform Qualtrics for online data collection (www.qualtrics.com). Participants were invited to an experiment on person description. Participants were shortly briefed on the concepts of traits and states. The exact formulation translates to: "In psychology, we distinguish between traits and states. Traits are characteristics of a person, that are stable over time and responsible for a person's general behavior. States are unstable, both temporally and situationally, so they are a temporary state of a person, that are responsible for a person's behavior in a specific situation. Many adjectives can be clearly classified as a trait or a state, others are ambiguous". The attention check described in Experiment 1 was used in this experiment as well. Participants were presented with a subset of 48 of the person-describing adjectives presented on one by one in individual random order and were asked to rate the variability of the adjectives on a slider with a continuous scale ranging from -100 (trait, stable) to +100 (state, variable). Finally, participants provided demographic information (age, sex, language proficiency, profession, and education).

A.1.4 Results

Results for each of the 323 adjectives are reported in Table A.1.

Tab. A.1.: Pilot Study: Variability ratings

Adjective	Translation	*M*	*SD*	*n*	Used in Experiments
introvertiert	introverted	-86.747	13.17	13	
sozial	social	-83.935	17.853	25	3 & 4
schlau	smart	-81.951	24.186	17	4
ehrgeizig	ambitious	-81.378	18.352	19	
intelligent	intelligent	-80.916	23.481	22	
klug	smart	-80.389	24.038	20	
autoritär	authoritarian	-78.271	23.954	16	
selbstbewusst	self-conscious	-77.486	24.897	21	
verantwortungsvoll	responsible	-75.178	22.016	25	
weise	clever	-74.975	27.257	17	
gewissenhaft	conscientious	-74.585	24.109	22	
traditionell	traditional	-74.118	26.931	17	
pflichtbewusst	dutiful	-73.44	21.05	23	
vertrauenswürdig	trustworthy	-72.912	32.848	22	
kreativ	creative	-72.314	30.881	20	1, 2, 3, & 4
musikalisch	musical	-72.079	39.982	20	
durchsetzungsfähig	assertive	-71.834	31.164	19	
perfektionistisch	perfectionistic	-71	35.648	20	
gläubig	religious	-70.972	45.922	15	
dickköpfig	stubborn	-70.369	22.754	13	
wissbegierig	inquisitive	-69.96	27.185	14	
konventionell	conventional	-69.422	40.319	17	
ehrlich	honest	-68.952	23.247	14	
zielstrebig	determined	-68.9	25.895	16	
clever	clever	-68.788	34.473	18	4
konservativ	conservative	-68.713	41.796	25	
tolerant	tolerant	-68.405	30.507	17	
impulsiv	impulsive	-68.194	23.448	17	4
geizig	stingy	-67.498	26.266	13	3
altmodisch	old-fashioned	-66.595	32.1	15	
verantwortungsbewusst	responsible	-65.66	33.695	17	
gütig	kind	-65.435	27.32	14	
konsequent	consistent	-64.967	36.986	18	
einfallsreich	imaginative	-64.663	30.53	22	
extravertiert	extraverted	-64.253	44.428	17	
organisiert	organized	-64.157	27.706	20	
großzügig	generous	-64.096	36.131	16	1 & 2
einfühlsam	compassionate	-64.082	44.112	16	
hartherzig	stone-hearted	-63.849	33.103	14	
dominant	dominant	-63.804	21.743	15	
weitsichtig	farsighted	-63.507	27.986	16	
unzuverlässig	unreliable	-63.503	31.673	19	
ordentlich	tidy	-63.434	45.622	21	
loyal	loyal	-63.086	41.549	17	
phantasievoll	visionary	-61.676	31.555	14	
exzentrisch	excentric	-60.937	30.383	17	
empathisch	empathetic	-60.84	52.816	20	
höflich	polite	-60.599	40.002	14	
umsichtig	careful	-59.745	39.294	16	
charakterschwach	weak	-59.291	46.345	16	
streng	strict	-58.935	28.587	17	1
sorgfältig	thorough	-58.213	37.603	15	
gesellig	sociable	-58.089	41.681	19	
emotional	emotional	-57.078	30.15	18	3 & 4
strebsam	ambitious	-55.92	40.162	20	
naiv	naive	-55.863	29.331	15	1 & 2
mitfühlend	compassionate	-55.38	48.928	18	
eingebildet	conceited	-55.158	38.177	19	1 & 2
nachsichtig	indulgent	-55.056	35.641	16	
tollpatschig	clumsy	-54.926	45.197	17	
cholerisch	choleric	-54.88	42.095	20	
risikofreudig	adventurous	-54.667	62.716	20	
tiefsinnig	profound	-54.193	43.755	16	
abenteuerlustig	adventurous	-54.086	49.829	20	3
selbstlos	selfless	-54.003	47.152	20	
arrogant	arrogant	-53.825	35.135	22	
optimistisch	optimistic	-53.697	46.795	19	
erfinderisch	innovative	-52.731	36.701	18	
hilfsbereit	cooperative	-52.551	58.976	18	1
chaotisch	chaotic	-52.471	43.468	18	3
flexibel	flexible	-52.453	40.629	20	
unorganisiert	unorganised	-52.422	33.109	21	3
verklemmt	uptight	-52.191	35.692	20	
neugierig	curious	-51.679	50.557	22	1, 2, 3, & 4
zielorientiert	goal-oriented	-51.239	50.881	21	
bestimmend	assertive	-51.155	37.07	19	
spirituell	spiritual	-51.033	45.256	21	

(continued)

Tab. A.1.: Pilot Study: Results (continued)

Adjective	Translation	*M*	*SD*	*n*	Used in Experiments
dumm	stupid	-50.869	44.014	19	
schüchtern	shy	-50.839	48.976	21	1, 2, 3, & 4
zurückhaltend	reluctant	-50.724	41.624	15	
gründlich	thorough	-50.346	44.736	26	
streitsüchtig	quarrelsome	-50.166	44.752	20	
gefühlsarm	unemotional	-49.229	37.387	16	
verschlossen	secretive	-48.465	43.64	14	
verständnisvoll	understanding	-48.419	49.222	21	
empfindlich	sensitive	-48.375	38.964	17	
prüde	prudish	-48.206	35.81	16	
berechnend	designing	-47.91	49.655	21	
cool	cool	-47.785	44.616	15	
ausgeglichen	balanced	-47.335	63.117	22	
geldgeil	greedy	-46.962	54.613	18	
kooperativ	cooperative	-46.334	44.31	22	
schreckhaft	jumpy	-46.289	53.062	20	
anspruchslos	undemanding	-46.201	46.695	22	
verantwortungslos	irresponsible	-45.952	52.664	17	4
kritisch	critical	-45.667	47.696	11	
spontan	spontaneous	-45.446	58.284	18	
sportlich	athletic	-44.138	64.575	13	1, 2, & 3
bedacht	thoughtful	-43.764	48.359	19	
fleißig	hard-working	-43.498	52.856	19	3
ignorant	ignorant	-42.721	44.567	15	
phantasielos	unimaginative	-42.047	41.127	22	
tüchtig	capable	-41.919	48.119	21	
wertschätzend	appreciative	-41.66	61.85	17	
still	calm	-41.379	43.497	14	
pedantisch	nit-picky	-41.27	47.398	16	
kränkbar	sensitive	-40.622	55.545	19	
besonnen	prudent	-40.566	39.801	14	
unehrlich	dishonest	-40.484	52.744	19	1
mutig	brave	-40.14	50.075	22	
unsportlich	unfit	-39.876	57.912	14	
angeberisch	cocky	-39.603	51.229	23	1, 2, & 4
verwöhnt	spoiled	-38.696	59.414	20	
gesundheitsbewusst	health conscious	-38.484	48.267	14	
verpeilt	sloppy	-38.057	39.818	16	
gesprächig	talkative	-37.841	52.062	23	
aktiv	active	-37.835	44.774	19	
vergesslich	forgetful	-37.766	52.771	18	
leidenschaftlich	passionate	-37.557	60.167	21	
willensschwach	weak-minded	-37.245	44.578	17	
vorsichtig	cautious	-37.196	45.542	17	1 & 2
kompromissbereit	dovish	-37.087	54.223	21	
misstrauisch	suspicious	-36.981	46.56	21	
gierig	greedy	-35.691	35.826	17	
unselbstständig	dependant	-35.281	47.838	21	
gelassen	calm	-34.366	49.99	21	
stark	strong	-33.864	64.887	13	
sentimental	sentimental	-33.848	64.438	20	
wehleidig	sniveling	-33.767	54.637	14	
leichtsinnig	careless	-33.758	43.273	26	3
genießerisch	appreciative	-33.506	45.918	20	
sicher	sure	-33.155	51.794	22	
nachtragend	resentful	-31.759	57.271	20	
dynamisch	dynamic	-31.639	50.768	17	
verantwortlich	responsible	-30.424	71.283	19	
alt	old	-30.213	76.135	20	
rücksichtslos	inconsiderate	-28.726	48.556	16	
voreingenommen	biased	-27.174	55.703	23	
ruhig	calm	-27.121	55.131	21	
unhöflich	impolite	-26.713	57.718	19	3 & 4
gemein	mean	-26.234	53.699	25	4
faul	lazy	-25.517	51.066	20	4
zwanghaft	compulsive	-25.213	48.978	20	
vorausschauend	foresighted	-24.476	66.238	8	
beeindruckbar	impressionable	-23.174	61.806	20	
grausam	cruel	-22.598	54.239	18	
korrupt	corrupt	-22.438	65.707	21	
gefräßig	voracious	-21.734	61.38	19	3 & 4
lustig	funny	-20.542	50.306	6	3
pünktlich	punctual	-19.925	58.785	17	
inkompetent	incompetent	-19.427	62.865	19	
aufdringlich	obtrusive	-19.124	48.486	16	
irritierbar	irritable	-18.168	58.386	16	
grüblerisch	contemplative	-17.265	50.212	17	3
begierig	eager	-17.064	55.122	16	

(continued)

Tab. A.1.: Pilot Study: Results (continued)

Adjective	Translation	*M*	*SD*	*n*	Used in Experiments
routiniert	experienced	-16.944	63.419	20	
kalt	cold	-16.819	59.611	12	
zuversichtlich	confident	-15.972	53.102	14	
energisch	vigorous	-15.676	55.754	17	
verletzlich	offendable	-14.725	59.497	18	
ungeduldig	impatient	-13.326	58.966	14	
aufbrausend	quick-tempered	-11.205	60.239	25	
gemütlich	cozy	-10.137	59.656	17	
schludrig	sloppy	-8.27	56.971	16	
unsicher	insecure	-8.131	40.619	15	
ängstlich	anxious	-7.913	45.583	12	3 & 4
gleichgültig	indifferent	-7.466	61.074	22	4
unbesorgt	unconcerned	-6.25	49.092	17	
zickig	bitchy	-4.946	74.274	14	
wachsam	vigilant	-4.319	54.416	16	
schwach	weak	-2.646	51.524	14	
skeptisch	sceptical	-1.898	53.901	20	
launisch	moody	-1.753	78.089	19	
kriminell	criminal	-0.758	59.294	17	
eifersüchtig	jealous	-0.533	53.757	23	
fröhlich	cheerful	-0.456	57.041	19	
eifrig	eager	0.456	70.206	13	3
nachlässig	clumsy	1.109	59.512	20	
aufmerksam	alert	1.313	66.602	19	
schnell	fast	1.629	60.668	16	
ungerecht	unfair	2.779	46.593	15	
entschlossen	dedicated	2.938	65.027	12	
strafend	punitive	4.995	63.671	22	
zweifelnd	doubtful	5.862	50.862	21	
aggressiv	agressive	6.23	56.24	18	
bissig	snappy	7.261	55.554	20	
reizbar	irritable	8.647	61.69	17	
ablenkbar	distractable	9.108	57.638	18	
gesund	healthy	9.359	63.086	14	
interessiert	interested	9.794	64.233	20	3 & 4
heiß	hot	10.134	64.045	18	
hoffnungsvoll	hopeful	10.341	60.527	23	
distanziert	distanced	11.935	60.464	22	
neidisch	envious	12.436	59.669	18	
nörgelig	cranky	13.924	62.563	14	
unentschlossen	indecisive	15.73	56.586	20	
melancholisch	melancolic	17.539	58.336	14	
unentschieden	undecided	17.654	65.853	16	
unbedacht	rash	19.001	59.307	19	
rastlos	restless	19.209	62.499	14	
motiviert	motivated	19.556	50.783	20	
süchtig	addicted	19.679	54.165	18	
depressiv	depressed	20.019	49.951	21	4
genießend	enjoying	20.558	57.841	22	
einsam	lonely	20.626	67.86	17	
informiert	informed	22.282	54.153	17	
energiegeladen	energetic	22.567	56.589	16	
zufrieden	contented	23.635	58.783	17	4
miserabel	miserable	25.956	70.973	19	
bewundernd	admiring	26.111	61.483	26	
schadenfroh	spiteful	27.056	55.984	18	
unruhig	restless	27.457	50.706	12	
bekümmert	sorry	28.936	58.086	19	
heiter	cheerful	29.189	53.061	24	
stolz	proud	29.837	55.535	16	
kränklich	sick	30.133	66.731	21	
besorgt	concerned	30.242	63.966	16	
ekstatisch	ecstatic	31.061	57.306	19	
sehnsüchtig	longing	31.388	57.311	23	
enthusiastisch	enthusiastic	32.992	56.067	17	3
apathisch	apathic	34.544	42.415	17	
desinteressiert	disinterested	34.87	55.667	14	
dankbar	thankful	34.958	63.01	13	
verschüchtert	intimidated	35.332	57.451	15	
bedauernd	regretful	35.576	48.802	19	
munter	blithely	39.34	61.414	22	
erwartungsvoll	expectant	40.886	60.235	19	
ermutigt	encouraged	41.525	59.224	12	
resigniert	resignated	41.636	42.18	20	
freudig	joyful	42.044	48.429	21	
bereuend	repentant	43.212	62.198	19	
entspannt	relaxed	43.721	51.252	15	
unmovitiert	unmotivated	45.158	38.611	19	

(continued)

Tab. A.1.: Pilot Study: Results (continued)

Adjective	Translation	*M*	*SD*	*n*	Used in Experiments
verpennt	dozy	45.326	55.033	22	
fieberhaft	feverish	46.117	42.341	22	
begeistert	thrilled	46.182	59.005	13	
leer	empty	46.813	54.647	14	
kraftlos	powerless	46.967	40.017	14	
uninteressiert	uninterested	50.133	54.52	19	
beunruhigt	worried	50.601	50.446	12	3 & 4
nervös	nervous	50.994	27.972	15	3
panisch	panicky	51.67	59.538	17	
beschwingt	exhilarated	54.166	22.388	19	
hellwach	awake	54.783	45.303	12	
aufgedreht	hyper	54.985	50.284	18	
nüchtern	sober	55.773	51.044	24	
konzentriert	focused	55.85	42.916	15	4
peinlich	embarrassed	56.712	37.29	20	
gehetzt	rushed	56.776	40.522	17	3
traurig	sad	56.998	40.272	14	1, 2, 3, & 4
gereizt	irritated	58.656	35.958	17	
inspiriert	inspired	58.806	49.731	14	1, 2, 3, & 4
vorfreudig	excited	59.331	47.046	13	
unmotiviert	unmotivated	59.579	37.05	18	
angespannt	tense	59.807	47.441	24	
verzaubert	enchanted	59.873	28.613	17	
erstaunt	astonished	60.56	38.255	15	
verzweifelt	desparate	60.694	39.732	14	
niedergeschlagen	downhearted	61.427	37.446	22	
verliebt	in love	61.552	54.237	22	1 & 2
unkonzentriert	unconcentrated	61.708	32.031	12	
euphorisch	euphoric	62.587	32.153	17	
glücklich	happy	62.861	45.006	16	1, 2, & 4
beflügelt	inspired	63.029	45.276	20	
eingeschüchtert	intimidated	63.356	27.403	17	
voll	full	64.328	37.189	18	
wach	awake	64.392	45.191	17	
beschämt	ashamed	64.613	42.638	19	3
ernüchtert	disillusioned	64.652	27.014	12	
gestresst	stressed	64.767	27.171	20	
gekränkt	hurt	65.539	35.366	20	
zornig	furious	65.741	34.841	14	
verwundert	puzzled	66.194	45.287	22	
genervt	annoyed	66.63	26.269	17	1, 2, & 4
unwohl	uneasy	66.976	32.903	16	
schläfrig	sleepy	67.954	44.32	17	
eingeschnappt	cross	68.448	34.772	22	
verärgert	upset	68.563	47.42	16	
verwirrt	puzzled	68.718	30.103	17	
pleite	broke	68.785	33.718	21	3
beleidigt	offended	69.065	25.899	15	
zittrig	shaky	69.292	28.427	23	
abgelenkt	distracted	70.746	42.634	16	4
berührt	touched	70.959	20.504	18	
enttäuscht	disappointed	71.47	28.453	21	
erschöpft	exhausted	71.674	23.782	17	
erfreut	delighted	72.171	21.369	20	1 & 2
krank	sick	72.499	29	16	1 & 2
verdutzt	puzzled	72.503	29.262	18	
betrübt	aggrieved	72.775	21.874	18	
betroffen	affected	73.131	19.203	15	
müde	tired	73.83	25.2	26	3 & 4
amüsiert	amused	73.854	19.832	22	
erkrankt	sick	73.981	23.573	19	
erschrocken	frightened	74.322	22.942	25	
überrascht	surprised	74.446	32.401	21	1, 2, & 4
erholt	relaxed	74.617	23.162	19	1 & 2
wütend	angry	75.531	25.962	18	1, 2, 3, & 4
aufgebracht	agitated	76.005	33.685	20	
aufgeregt	excited	76.12	26.337	18	1, 2, 3, & 4
verführt	seduced	76.169	23.389	17	
bedroht	threatened	76.276	25.235	17	
übernächtigt	bleary-eyed	76.582	30.323	25	
trauernd	mourning	77.558	24.114	14	
gelangweilt	bored	77.616	16.904	14	1 & 2
erleichtert	relieved	78.473	25.171	18	
ausgeruht	rested	78.77	29.548	20	
irritiert	irritated	79.137	36.879	20	
gedemütigt	humiliated	79.506	33.041	20	
verletzt	hurt	80.879	14.701	20	
gerührt	touched	81.722	23.651	20	1 & 2

(continued)

Tab. A.1.: Pilot Study: Results (continued)

Adjective	Translation	*M*	*SD*	*n*	Used in Experiments
gespannt	eager	81.855	17.965	12	
hungrig	hungry	82.66	23.098	14	1, 2, 3, & 4
satt	full	83.127	25.406	20	1, 2, & 3
schockiert	shocked	83.662	18.469	15	
angeekelt	disgusted	84.3	23.655	20	
erregt	agitated	84.488	21.959	15	
angetrunken	drunk	84.532	17.91	16	
verknallt	infatuated	84.845	16.464	11	
high	high	85.874	22.561	17	
betrunken	drunk	87.235	17.79	21	
durstig	thirsty	88.354	24.41	18	1, 2, & 3
verblüfft	bewildered	93.522	7.718	13	

A.2 Experiments 1 and 2: Material Pretest

A.2.1 Participants

Pretest analyses rely on valid data from a total of $N = 106$ participants (25 female; average age $M = 43.4$ years, $SD = 24$, ranging from 21 to 83). Participants were recruited via Figure Eight (www.figure-eight.com) and received monetary compensation of 2.30 USD for the duration of nine minutes. We excluded data from four participants because they did not pass an initial attention test and one further participant because they self-reported insufficient language proficiency.

A.2.2 Materials

We generated 96 single-implication statements that implied, but not explicitly mentioned a trait and 96 statements that implied, but not explicitly mentioned a state. Statements described a behavior and were written in first-person perspective in German language.

A.2.3 Procedure

Participants were invited to an experiment on spontaneous person associations and randomly assigned to one of four subsets of 24 state-implying and 24 trait-implying statements. Participants were asked to write down the first adjective that came to mind upon reading each statement, that could describe the person making the statement. Subsequently, participants provided demographic information.

A.2.4 Results

For each statement, we grouped synonyms and calculated the proportions of each trait or state word as a consensus score. We selected 18 trait-implying statements and 18 state-implying statements based on the following criteria: (a) the implied trait or state reached at least 50% consensus, and (b) the implied trait or state did not semantically overlap with a trait or state implied by another selected statement. Selected statements and their pretest results are displayed in Tables A.2 and A.3, translated statements are listed in Tables A.6 and A.7.

The selected trait-implying and state-implying statements did not differ in agreement rate (M_{trait} = 80%, SD = 10; M_{state} = 78%, SD = 9, $t(34)$ = 0.711, p = .49, d_s = 0.237, 95% CI [-0.442, 0.917], or statement length (M_{trait} = 13.8 words, SD = 4.1; M_{state} = 13.3 words, SD = 4, $t(32.964)$ = 0.389, p = .70, d_s = 0.31, 95% CI [-0.381, 1]). Additionally, we selected nine filler statements that described a behavior and explicitly mentioned a trait or a state.

Tab. A.2.: Experiments 1 and 2: Trait-implying statements (German)

Nr	Statement	Implied Trait	Other Trait (c)	Other Trait (ic)	New Trait (c)	New Trait (ic)	%	*n*
1	heute in der bahn hat sich eine total schöne Frau neben mich gesetzt. Hab sie aber nicht angesprochen…	schüchtern	streng	sportlich	exzentrisch	harmonisch	87	23
2	Ich kann morgen ein paar Termine verschieben und spring' für dich ein, alles gut!	hilfsbereit	sportlich	stur	weise	exzentrisch	81.8	22
3	die bescheinigung müsste ich irgendwo liegen haben, frag mich aber nicht, wo	unordentlich	angeberisch	ehrlich	misstrauisch	leidenschaftlich	63.6	22
4	nein, es ist alles wieder gut. er hat erklärt, der lippenstiftfleck an seinem hemd kommt weil er erste Hilfe bei jemandem geleistet hat.	naiv	stur	hilfsbereit	irritierbar	furchtlos	63.6	22
5	sie gefiel mir, da hab ich 'unauffällig' meinen porsche-schlüssel auf den tisch gelegt	angeberisch	ungeschickt	fleißig	machthungrig	gastfreundlich	80	20
6	War gestern tanzen mit Marisa. Bin ihr die ganze Zeit auf die Füße getreten!	ungeschickt	unordentlich	vorsichtig	fordernd	kooperativ	63.6	22
7	klar krieg ich den rum, wie soll der MIR widerstehen?	eingebildet	penibel	sportlich	resigniert	nachsichtig	78.3	23
8	Meinen Schülern lasse ich während der Stillarbeitsphase keinen einzigem Mucks durchgehen!	streng	eingebildet	neugierig	leichtsinnig	bodenständig	90.9	22
9	auch wenn die straße leer ist, hab ich lieber nochmal geschaut, ob kein rechtsabbieger kommt.	vorsichtig	ehrlich	unordentlich	gastfreundlich	machthungrig	86.4	22
10	hab dem obdachlosen fünf euro zugesteckt	großzügig	vorsichtig	ungeschickt	leidenschaftlich	misstrauisch	73.9	23
11	ich hab für meine weihnachtsgeschenke nur bonusmeilen ausgegeben	sparsam	hilfsbereit	streng	kooperativ	beeinflussbar	78.3	23
12	habe eben zuviel wechselgeld bekommen. Bin dann nochmal zurück und habs der kassiererin zurückgegeben.	ehrlich	fleißig	angeberisch	nachsichtig	fordernd	81.8	22
13	du kannst soviel argumentieren wie du willst, Ich bleibe bei meiner meinung.	stur	schüchtern	kreativ	beeinflussbar	verantwort-ungsbewusst	86.4	22
14	bin extra noch zum anderen blumenladen gefahren, jetzt passen deko, geschirr und tischdecke genau zueinander	penibel	naiv	sparsam	missmutig	weise	68.2	22
15	wir hatten alle bücher schon durch, da habe ich mir einfach selber eine neue geschichte für meine tochter ausgedacht.	kreativ	sparsam	schüchtern	furchtlos	irritierbar	81.8	22
16	ich trainiere jetzt dreimal die woche, dieses jahr schaffe ich den 10k-lauf unter 50 minuten!	sportlich	neugierig	penibel	bodenständig	missmutig	71.4	21
17	ich büffele am wochenende und zur not auch in den ferien	fleißig	kreativ	eingebildet	harmonisch	leichtsinnig	95.2	21
18	Heute lerne ich endlich den neuen Freund meiner Schwester kennen, mal gucken was der so für ein Typ ist!	neugierig	großzügig	naiv	verantwort-ungsbewusst	resigniert	73.9	22

Tab. A.3.: Experiments 1 and 2: State-implying statements (German)

Nr	Statement	Implied State	Other State (c)	Other State (ic)	New State (c)	New State (ic)	%	*n*
1	hab grad ne ganze flasche wasser in einem zug leergetrunken	durstig	genervt	inspiriert	neidisch	geborgen	86.9	23
2	Kann nicht aufhören an sie zu denken	verliebt	erfreut	krank	überzeugt	angeekelt	77.3	22
3	Musste sogar den gratis Bagel ablehnen, krieg einfach nichts mehr rein.	satt	überrascht	genervt	wach	besessen	90.9	22
4	mein magen knurrt, schaue schon die ganzezeit auf die uhr, wann ist denn endlich mittagspause?	hungrig	traurig	überrascht	spöttisch	engagiert	72.7	22
5	Als meine Schwester und ihr Mann die Ringe getauscht haben konnte ich die Tränen einfach nicht mehr aufhalten.	gerührt	erholt	beschämt	informiert	leer	63.6	22
6	das wochenende war echt schön, lag die ganze zeit am see. Jetzt kann ich auch wieder in die Arbeitswoche starten	erholt	aufgeregt	wütend	locker	unartig	76.2	21
7	Ich hab gar nicht damit gerechnet, dass du heute vorbeikommst!	überrascht	verliebt	traurig	sicher	pleite	85.7	20
8	bin heute früh nach hause, mein kopf tat total weh und ich hatte halsweh.	krank	hungrig	gerührt	pleite	sicher	86.4	22
9	Meine Kollegin kam heute mit einem wunderschönen Blumenstrauß in meinem Büro vorbei, ich habe sie gleich umarmt	erfreut	satt	durstig	friedlich	distanziert	77.3	22
10	Dieses Wochenende ist echt nix los, zappe durchs Fernsehen	gelangweilt	gestresst	glücklich	distanziert	überzeugt	94.1	20
11	jetzt trag ich den ganzen kram schon zum vierten mal in dieses blöde formular ein. . . .	genervt	beschämt	aufgeregt	angeekelt	friedlich	72.7	22
12	Die Prüfung ist richtig gut gelaufen. Bin auf dem Nachhauseweg tatsächlich ein bisschen auf und ab gehüpft.	glücklich	gerührt	gelangweilt	gesund	zögerlich	68.2	22
13	laufe schon den ganzen morgen wie ein tiger im käfig durch die wohnung, weil Jonas heute endlich wiederkommt!	aufgeregt	inspiriert	beschämt	geborgen	neidisch	86.4	22
14	Ich hab stundenlang vor dem Gemälde gesessen und die Farben auf mich wirken lassen.	inspiriert	glücklich	gestresst	engagiert	spöttisch	66.7	22
15	hab den typ beim bäcker total angeschnauzt	wütend	durstig	erfreut	zögerlich	wach	71.4	22
16	bin gleich nach hause und habe einfach nur geweint.	traurig	gelangweilt	erholt	unartig	gesund	95.2	22
17	habe keine zeit auf nen kaffee zu bleiben, muss schnell weiter!	gestresst	wütend	verliebt	leer	informiert	95.2	21
18	würde am liebsten im boden versinken, kann nicht glauben dass ich das gesagt habe!	beschämt	krank	inspiriert	besessen	locker	81	21

Tab. A.4.: Experiments 1 and 2: Filler statements (Trait, German)

Nr	Statement	Presented Trait
1	Hab meiner Frau heut ganz romantisch Blumen mitgebracht	romantisch
2	Fühle mich in letzter Zeit schon öfters einsam	einsam
3	Ich bin zumindest gläubig genug, um in die Kirche zu gehen	gläubig
4	was den urlaub angeht, bin ich total optimistisch.	optimistisch
5	Ich bin viel zu ablenkbar, bin nur am Handy!	ablenkbar
6	Hab heute den ganzen Tag faul in der Sonne gelegen.	faul
7	Vorsicht, auf so etwas reagiere ich empfindlich!	empfindlich
8	Ich war schlau und hab mir nen Sitzplatz reserviert. Brechend voll hier!	schlau
9	Ich hab mich nicht aus der Ruhe bringen lassen und bin ganz höflich geblieben.	höflich

Tab. A.5.: Experiments 1 and 2: Filler statements (State, German)

Nr	Statement	Presented State
1	mellie bekommt jetzt ne abmahnung, bin ja schon ein bisschen schadenfroh!	schadenfroh
2	Hab meinen Freund eben total angeschrien, jetzt fühle ich mich schuldig...	schuldig
3	Wie lange geht die Lernphase noch. Ich fühle mich total ausgelaugt!	ausgelaugt
4	Sitz gemütlich in der Sonne, les mein Buch, nichts kann mich stören.	gemütlich
5	Das war ein Bier zuviel gestern, übel verkatert!	verkatert
6	Bin euch sooo dankbar, dass ihr alle beim Umzug geholfen habt!!	dankbar
7	Ich fahre jeden Morgen mit dem Fahrrad zu Arbeit, man muss aktiv sein!	aktiv
8	Die Schlange vor dem Club war mega lang, hab mich einfach dreist dran vorbeigemogelt	dreist
9	Mein Sitznachbar meinte ich wäre unhöflich, nur weil ich mich in der Vorlesung ein bisschen unterhalten hab.	unhöflich

Tab. A.6.: Experiments 1 and 2: Trait-implying statements (translated)

Nr	Statement	Implied Trait	Other Trait (c)	Other Trait (ic)	New Trait (c)	New Trait (ic)	%	*n*
1	A beautiful woman sat next to me today in the metro. I didn't talk to her though...	shy	strict	athletic	eccentric	harmonic	87	23
2	I can move some appointments tomorrow and fill in for you, no problem!	helpful	athletic	stubborn	wise	eccentric	81.8	22
3	The certificate must be somewhere around here, don't ask me where though.	messy	bragging	honest	suspicious	passionate	63.6	22
4	No, it's all fine. He explained that the lipstick on his shirt came from him providing first aid to someone.	naive	stubborn	helpful	irritable	fearless	63.6	22
5	I liked her, so I subtly laid my Porsche keys on the table.	bragging	clumsy	diligent	power-hungry	hospitable	80	20
6	I went dancing with Marisa yesterday. I kept on stepping on her feet!	clumsy	messy	cautious	demanding	cooperative	63.6	22
7	Of course I can win him over, how could he resist ME?	conceited	meticulous	athletic	resignated	lenient	78.3	23
8	I don't give my pupils a pass on a single noise during the quiet working phase!	strict	conceited	curious	careless	grounded	90.9	22
9	Even though the street was empty, I rather checked twice for anybody turning right.	cautious	honest	messy	hospitable	power-hungry	86.4	22
10	I gave the homeless guy five euros.	generous	cautious	clumsy	passionate	suspicious	73.9	23
11	I only spent bonus miles for my christmas presents.	economical	helpful	strict	cooperative	influenceable	78.3	23
12	I've just received too much change, so I went back and gave it back to the cashier.	honest	diligent	bragging	lenient	demanding	81.8	22
13	You can argue as much as you like, I'm not changing my mind.	stubborn	shy	creative	influencable	responsible	86.4	22
14	I even went to another florist, now the decoration, the dishes and the tablecloth are all matching!	meticulous	naive	economical	discontented	wise	68.2	22
15	We already finished all our books, so I just made up my own story for my daughter.	creative	economical	shy	fearless	irritable	81.8	22
16	I'm training three times a week, this year I'll manage the 10k-run in less than 50 minutes!	athletic	curious	meticulous	grounded	discontented	71.4	21
17	I'm studying on the weekends and if necessary even during the holidays.	diligent	creative	conceited	harmonic	careless	95.2	21
18	Today I'll finally get to know the new boyfriend of my sister, let's see what kind of a guy he is.	curious	generous	naive	responsible	resignated	73.9	2

Tab. A.7.: Experiments 1 and 2: State-implying statements (translated)

Nr	Statement	Implied State	Other State (c)	Other State (ic)	New State (c)	New State (ic)	%	*n*
1	Just finished an entire bottle of water in one gulp.	thirsty	annoyed	inspired	jealous	sheltered	86.9	23
2	I just can't stop thinking about her.	in love	pleased	sick	convinced	disgusted	77.3	22
3	I even rejected the free Bagel, I just can't eat any more.	full	surprised	annoyed	woke	obsessed	90.9	22
4	My tummy rumbles and I keep checking the clock. When are we finally having lunch break?	hungry	sad	überrascht	mocking	engaged	72.7	22
5	When my sister and her husband exchanged rings, I just couldn't hold back the tears.	touched	recovered	ashamed	informed	empty	63.6	22
6	Such a lovely weekend, I spent all day by the lake. Ready for a new working week!	recovered	excited	angry	loose	naughty	76.2	21
7	I wasn't expecting you coming over today!	surprised	in love	sad	secure	broke	85.7	20
8	Went home early, my head hurt and I had a sore throat.	sick	hungry	touched	broke	secure	86.4	22
9	My coworker came into my office with a beautiful bouquet of flowers, I hugged her right away.	pleased	full	thirsty	peaceful	distanced	77.3	22
10	There really is nothing going on this weekend. I am just channel surfing on the tv.	bored	stressed	happy	distanced	convinced	94.1	20
11	Now I need to fill out this stupid form for the fourth time.	annoyed	ashamed	excited	disgusted	peaceful	72.7	22
12	The exam went really well. On my way home I even bounced around.	happy	touched	bored	healthy	hesitant	68.2	22
13	Spent all morning running around in my flat like a tiger in a cage, because Jonas is finally coming back today!	excited	inspdire	ashamed	sheltered	jealous	86.4	22
14	I spent hours sitting in front of the painting, letting the colors sink in.	inspired	happy	stressed	engaged	mocking	66.7	22
15	I just ranted at that guy at the bakery.	angry	thirsty	pleased	hesitant	woke	71.4	22
16	I went home right away and cried.	sad	bored	recovered	naughty	healthy	95.2	22
17	I don't have time to stay for a coffee, I have to run.	stressed	angry	in love	empty	informed	95.2	21
18	I'd like to sink into the ground, can't believe that I said this.	ashamed	sick	inspired	obsessed	loose	81	21

Tab. A.8.: Experiments 1 and 2: Filler statements (Trait, translated)

Nr	Statement	Presented Trait
1	Today I was feeling romantic and bought my wife some flowers.	romantic
2	Lately I'm feeling lonely.	lonely
3	At least I'm religious enough to go to church.	devout
4	I am completely optimistic about the vacation.	optimistic
5	I am way too easy distractable, keep checking the phone all the time.	distractable
6	Spent all day laying in the sun lazily.	lazy
7	Be careful, I'm sensitive in these matters!	sensitive
8	I was smart and reserved a seat. It's so full here!	smart
9	I didn't let him bother me and remained polite.	polite

Tab. A.9.: Experiments 1 and 2: Filler statements (State, translated)

Nr	Statement	Presented State
1	Melli is about to get a warning letter, I have to admit, I'm a little gleeful.	gleeful
2	I just shouted at my boyfriend, now I'm feeling guilty. . .	guilty
3	When will this exam period be over? I'm completely exhausted!	exhausted
4	Sitting comfortably in the sun, reading my book. Nothing can bother me.	comfortable
5	That was one beer too many, I'm completely hungover!	hungover
6	I'm sooo thankful to all of you for helping me move!	thankful
7	Everyday I'm riding my bike to work, you need to stay active!	acitve
8	The queue in front of the club was so long, I brazenly passed it.	brazenly
9	The guy sitting next to me told me I was rude, just because I was having a chat during the lecture.	rude

A.3 Experiments 3 and 4: Material Pretest

A.3.1 Statement Generation

A.3.1.1 Participants

We collected data from 110 participants (54 female; average age $M = 37.2$ years, $SD = 11$, ranging from 19 to 67). Participants were recruited via Amazon Mechanical Turk and received monetary compensation of 1.50 USD for the duration of seven minutes.

A.3.1.2 Procedure

Participants were briefed shortly about definitions of states and traits and were presented with a randomly selected subset of seven out of 77 state words. Participants were asked to describe one or more behaviors that imply the presented state, but simultaneously might be explained by a stable personality characteristic.

A.3.1.3 Results

Participants generated 819 dual-implication behavioral statements. After a quality check and de-duplication, 288 statements were selected and proofread and corrected by a professional lector, who was a native English speaker.

A.3.2 Pretest

A.3.2.1 Participants

Pretest analyses rely on valid data from a total of $N = 172$ participants (73 female, 1 other, 1 unspecified; average age $M = 37$ years, $SD = 11$, ranging from 19 to 69). Participants were recruited via Amazon Mechanical Turk and received monetary compensation of 2.30 USD for the duration of nine minutes.

A.3.2.2 Procedure

Participants were randomly assigned to one of six subsets of 48 dual-implication statements. Participants were asked to write down both a stable trait and a temporal state that came to mind upon reading each statement, that could describe the person making the statement. The questions were displayed on the same page. The order of questions was randomized between subjects. Subsequently, participants provided demographic information.

A.3.2.3 Results

For each statement, we grouped synonyms and calculated the proportions of each trait and state word as a consensus score. We selected 41 trait- and state-implying statements four use in our main studies based on the criteria that (a) the implied trait and state both reached at least 50% consensus, and (b) the implied trait and state had similar consensus ratings. Selected statements and their pretest results are displayed in Table A.10.

A.3.2.4 Stimulus selection

The statements selected for Experiment 3 did not differ in agreement rate (trait: M = 62%, SD = 10.7; state: M = 67%, SD = 13.7, $t(43.612) = 0.711$, $p = .19$, $d_s = 0.384$, 95% CI [-0.971, 0.202]).

Tab. A.10.: Experiments 3 and 4: State- and trait-implying statements

Nr	Statement	Implied Trait	%	Implied State	%	Experiment
1	Josh cursed at the passerby.	rude	73.9	angry	78.3	
2	Nick jumped up and down when his team scored.	enthusiastic	57.7	excited	61.5	
3	After the date with Paul, Susan broke her rule of not kissing on the first date.	adventurous	53.8	attracted	58.3	
4	When Erin came home from the volleyball tournament, she dropped onto the couch.	athletic	52	tired	88.5	
5	Aaron didn‘t take the time to look over his work after he was done.	careless	57.1	rushed	50	
6	Laura didn’t buy the top cut of beef.[1]	cheap	62.5	broke	54.2	
7	Sarah rubbed her chin while looking at the painting.	contemplative	56	interested	62.5	
8	Pete told the waiter the food tasted good.[2]	friendly	56.5	pleased	79.2	
9	Ally stared at the stage after the magic trick was performed.	gullible	50	amazed	65.2	
10	By five in the morning, Mike had already been working on the project.	hardworking	79.2	eager	87.5	
11	Mark scrambled through his papers looking for an assignment.	disorganized	64	nervous	60	
12	Gina did not finish the food on her plate last night at dinner.	picky	52.2	full	72.7	
13	Tom cut to the front of the water fountain line.	rude	58.3	thirsty	66.7	
14	Daniel laughed at the joke.	humorous	62.5	amused	100	
15	Tracy refused to make eye contact with her friend.	shy	50	ashamed	50	
16	Maggie sighed as her unread emails piled up.	unorganized	64	overwhelmed	54.2	
17	Stephen stuffed himself at dinner.	gluttonous	60	hungry	60	
18	Elle stared at the door after hearing three loud knocks.	fearful	87	scared	69.2	3 & 4
19	After John attended the workshop, his ideas just kept coming.	creative	68	inspired	58.3	3 & 4
20	Tim’s (Eric´s) jaw dropped to the floor.	dramatic	52	surprised	83.3	3 & 4
21	Lily cried watching the tearjerker.	emotional	86.4	sad	91.3	3 & 4
22	Andrea watched her neighbor’s house to see who came and went.	nosy	66.7	curious	62.5	3 & 4
23	Greg read the book until late at night.	studious	62.5	interested	69.6	3 & 4

(continued)

Tab. A.10.: Experiments 3 and 4: State- and trait-implying statements (continued)

Nr	Statement	Implied Trait	%	Implied State	%	Experiment
24	Nikki sat by her boyfriend's hospital bed and held his hand.	caring	75	worried	50	4
25	Ryan told all his colleagues about his upcoming vacation to the Bahamas.	boastful	53.8	excited	62.5	4
26	Josh sat in the park, hanging his head.	depressed	69.6	sad	69.6	4
27	Mark was trying to solve an equation.	smart	52	focused	73.9	4
28	Rebecca smiled when she greeted the new coworker.	friendly	58.3	content	60	4
29	Brandy woke up in the middle of the night for a snack.	gluttonous	54.2	hungry	84	4
30	Jane's eyes widened as she saw the demo in real-time.	interested	53.8	amazed	66.7	4
31	Rick gambled his paycheck at the casino.	irresponsible	53.8	Impulsive	46.2	4
32	Michelle complimented the chef on the food.	kind	75	satisfied	60.9	4
33	After Jayne's phone rang, she didn't bother finishing her part of the work project.	lazy	56	distracted	52	4
34	Bill didn't get out of bed.	lazy	65.2	tired	70.8	4
35	Peter hugged his mother after hearing the good news of her cancer remission	loving	65.4	happy	61.5	4
36	Anna hit the cat with a rolled-up magazine.	mean	64	angry	70.8	4
37	When Dale tasted the overcooked broccoli, he spat it out.	picky	56.5	disgusted	83.3	4
38	Ashley rolled her eyes when her colleague talked to her on her coffee break.	rude	53.8	annoyed	53.8	4
39	Kim didn't report the sexual harassment at work.	timid	60.9	scared	66.7	4
40	When his daughter threw up, Paul left the bathroom.	uncaring	53.8	disgusted	81.8	4
41	Jim smiled at the pretty girl.	kind	55.6	attracted	52	4

Note. [1]adapted from Lee, Shimizu, Masuda, & Uleman (2017), [2]adapted from Ham & Vonk (2003)

Tab. A.11.: Experiments 3 and 4: Filler statements

Nr	Statement	Presented Trait/State	Experiment
1	Zach confidently walked into the interview room.	confidently	3 & 4
2	Kyle was so drunk that he put a tasseled lamp shade on his head at the party.	drunk	3 & 4
3	Christina remained strict when her child threw a temper tantrum.	strict	3 & 4
4	Jamie went jogging at least three times a week to stay fit.	fit	3 & 4
5	Sam bumped into the closet door and clumsily broke her glasses.	clumsy	3 & 4
6	Vanessa prepared herself a cozy bath on the long weekend.	cozy	3 & 4
7	Kevin had something romantic planned for Valentine's Day.	romantic	3 & 4
8	Kelly got very competitive when her older brother was around.	competitive	3 & 4
9	When Heather woke up from a scary nightmare, she was wide awake.	awake	3 & 4
10	Patrick felt dumb when he couldn't answer a single question on the easy test.	dumb	3 & 4
11	Amy remained calm when her boss yelled at her.	calm	3 & 4
12	Brandon was lucky and won first prize in the church raffle.	lucky	3 & 4
13	Ben's stories were so funny, they made people laugh so hard they held their sides.	funny	4
14	Sara felt sad as she thought about how bad her day had gone.	sad	4
15	Linda was fanatic about shopping whenever she heard about another big sale.	fanatic	4
16	Markus was so angry that he yelled at his children for no apparent reason.	angry	4
17	Austin was very bossy and pushed the other guy because he wouldn't be quiet.	bossy	4
18	Michael was disappointed when his date stood him up.	disappointed	4
19	Nick felt awkward introducing his ex-girlfriend to his new one.	awkward	4
20	They thought Tracy was crazy because she talked to her car to get it to start.	crazy	4
21	David was picky and sent the wine back because it was from a poor year.	picky	4
22	Melissa left the room because she was offended by the joke.	offended	4
23	Sue bumped into the open closet door and clumsily broke her glasses.	clumsily	4
24	Sharon was naive enough to think that superman really could fly.	naive	4

A.4 Experiments 1-3: Response Latency Analyses

In Experiments 1-3, we used the time participants took to press the respective response key as a measure of response latencies. We only included RTs of correct rejections into analyses. Note that the accuracy of this measurement is restricted, as Experiments 1 and 3 were conducted online, on participants' individual devices.

Tab. A.12.: Means and standard deviations of response latencies for correct rejections on target trials in the false recognition task in Experiments 1-3

Inference	Condition	Experiment 1	Experiment 2	Experiment 3
Trait	Implied	2742 (1703)[a]	2435 (1031)	2570 (2308)[c]
	Implied-Other	2741 (1951)	2405 (1051)	2324 (1494)[c]
	New	2228 (1485)[a]	2256 (987)	-
State	Implied	3488 (6125)[b]	2682 (1192)	2753 (2153)
	Implied-Other	3210 (5257)	2592 (1133)	2690 (2192)
	New	2452 (2547)[b]	2474 (961)	-

Note. [a] $t(43) = 3.762$, $p = .001$ (one-tailed), $d_z = 0.43$, 95% CI [0.192, 0.667]; [b] $t(36) = 3.271$, $p = .002$ (one-tailed), $d_z = 0.303$, 95% CI [0.114, 0.493]; [c] $t(363) = 2.712$, $p = .007$ (one-tailed), $d_z = 0.104$, 95% CI [0.029, 0.180].

A.5 Experiment 2: Supplemental Analyses

A.5.1 Analyses of Error Rates for the First 68 Valid Data Sets

Due to imbalanced randomization and an unusually high number of exclusions, we exceeded the pre-registered sample size for the interim analyses to achieve a balanced assignment of participants. We report analyses for the first 68 valid data sets below. Data of the first 68 valid data sets were submitted to a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) within-subject ANOVA. Results show a main effect for condition, $F(2, 132) = 35.07$, $p < .001$, $\eta_p^2 = .35$, 90% CI [0.36, 0.61], and for inference, $F(1, 66) = 5.92$, $p = .018$, $\eta_p^2 = .08$, 90% CI [0.01, 0.2]. We also observed an interaction between inference and condition, $F(2, 132) = 6.58$, $p = .002$, $\eta_p^2 = .09$, 90% CI [0.04, 0.28]. Separate analyses confirmed spontaneous person inference effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied condition (M = .29, SD = .25) than in the implied-other trait condition (M = .20, SD = .22), $t(39) = 2.36$, $p = .023$ (one-tailed), $d_z = 0.373$, 95% CI [-0.08, 0.82], and in the new trait condition (M = .13, SD = .17), $t(39) = 4.016$, $p < .001$ (one-tailed), $d_z = 0.635$, 95% CI [0.18, 1.09]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition (M = .50, SD = .23) than in the implied-other state condition (M = .28, SD = .20), $t(27) = 4.221$, $p < .001$ (one-tailed), $d_z = 0.798$, 95% CI [0.24, 1.35], and in the new state condition (M = .14, SD = .19), $t(27) = 7.367$, $p < .001$ (one-tailed), $d_z = 1.392$, 95% CI [0.8, 1.99].

A.5.2 Pre-Registered Contrasts

We had pre-registered a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) within-subject ANOVA with planned contrasts as main analyses for this Experiment. However, we decided to report t-tests in the paper, to be consistent with the reporting of results of Experiments 1, 3, and 4. Both the t-tests and the planned contrasts yield significant results:

In the trait inference condition, participants showed higher false recognition rates in the implied condition (M = .29, SD = .24) than in the implied-other trait condition (M = .20, SD = .21), $F(1,41) = 18.798$, $p < .001$, $\eta_p^2 = .31$, 90% CI [.13, .47], and in the new trait condition (M = .14, SD = .18), $F(1,41) = 64.981$, $p < .001$, $\eta_p^2 = .61$, 95% CI [.44, .71]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition (M = .51, SD = .22) than in the implied-other state condition (M = .31, SD = .21), $F(1,43) = 7.721$, $p = .008$, $\eta_p^2 = .15$, 95% CI [.02, .31], and in the new state condition (M = .18, SD = .20), $F(1,43) = 17.842$, $p < .001$, $\eta_p^2 = .29$, 95% CI [.11, .45].

A.5.3 Navon Task

A.5.3.1 Method

The Navon task (Navon, 1977) was programmed in Inquisit 4 (www.millisecond.com). In this task, participants saw stimuli that appeared like the capital letters F, H, L, or T in large print, comprised of the same letters in smaller print - such that each stimulus presented one of the target letters H or L in either large or small print, respectively. Participants were asked to indicate via key-press which of the letters ‘H’ or ‘L’ were present in a stimulus. The task consisted of 50 trials presented in individual random order after completion of ten practice trials. We aimed to use this task as an indirect measure of global versus local processing dominance, in order to explore whether this would be associated with the strength of state and trait inferences. We thus computed a Navon index as the average response latency difference between trials in which the target letter was presented as large letter (global condition) or as small letter (local condition).

A.5.3.2 Results

We calculated a Navon index by substracting participant’s average response latency in the Local condition from their average response latency in the Global condition. We correlated the Navon index with the mean differences between implied and implied-other as well as new trials. The Navon index did not correlate with the mean difference between implied traits and implied-other traits, $r(41) = .12$, $p = .449$; nor with the mean difference between implied traits and new traits, $r(41) = -.03$, $p = .850$; nor with the mean difference between implied states and implied-other states, $r(40) = .19$, $p = .233$; nor with the mean difference between implied states and new states, $r(40) = .26$, $p = .099$.

A.5.4 Implicit Personality Theories

We derived a continuous measure of participants’ implicit personality theories by substracting mean responses to items measuring incremental beliefs from mean responses to items measuring entity beliefs (see Table A.13). We then correlated this measure of participants’ implicit personality theory with the mean differences between implied and implied-other as well as new trials. While the correlation between implicit personality theories and the mean difference between implied traits and new traits approached significance, $r(42) = .29$, $p = .056$, no other correlations were significant, neither the correlation with the mean difference between implied traits and implied-other traits, $r(42) = .07$, $p = .633$; nor with the mean difference between implied states and implied-other states, $r(40) = .17$, $p = .290$; nor with the mean difference between implied states and new states, $r(40) = .14$, $p = .387$.

Tab. A.13.: Items measuring implicit personality theories, original and translated.

Nr	Mindset	Item
1*	E	The kind of person someone is something very basic about them and it can't be changed very much.
		Die Art von Mensch die jemand ist, ist ein sehr elementarer Teil einer Person und lässt sich nicht ohne Weiteres verändern.
2*	E	People can do things differently, but the important parts of who they are can't really be changed.
		Menschen können sich unterschiedlich verhalten, aber ihre wesentlichen Eigenschaften lassen sich nicht wirklich ändern.
3	I	Everyone, no matter who they are, can significantly change their basic characteristics.
		Jede/r, ganz egal wer er oder sie ist, kann sich von Grund auf verändern.
4*	E	As much as I hate to admit it, you can't teach an old dog new tricks. People can't really change their deepest attributes.
		So ungern ich es zugebe, was Hänschen nicht lernt, lernt Hans nimmermehr. Menschen können ihre tiefsten Eigenschaften nicht wirklich verändern.
5	I	People can always substantially change the kind of person they are.
		Menschen können sich immer von Grund auf verändern.
6*	E	Everyone is a certain kind of person and there is not much that can be done to really change that.
		Jede/r ist eine bestimmte Art von Mensch, und daran lässt sich nicht viel ändern.
7	I	No matter what kind of person someone is, they can always change very much.
		Egal, was für eine Art Mensch jemand ist – sie/er kann sich immer stark ändern.
8	I	All people can change even their most basic qualities.
		Alle Menschen können sich ändern – selbst in ihren elementarsten Eigenschaften.

Note. * reverse-coded item. E: entity mindset. I: incremental mindset. German translations are italicized.

A.6 Experiment 3: Supplemental Analyses

A.6.1 Correlation Analyses

We calculated the correlation between the state inference and trait inference effects on the by-participant and the by-item level. On the participant level, we observed no significant correlation, $r(363) = .00$, 95%CI [-0.10, 0.10], $p = .966$. On the item level, we observed a significant positive correlation, $r(22) = .43$, 95%CI [0.03, 0.71], $p = .037$. These results indicate that participants had no individual response tendencies towards either trait or state inferences, nor that individual behavioral statements prompted either state or trait inferences.

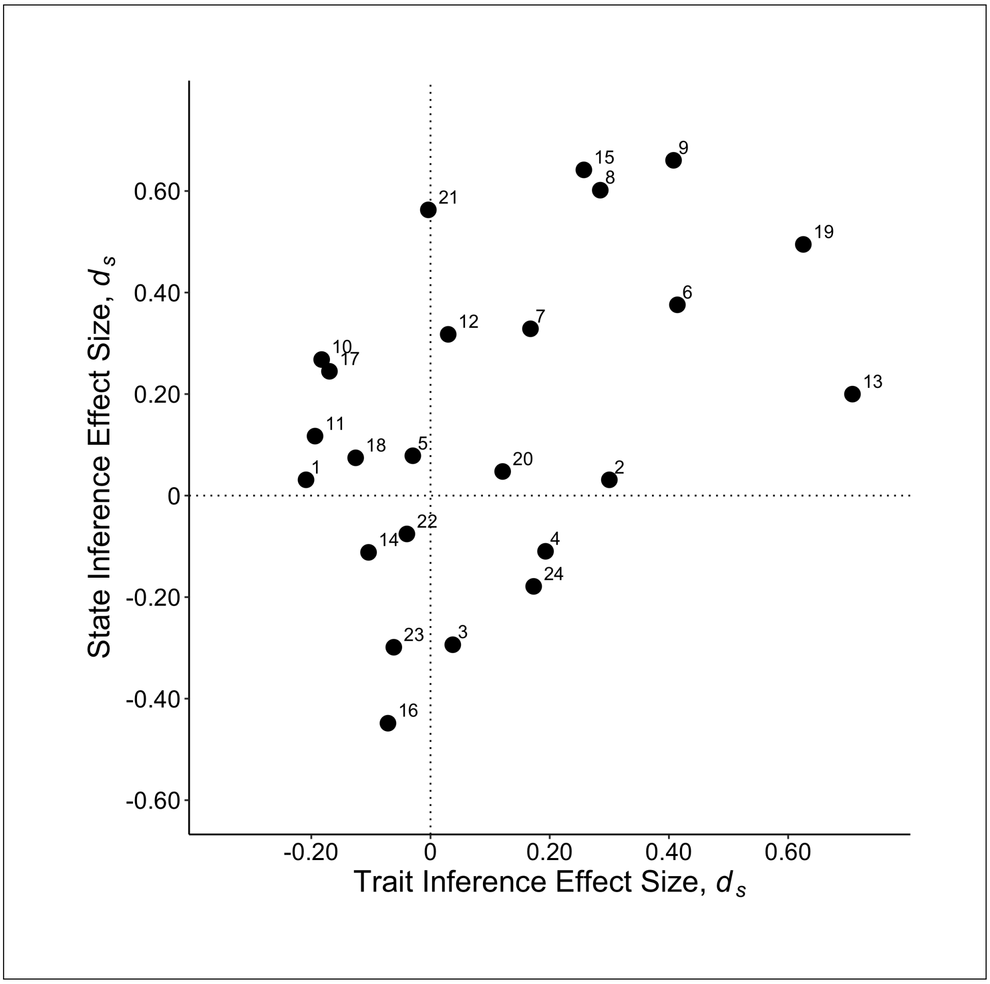

Fig. A.1.: Experiment 3: Effect sizes for state and trait effects per item

A.6.2 Analyses of Error Rates

False recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject ANOVA. Results show significant main effects for Inference, $F(1, 340) = 102.39$, $p < .001$, $\eta_p^2 = .23$, 90% CI [0.17, 0.29], and for Condition, $F(1, 340) = 197.10$, $p < .001$, $\eta_p^2 = .37$, 90% CI [0.3, 0.42], with no significant interaction, $F(1, 340) = 111.74$, $p < .001$, $\eta_p^2 = .25$, 90% CI [0.18, 0.31]. Separate analyses confirmed spontaneous person inference effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied trait condition ($M = 0.034$, $SD = 0.947$) than in the implied-other trait condition ($M = 0.019$, $SD = 0.962$), $t(340) = 5.674$, $p < .001$ (one-tailed), $d_z = 0.304$, 95% CI [0.196, 0.412]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition ($M = 0.079$, $SD = 0.917$) than in the implied-other state condition ($M = 0.02$, $SD = 0.953$), $t(340) = 15.276$, $p < .001$ (one-tailed), $d_z = 0.818$, 95% CI [0.696, 0.939]. Effect sizes differed significantly between state and trait condition, $t(340) = 10.893$, $p < .001$ (one-tailed), $d_z = 0.729$, 95% CI [0.581, 0.876].

A.6.3 Correlation Analyses

We calculated the correlation between the state inference and trait inference effects on the by-participant and the by-item level. We observed no significant correlations, neither on the participant level, $r(363) = .00$, 95%CI [-0.10, 0.10], $p = .966$, nor on the item level, $r(339) = .06$, 95%CI [-0.04, 0.17], $p = .263$. These results indicate that participants had no individual response tendencies towards either trait or state inferences, nor that individual behavioral statements prompted only either state or trait inferences.

A.6.4 By-Item Analyses

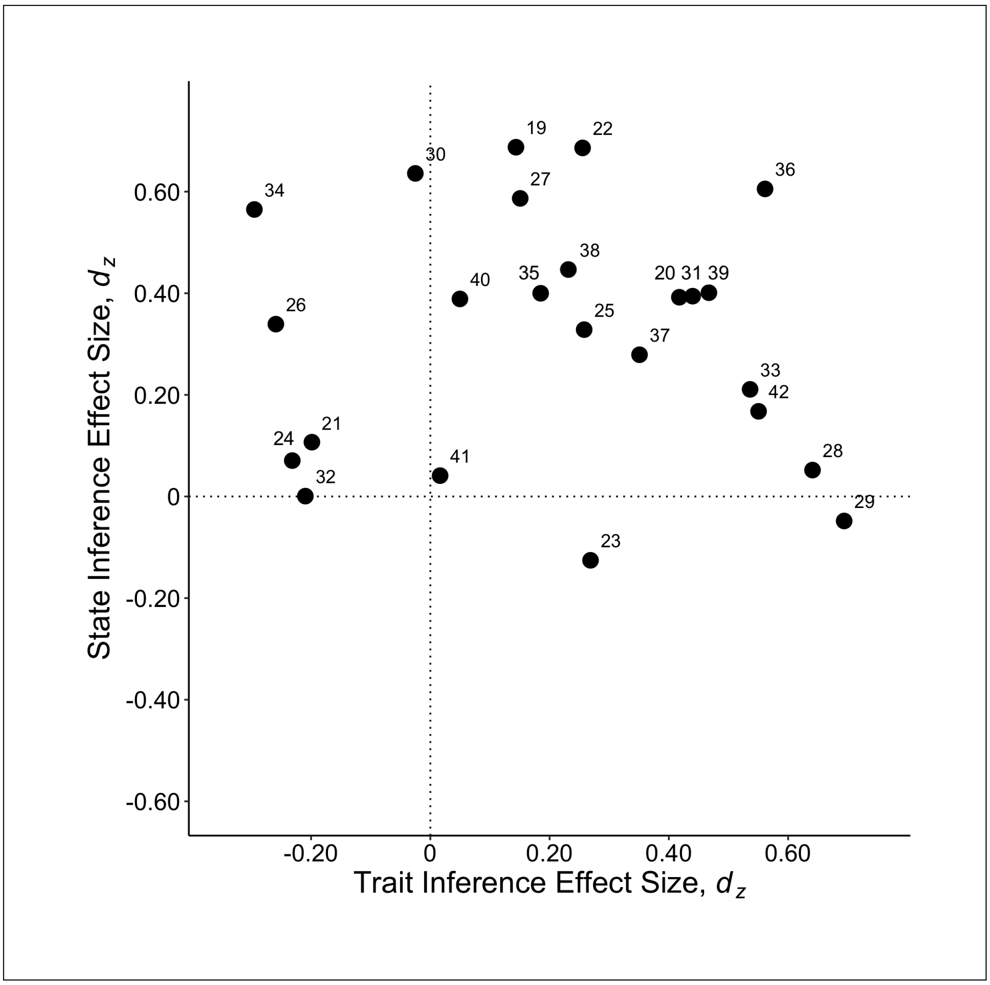

Fig. A.2.: Experiment 4: Effect sizes for state and trait effects per item

A.7 Experiment 4: Supplemental Analyses

A.7.1 Analyses by Trimming Criteria

Tab. A.14.: Experiment 4: Trimming criteria and effect sizes

Trimming Criterion	Trait Inference (Cohen's d_z)	State Inference (Cohen's d_z)	Interaction (η_p^2)
No Transformation			
No Trim	0.580	0.790	.03
$\pm 2.5SD$	0.529	0.887	.04
2500 > RT > 0	0.663	0.982	.07
2000 > RT > 0	0.699	0.973	.06
1500 > RT > 0	0.786	1.075	.07
Log transformation			
No Trim	0.724	0.973	.05
$\pm 2.5SD$	0.668	1.002	.05
2500 > RT > 0	0.716	1.073	.07
2000 > RT > 0	0.741	1.045	.06
1500 > RT > 0	0.806	1.117	.06
Inverse transformation			
No Trim	0.763	1.072	.06
$\pm 2.5SD$	0.716	1.051	.04
2500 > RT > 0	0.736	1.110	.07
2000 > RT > 0	0.754	1.072	.06
1500 > RT > 0	0.805	1.116	.05

Note. All effects reached statistical significance with all trimming criteria ($p < .01$). Effect sizes d_z stem from t-tests implied vs. implied-other conditions, η_p^2 stems from the 2 x 2 interaction between inference and condition.

Tab. A.15.: Experiment 6 & 9: Target statements

Nr	Statement	Implied Trait	%	Implied State	%
1	He bought the new car straight away.	impulsive	64	excited	50
2	She hit the cat with a rolled-up magazine.	mean	64	angry	70.8
3	She rolled her eyes when her colleague talked to her on her coffee break.	rude	53.8	annoyed	53.8
4	He didn't get out of bed.	lazy	65.2	tired	70.8
5	He didn't take the time to look over his work after he was done.	careless	57.1	rushed	50
6	She stared at the door after hearing three loud knocks.	fearful	87	scared	69.2
7	He rapidly sketched portraits in his notebook.	artistic	80	inspired	48
8	He didn't sweat at all when questioned by the cop.	confident	76.9	calm	61.5
9	A wry smile spread across her lips as she jumped from the airplane.	adventurous	52	thrilled	48
10	He read the book until late at night.	studious	62.5	interested	69.6
11	She hugged her mother after hearing the good news of her cancer remission	loving	65.4	happy	61.5
12	She complimented the chef on the food.	kind	75	satisfied	60.9
13	He smiled when he greeted the new coworker.	friendly	58.3	content	60
14	She laughed at the joke.	humorous	62.5	amused	100
15	When she came home from the volleyball tournament, she dropped onto the couch.	athletic	52	exhausted	88.5
16	Faced with so many possibilities, he didn't know what to do.	indecisive	54.2	confused	52.2
17	He sat in the park, hanging his head.	depressed	69.6	sad	69.6
18	She did not finish the food on her plate last night at dinner.	picky	52.2	full	72.7
19	She covered her face with her hand.	shy	62.5	embarrassed	59.1
20	He stumbled out of the bar and started a fight.	violent	60	drunk	71.4
21	He stepped on his girlfriend's toes during the foxtrot.	clumsy	75	nervous	52.5
22	She finished more tasks than he was assigned.	dilligent	61.5	motivated	69.6
23	She sat by her boyfriend's hospital bed and held his hand.	caring	75	worried	50
24	She was trying to solve an equation.	smart	52	focused	74

Tab. A.16.: Experiment 6 & 9: Filler statements

Nr	Statement	Presented Adjective
1	She was very competent at her job.	competent
2	He helped him because he was social.	social
3	He was clumsy, he bumped into the closet door and broke her glasses.	clumsy
4	She was very competitive, especially when her brother was around.	competitive
5	She was naive enough to think that superman really could fly.	naive
6	He felt proud.	proud
7	She felt loved.	loved
8	He was disappointed.	disappointed
9	He was lucky and won first prize in the church raffle.	lucky
10	She left the room because she was offended by the joke.	offended
11	She got distracted and stepped on her boyfriend's feet at the posh dance event.	distracted
12	He confidently walked into the interview room.	confidently

Tab. A.17.: Experiment 7: Target statements

Nr	Statement	Implied Trait	%	Implied State	%
1	After John attended the workshop, his ideas just kept coming.	creative	68	inspired	58.3
2	After the date with Paul, Susan broke her rule of not kissing on the first date.	adventurous	53.8	attracted	58.3
3	Dave bought the new car straight away.	impulsive	64	excited	50
4	Anna hit the cat with a rolled-up magazine.	mean	64	angry	70.8
5	Ashley rolled her eyes when her colleague talked to her on her coffee break.	rude	53.8	annoyed	53.8
6	Ben didn't sweat at all when questioned by the cop.	confident	76.9	calm	61.5
7	Bill didn't get out of bed.	lazy	65.2	tired	70.8
8	Frank honked his horn at the yellow light.	impatient	79.2	frustrated	60.9
9	Elle stared at the door after hearing three loud knocks.	fearful	87	scared	69.2
10	Gina did not finish the food on her plate last night at dinner.	picky	52.2	full	72.7
11	Greg read the book until late at night.	studious	62.5	interested	69.6
12	Jane's eyes widened as she saw the demo in real-time.	interested	53.8	amazed	66.7
13	Josh sat in the park, hanging his head.	depressed	69.6	sad	69.6
14	Michelle complimented the chef on the food.	kind	75	satisfied	60.9
15	Nikki sat by her boyfriend's hospital bed and held his hand.	caring	75	worried	50
16	Peter hugged his mother after hearing the good news of her cancer remission	loving	65.4	happy	61.5
17	Hank finished more tasks than he was assigned.	dilligent	61.5	motivated	69.6
18	Rebecca smiled when she greeted the new coworker.	friendly	58.3	content	60
19	Sam read the menu for the third time.	indecisisve	50	undecided	52.2
20	Steve laughed at the joke.	humorous	62.5	amused	100
21	Sarah covered her face with her hand.	shy	62.5	embarrassed	59.1
22	Tom stumbled out of the bar and started a fight.	violent	60	drunk	71.4
23	Tracy scrambled through her papers looking for an assignment.	disorganized	64	nervous	60
24	When Erin came home from the volleyball tournament, she dropped onto the couch.	athletic	52	exhausted	88.5

Tab. A.18.: Experiment 7: Filler statements

Nr	Statement	Presented Adjective
1	When all her friends surprised her for her birthday, she felt loved.	loved
2	When his date stood him up, he was disappointed.	disappointed
3	He was lucky and won first prize in the church raffle.	lucky
4	She left the room because she was offended by the joke.	offended
5	She got distracted and didn't hear her phone ringing.	distracted
6	He was late to the meeting.	late
7	She was very competent at her job.	competent
8	When his colleague asked, he helped him because he was social.	social
9	He prepared himself a cozy bath on the weekend.	cozy
10	She was very competitive, especially when her brother was around.	competitive
11	She was naive enough to think that superman really could fly.	naive
12	When he could answer all the questions on the test, he felt proud.	proud

A.10 Experiments 6 & 7: Supplemental Analyses

Tab. A.19.: Means and standard deviations of response latencies for correct rejections on target trials in the false recognition task in Experiments 6 & 7

Inference	Information	Condition	Experiment 6	Experiment 7
Trait	Actor	Implied	2693 (2403)[a]	2455 (1622)
		Implied-Other	2644 (2191)[a]	2098 (1008)
	Communicator	Implied	2672 (2401)	2443 (2038)
		Implied-Other	2636 (2707)	2366 (1392)
State	Actor	Implied	3103 (5028)	2466 (1423)
		Implied-Other	2773 (2829)	2433 (1172)
	Communicator	Implied	2769 (3044)[b]	2526 (1026)
		Implied-Other	2891 (4231)[b]	2427 (1246)

Note. [a] $t(170) = 0.24$, $p = .811$ (one-tailed), $d_z = 0.012$, 95% CI [-0.09, 0.114]; [b] $t(180) = 0.188$, $p = .851$ (one-tailed), $d_z = 0.008$, 95% CI [-0.077, 0.094]

A.11 Experiments 8 & 9: Supplemental Analyses

Tab. A.20.: Means and standard deviations of response latencies (in ms) for correct rejections on target trials in the false recognition task in Experiments 8 & 9

Inference	Information	Condition	Experiment 8	Experiment 9
Trait	No Load	Implied	2079 (638)	-
		Implied-Other	2207 (638)	-
	Low Load	Implied	2163 (778)	1944 (795)
		Implied-Other	2101 (733)	1871 (752)
	High Load	Implied	1969 (910)	1863 (831)
		Implied-Other	2075 (1040)	1899 (895)
State	No Load	Implied	2151 (909)	-
		Implied-Other	2124 (676)	-
	Low Load	Implied	2197 (849)	2030 (982)
		Implied-Other	2085 (654)	2078 (1025)
	High Load	Implied	2072 (1115)	1999 (995)
		Implied-Other	2019 (797)	1946 (900)

Note. No significant differences between any conditions.

Appendix B: Kruse & Degner (2021)

B

Parts of this thesis were published in Kruse, F. & Degner, J. (2021). Spontaneous state inferences. *Journal of Personality and Social Psychology, 121*(4), 774-791.

Spontaneous State Inferences

Felix Kruse and Juliane Degner
Universität Hamburg

Perceivers routinely draw inferences about others from their behavior in an attempt to make sense of the world. Previous research has established that spontaneous inferences include stable characteristics such as traits and a number of variable person-related concepts such as goals, intentions, and motivations. The current research investigated the occurrence of more general spontaneous state inferences. In a series of four pre-registered studies (N = 883), we adapted two established experimental paradigms frequently used in spontaneous social inference research to the investigation of spontaneous trait and state inferences. In Studies 1 and 2, we observed evidence for the occurrence of spontaneous state inferences from state-implying statements. In Studies 3 and 4, we observed the simultaneous occurrence of spontaneous trait and state inferences from statements that allowed for both inferences. In a fifth study (N = 97), we provide evidence that people represent states and traits as functionally different: Participants judged the likelihood of behavioral repetition higher when the same behavior was related to a trait-inference than a state-inference. The observation of multiple simultaneous spontaneous inferences in the current research suggests that further theory building regarding the underlying mechanisms and processes of spontaneous impression formation in person perception from behavior is warranted.

People routinely attempt to make sense of the world they perceive around them. In order to do so, perceivers often go beyond the information that is available to them. One percept of particular interest to lay people and psychologists alike are other people. For our everyday life, making sense of the people around us is paramount: It helps us deal with the sheer complexity of our social environment by attributing causes to observed events (e.g., Kelley, 1967) and predicting future behavior (e.g., Dunning, Griffin, Milojkovic, & Ross, 1990; McCarthy & Skowronski, 2011b; Nussbaum, Trope, & Liberman, 2003), thereby reducing uncertainty (e.g., Heider, 1958; Trope & Gaunt, 2000) and maintaining the illusion of controllability (Biner, Angle, Park, Mellinger, & Barber, 1995; Bruner, 1957).

When trying to make sense of peoples' behavior, observers have a tendency to disporportionally focus on personal dispositions as presumed causes of behavior and disregard situational factors (Heider, 1958; Malle, 2008; Olcaysoy Okten & Moskowitz, 2018), a phenomenon known as the correspondence bias (Jones & Davis, 1965), previously termed fundamental attribution error (Ross, 1977). In psychological research and theorizing, these personal dispositions have frequently been equated with personality *traits*: internal, enduring, and invariant qualities underlying the individual differences in peoples' experiences and behaviors across contexts (Cantor & Mischel, 1977, 1979). Classic theories on impression formation from behavior focused entirely on traits as dispositional attributions (e.g., Gilbert, Pelham, & Krull, 1988; Trope, 1986). Some even assumed that trait attributions occur spontaneously and effortlessly once behavior is identified, whereas the inclusion of further factors that might have caused or constrained the actor's behavior requires a more effortful, deliberate information search and processing (e.g., Krull & Erickson, 1995; Trope & Gaunt, 2000).

Spontaneous Trait Inferences

A large body of research has indeed supported the assumption that people form trait inferences *spontaneously* when confronted with an actor's behavior. If we learn, for instance, that Daniel laughed at a joke, we may infer that Daniel is a *jolly* person, thus assuming a trait as cause of the observed behavior - a *spontaneous trait inference* (STI; for recent reviews, see Moskowitz & Olcaysoy Okten, 2016; Uleman, Rim, Saribay, & Kressel, 2012). When drawing such inferences, we go beyond the given information and supplement self-generated information, for example by assuming that Daniel's single behavior was representative of an underlying stable disposition to be good-humoured. Research has demonstrated that we draw these inferences instantaneously when encoding behavioral information (Uleman, Hon, Roman, & Moskowitz, 1996), and require only a minimum of our attention (Todorov & Uleman, 2003) and cognitive resources to do so (Wells, Skowronski, Crawford, Scherer, & Carlston, 2011). More importantly, we draw these inferences independent of whether or not we have an explicit intention to form an impression of an actor (e.g., Carlston, Skowronski, & Sparks, 1995; Carlston & Skowronski, 1994; Ferreira et al., 2012; McCarthy & Skowronski, 2011a), and even when it counteracts our current processing goals (e.g., interfering with current task performance, Ham & Vonk, 2003; Todd, Molden, Ham, & Vonk, 2011; Todorov & Uleman, 2004, 2002, 2003), or we actively try to avoid impressions (e.g., Shimizu, 2017). This spontaneity and uncontrollability of trait inferences may offer one explanation why perceivers so often fall prone to the correspondence bias (e.g., Moskowitz, 2005).

Multiple Social Inferences

Spontaneous trait inferences have been investigated intensively, and the evidence in their support has repeatedly been called ubiquitous (Schneid, Carlston, & Skowronski, 2015; Uleman, 2005). Unfortunately, it seems that the ubiquity of evidence for the occurrence of spontaneous trait inferences is often misconceived as evidence for the ubiquity of trait inferences in impression formation from behavior. For example, many established textbooks in social psychology refer only to trait inferences when discussing spontaneous person perception (e.g., Hewstone, Stroebe, & Jonas, 2016; Myers & Twenge, 2018). This representation, however, overlooks two important aspects: On the one hand, researchers have repeatedly discussed and empirically documented multiple spontaneous inference activation, especially simultaneously occurring person and situation attributions (e.g., Lupfer, Clark, & Hutcherson, 1990; Todd et al., 2011). On the other hand, person dispositions as attributions of behavior can encompass more than only personality traits. For example, studies on deliberate impression formation have documented a wide array of different person inferences, ranging from intentions and desires to values and beliefs (e.g., Malle & Holbrook, 2012), which questions the presumed inevitability or priority of trait inferences. In this line, a number of studies on spontaneous impression formation demonstrated spontaneous goal inferences (e.g., Hassin, Aarts, & Ferguson, 2005; Olcaysoy Okten & Moskowitz, 2019; Van Overwalle, Van Duynslaeger, Coomans, & Timmermans, 2012), and motivational inferences (e.g., Reeder, 2009a, 2009b; Reeder, Hesson-McInnis, Krohse, & Scialabba, 2001). Thus, when making sense of others' behavior, perceivers are able to spontaneously draw multiple inferences: I may infer that Daniel is a *jolly person*, while simultaneously inferring that the *joke* he is laughing about may be very *funny*, and/or that Daniel may have the current goal of *being entertained* and having a good time.

In light of these more recent findings, it seems clear that spontaneous impressions of others are not limited to traits, but may encompass a variety of person attributions. Evidence for some of these attributions already exists in the literature. We argue that these may not be limited to goals and motivations, but more generally include inferences about an actor's current mental states, such as affective or emotional states (e.g., happy, surprised, anxious, ashamed), cognitive and attentional states (e.g., focused, interested, distracted, bored) and physiological states (e.g., hungry, full, sick, tired). We not only presume that people *spontaneously* draw multiple inferences from behaviors but also that these multiple inferences occur *simultaneously* upon encoding others' behaviors.

Spontaneous State Inferences

Psychological theorizing and research has separated traits from states in a number of different ways. A common theoretical account defines traits as stable, inter-individual differences in peoples' proneness, tendency, style, or disposition to behave, feel, or think in certain ways (e.g., Hamaker, Nesselroade, & Molenaar, 2007). States also describe person dispositions for thoughts, feelings, or behaviors. However, states refer to transient *intra*-individual differences that reflect peoples' continuous adaptation to situational demands. Thus, whereas traits are typically conceptualized by their relative stability, consistency, and in-variance over time and across situations, states are characterized by their relative instability, inconsistency, and variance (Hamaker et al., 2007). In our research, we thus operationalize traits and states by their relative temporal stability: Although by no means a binary criterion, traits tend to endure over time, while states are rather short-lived phenomena. This implies that trait and state inferences could hold a different functional value for the perceiver because they may influence predictions about future behavior in different ways.

There are several theoretical and empirical perspectives that support the assumption of spontaneous state inferences. First, and most importantly, some theoretical accounts of spontaneous impression formation actually take state inferences as a given (e.g., Uleman, 2005), at times even regard them as the "default mode" of understanding behavior (Korman & Malle, 2016; Malle & Holbrook, 2012). For example, Heider (1958) originally assumed in his considerations of peoples' naive analyses of action that person inferences for causal attributions of behavior include contemplations of transitory person states such as fatigue and mood, attitudes and needs, or social and legal status. Similarly, in their considerations of correspondence inferences, Jones & Davis (1965) argue that trait inferences about actors rely on inferences about actors' intentional states. A similar argument is brought forward in Reeder's multiple inference model (MIM; Reeder, 2009a, 2009b). It is striking that although several of the most seminal theoretical accounts of interpersonal impression formation and attribution have explicitly addressed state inferences, research on spontaneous impression formation has not given these inferences that much attention. Indeed, the need to investigate whether and to what extent people spontaneously infer states from behavior has been stated repeatedly (e.g., Lillard & Skibbe, 2006; Uleman, Saribay, & Gonzalez, 2008).

Mental states have seen some more attention in research on deliberate impression formation. There is indeed empirical evidence that people infer mental states from behavior when explicitly prompted to form impressions (e.g, Ames, 2004) or when asked to write about their impressions of others (McClung & Reicher, 2018), and readily explain intentional actions in terms of beliefs, desires, values, and internal states (e.g., Malle, 2004; Malle & Holbrook, 2012; Olcaysoy Okten & Moskowitz, 2019, Experiment 3). Related research in the field of developmental psychology has repeatedly documented children's and adults' ability to deliberately infer and use others' affective and cognitive states, termed *mentalizing* or *theory of mind* (e.g., Kiley Hamlin, Ullman, Tenenbaum, Goodman, & Baker, 2013; Ruffman, 2014; Scott & Baillargeon, 2017). However, we lack empirical research investigating if and to what extent any of the above listed state inferences occur spontaneously (e.g., Apperly, Riggs, Simpson, Chiavarino, & Samson, 2006; Lillard & Skibbe, 2006).

There is another field of research, which may provide important insights for the current research question, namely the field of text comprehension. Given that most STI research relies on written statements about others' behavior (but see Fiedler & Schenck, 2001; Fiedler, Watling, Menges, & Schenck, 2005), theorizing and research on text comprehension may provide helpful insights to the question which inferences people draw spontaneously from written behavior statements. There are indeed several theoretical accounts of text comprehension that describe spontaneous inferences about mental states from text (e.g., Cook & O'Brien, 2017; Graesser, Singer, & Trabasso, 1994; Kintsch, 1998), and empirical evidence for their occurrence has been provided, at least for emotional states (e.g., Diergarten & Nieding, 2016; Gernsbacher, Goldsmith, & Robertson, 1992). However, it remains an open question if and to what extent text comprehension effects actually represent inferences about actors of behavior or merely represent mere activation of semantic associations

There are good reasons why mental state inferences could and should occur *spontaneously* in the social inference process. First, research on person perception has already documented that people are generally able to spontaneously process situational information when forming impressions (e.g., Reeder et al., 2001). Trope's (1986) model of impression formation and attribution even assumes that situational information is initially processed and used in order to identify observed behavior (i.e., situational inducement). Technically, we cannot assume that situations directly *cause* behavior but have to assume that situations or their appraisals impact people who in turn respond with behavior. Previous research has indeed documented that people process others' mental states when thinking about eliciting situations (e.g., Thornton & Tamir, 2020).

Second, it has been repeatedly argued that the presumed dominance of trait inferences in impression formation results from their high functional value: Knowing a person's stable dispositions allows predicting their future behavior (e.g., Heider, 1958; Hoffman, Mischel, & Mazze, 1981; McCarthy & Skowronski, 2011b). However, drawing state inferences can be equally informative. On the one hand, inferring others' mental states does enable a perceiver to derive situation-specific short-term behavioral expectations and tailor their own responses accordingly (e.g., Thornton & Tamir, 2020). On the other hand, state inferences can also signal that the current observation may not warrant predicting future behavior, or that behavior predictions should be limited to the very short-term, thus preventing erroneous over-generalizations about others. Consequently, spontaneous state inferences may turn out to be equally functional and adaptive as spontaneous trait inferences.

Finally, recent theorizing in personality psychology actually defines the person-descriptive aspects of traits as density distributions of states (i.e., whole trait theory; Fleeson & Jayawickreme, 2015; Jayawickreme, Zachry, & Fleeson, 2019). In this understanding, a trait-ascription is related to a person's frequent manifestations of trait-related states. For example, a person who is described with the trait label *shy* is expected to *feel and act in in shy ways* frequently and in many situations, an idea originating with Mischel (1968). Applying whole trait theory to person perception suggests that in order to ascribe a trait to an actor, observers might need to recognize the actor's current state, infer that the actor experiences this state frequently and in many situations, and thus generalize that this current state is a representative manifestation of an underlying trait.

In summary, there are several theoretical and empirical reasons that support our notion that perceivers may spontaneously infer mental states from behaviors. When referring to mental states, we explicitly include *any* behavior-related person condition that is temporally limited. We thus propose a wider range of inferences than the considerations of intentionality, desire, and belief proposed by Malle (e.g., Malle, 2005; Malle & Holbrook, 2012), and further include affective, emotional, cognitive, attentional, and physiological states.

The goal of our research was (a) to provide empirical evidence that state inferences can occur spontaneously, (b) to investigate whether perceivers can simultaneously draw trait and state inferences from the same behaviors or whether trait and state inferences from the same behavior would be mutually exclusive, and (c) to demonstrate that state and trait inferences have a different functional value for perceivers.

The Current Research

We conducted a series of five studies using different stimulus materials and established experimental paradigms. In the first set of two studies, we employed single-implication behavioral descriptions that allow for either unambiguous trait inferences (e.g., "I gave the homeless man five euros"–generous) or unambiguous state inferences (e.g., "When my sister and her husband exchanged rings, I just couldn't hold back the tears"–touched). We used a false recognition paradigm (Todorov & Uleman, 2002), collected data online (Study 1), and compared results with data collection in the lab (Study 2). We then developed a second set of dual-implication behavioral stimuli that allowed for simultaneous trait and state inferences (e.g., "Vanessa read the book until late at night"–studious, interested) which we employed in a false recognition paradigm (Study 3) and a probe recognition paradigm (Study 4; Todd et al., 2011). To foreshadow results: All four studies provided robust evidence for the spontaneous and simultaneous occurrence of state and trait inferences. Finally, in Study 5, we explored if and to what extend people functionally distinguish between state and trait inference in impression formation, focusing on the perceived predictability of future behavior. All studies were preregistered with the Open Science Framework (https://osf.io/v5j78/) where we also provide open access to materials, raw data, and analyses codes. We report how we determined our sample sizes, all data exclusions, all manipulations, and all measures in the description of each study and the online supplemental materials. The faculty's local ethics committee approved all procedures (protocol number: 2018_180).

Material generation

For the current research endeavor, it was paramount to create stimulus materials that allow for a valid differentiation between state and trait inferences without favouring either inference. We therefore conducted a pilot study and a series of pretests in order to develop one set of single-implication behavioral statements that exclusively allow for either trait or state inferences (to be used in Studies 1 and 2) and another set of dual-implication behavioral statements that allow for simultaneous trait and state inferences (to be used in Studies 3-5). We summarize this extensive preparatory research here and provide detailed descriptions and results of all pretests in the online supplemental materials.

Given that our studies relied on established experimental paradigms that use adjective probes, we first conducted a pilot study to establish a set of person-describing adjectives that people use unequivocally to either refer to traits or states. We briefed a sample of participants about trait and state concepts and asked them to rate the relative stability vs. variability of a list of 323 adjectives. Based on the resulting ratings, we chose those 69 adjectives that were rated as the most stable as traits (e.g., introverted, smart, ambitious) and the 72 most variable as states (e.g., bewildered, thirsty, disgusted) as basis for further stimulus generation. In a next step, we generated a set of 196 statements in German language that described single behaviors implying these selected states or traits (without explicitly mentioning the respective adjectives). We ensured that behavioral statements focused on relatively mundane behaviors and did not contain any additional linguistic markers that might elicit trait or state inferences. We submitted these statements to a pretest, in which we asked participants to name a person-describing adjective that came to mind when reading a statement. For Study 1 and 2, we selected only statements that reached a consensus score of 50% or higher.[1] The selected trait-implying and state-implying statements did not differ significantly in agreement rate or statement length (see online supplemental materials).

For Studies 3 to 5, we created a second set of dual-implication behavioral statements in English language that described more ambiguous behaviors potentially allowing for simultaneous trait and state inferences. For statement creation, we relied on a large sample of crowd-workers. We provided participants with a list of person describing states selected from our pilot study and asked them to think of a behavior that could be indicative of this state while at the same time being attributable to a person's trait (see online supplemental materials for detailed instructions). Based on these responses, we selected a set of 288 statements that we submitted to a further pre-test. We therefore briefed an independent sample of participants (N = 171) about the trait and state concepts as stable vs. variable actor characteristics and asked them to name both a state and a trait that may come to mind when reading a behavioral statement. Again, we only selected statements to be included in Studies 3, 4, and 5 that reached a consensus score of 50% or higher for both, trait and state, with the further constraint that mean consensus scores for traits and states did not differ across all selected stimuli.

Study 1

In the first study, our goal was to establish whether people show indications of spontaneous state inferences when presented with single-implication state-implying behavioral statements. We adapted the established false recognition paradigm (Todorov & Uleman, 2002). In this paradigm, participants are first asked to memorize a series of statements describing individual actors' behaviors. In a later recognition phase they are presented with the images of the actor and indicate whether a probe word occurred in the statement or not. In the crucial conditions, these probes are adjectives implied by the presented behavior. Spontaneous trait inferences are inferred from higher false recognition rates in the implied-trait condition (e.g., erroneously responding "yes"

[1]Note that in some cases the adjectives named by participants were not identical with the pre-tested adjectives on which we had based stimulus generation. Separate analyses, however, revealed no significant differences in the pattern of results

to an adjective that was implied but not presented in a behavioral statement) as compared to a control condition. In the present study we implemented the paradigm using behavioral statements either unambiguously implying states or traits in order to (a) establish whether significant indicators of state inferences can be observed and to (b) provide a first effect size estimate for state inferences. We used a similar approach as Levordashka and Utz (2017), presenting behavioral statements to appear like ostensible posts on a social media platform.

Method

Sample size determination

Based on a minimum effect size of interest of Cohen's d_z = 0.20 for crucial one-tailed within-sample *t*-tests comparing between implied and implied-other conditions (see *Procedure*) for the trait and the state condition, both conditions required a minimum of n = 156 participants, respectively, to provide enough statistical power (1 - β = .80) with α = .05, calculated using G*Power 3.1 (Faul, Erdfelder, Lang, & Buchner, 2007). However, to avoid unnecessary spending, we relied on a sequential testing procedure (Lakens, 2014) with one interim analysis planned at 37 valid data sets or time = .26 for each condition, using a Pocock-type spending function calculated with the GroupSeq package for R (Pahl, 2018). We pre-registered to stop data collection if the observed effects were significant at the interim analysis at α_1 = .018. If they were not, we planned to continue data collection until N = 184 valid data sets would have been collected and perform the final analysis with α_2 = .032. In order to secure the required numbers of valid data sets after applying the pre-registered exclusion criteria, we overpowered both studies by 12.5%, thus planning to collect data from n = 41 and n = 207 participants per condition for the interim and final analyses, respectively.

Participants

The current study relies on valid data from a total of N = 86 participants (45 female; average age M = 28.9 years, SD = 8.7, ranging from 18 to 58). The majority of participants indicated being native speakers of German (87%), while 13% indicated to speak German as one of their native languages. Participants were recruited via Prolific (www.prolific.ac) and received monetary compensation of 1.49 GBP (approx. 1.91 USD) for the duration of ten minutes. Participants were randomly assigned to either the trait condition (n = 45) or the state condition (n = 41). Data of three initial participants who had received erroneous instructions due to a programming error were not included into analyses. Following our pre-registered exclusion criteria, we excluded three further participants because they did not pass the initial attention test and nine further participants because they aborted the experiment before debriefing and thus did not give informed consent for data analyses.

Materials

We used a set of 18 single-implication statements that implied, but not explicitly mentioned a trait, as well as 18 single-implication statements that implied, but not explicitly mentioned a state. Each statement was presented together with a person's name and a portrait picture (selected from the 10k US adult faces database; Bainbridge, Isola, & Oliva, 2013), designed such that they appeared like messages from an instant messaging service application (Levordashka & Utz, 2017, see Figure 1 for an example).

Figure 1

Example Stimulus (Studies 1 and 2). Original Stimuli were in German language.

Design

We had pre-registered this research prior to data collection as two independent studies, one on trait inferences and one on state inferences. However, because both studies were conducted simultaneously and participants were randomly assigned to either of the studies, we collapsed data and treated both studies as a between-subjects condition in the analyses. Thus, the present study followed a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) mixed design with the factor Inference varying between participants and Condition varying within participants. We implemented an additional between-subjects factor based on stimulus-set assignment: Using a counter-balanced Latin square design, stimuli were assigned to separate sets of six to be presented equally often in the implied-, implied-other, and new condition.

Procedure

Data collection was conducted using the platform Qualtrics for online data collection (www.qualtrics.com). The study started with a welcome page that contained an initial attention check requiring participants to click on a logo instead of the continue button (Oppenheimer, Meyvis, & Davidenko, 2009). If participants failed the attention check, they received a notice and were asked to re-read instructions on the welcome page. If they failed the attention check a second time, they were excluded from participation and directed back to Prolific.

Participants then started the learning phase of the false recognition paradigm (Todorov & Uleman, 2002), in which they were presented with one of two sets of stimuli (trait vs. state) containing 18 targets and nine fillers, presented in individual random order. Each individual stimulus was displayed for five seconds. Participants were instructed to read the stimuli carefully in preparation of a memory test.

Directly following the learning phase, participants completed the recognition task. In each recognition trial, an actor's name and portrait from the learning phase were presented together with a single probe adjective. Participants were instructed to click on a "yes" button when they recognized this probe word to have been presented earlier in the statement of the same actor and to click on a "no" button when they did not recognize the probe word to have been presented in the statement. The 18 portraits from the target trials of the learning phase were split into three recognition conditions, such that within each participant, six portraits each were presented with (a) the specific trait/state adjective previously implied by the respective actor's statement (implied trait/state probe), (b) with a trait/state adjective implied by a *different* actor's statement (implied-other trait/state probe), and (c) with a new trait/state adjective, which had not been implied in any of the statements (new trait/state probe). We further balanced the valence of implied-other and new control probes, so that half were of opposite valence, and half were of matching valence with regard to the implied trait or state (e.g., Schneid, Carlston, & Skowronski, 2015; Schneid, Crawford, Skowronski, Irwin, & Carlston, 2015).[2] All target probes required a "no" response. To avoid response biases, the recognition test included nine filler trials in which we presented the actors with the trait or state adjective that was explicitly mentioned in the filler statements during the learning phase, thus requiring "yes" responses. We recorded responses and response latencies. After the test phase, participants provided demographic information (age, gender, education, profession). Participants also self-reported language proficiency using a 7-point scale (1: "German is my only native language" to 7: "My German is not good enough to understand this question"). Finally, participants were fully debriefed about the purpose of the study and once again asked for consent for data storage and analyses. We had pre-registered an exploratory measure of Implicit Personality Theories (Dweck, Chiu, & Hong, 1995), which, however, was not included in this study because of a programming oversight. This measure was included in Study 2 instead.

Results

We conducted all data analyses using R 3.6.3 (R Core Team, 2020). As customary with the false recognition paradigm, we focused analyses on false response rates (Todorov & Uleman, 2003). Analyses of response latencies for correct rejections and further exploratory analyses are reported in the online supplemental materials. The pre-registered main analyses use ANOVA to analyze differences between participants' scores averaged per condition, aggregated across statements. Data were submitted to a 2 (Inference: trait vs. state) x 3 (Condition: implied vs. implied-other vs. new) ANOVA with Inference varying between participants and Condition varying within participants. Results show a significant main effect for condition, $F(2, 168) = 28.56$, $p < .001$, $\eta^2_G = .13$, 90% CI [.06, .21], and a significant main effect for inference, $F(1, 84) = 11.45$, $p = .001$, $\eta^2_G = .06$, 90% CI [.01, .17], but no significant interaction, $F(2, 168) = 0.46$, $p = .634$, $\eta^2_G < .01$, 90% CI [0, .03], see upper left panel of Figure 2. Separate analyses confirmed expected effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied trait condition ($M = .30$, $SD = .27$) than in the implied-other trait condition ($M = .21$, $SD = .22$), $t(44) = 2.246$, $p = .03$ (one-tailed), $d_z = 0.376$, 95% CI [0.03, 0.72] and in the new trait condition ($M = .11$, $SD = .15$), $t(44) = 4.818$, $p < .001$ (one-tailed), $d_z = 0.868$, 95% CI [0.45, 1.29]. More importantly, in the state inference condition, participants also showed higher false recognition rates in the implied state condition ($M = .45$, $SD = .27$), than in the implied-other state condition ($M = .33$, $SD = .22$), $t(40) = 2.508$, $p = .016$ (one-tailed), $d_z = 0.474$, 95% CI [0.07, 0.87] and in the new state condition ($M = .20$, $SD = .24$), $t(40) = 5.124$, $p < .001$ (one-tailed), $d_z = 0.97$, 95% CI [0.51, 1.43].

Discussion

The results of this first study provided initial evidence that people *can* spontaneously draw both trait and state inferences when processing behavioral information with clear trait- or state-implications: In our adaption of the false recognition paradigm, participants were more likely to falsely recognize state and trait adjectives that were implied by behavioral descriptions as compared to non-implied state or trait words. Furthermore, effect sizes were in a comparable range for both the state inference effects as well as the trait inference effects.

The interpretation of our results is, however, limited by the following caveat: When conducting the interim analyses according to the sequential testing procedure, we had erroneously accepted the difference between implied and implied-other conditions in the trait-inference condition as

[2] Exploratory analyses demonstrated that valence (in)congruence of implied-other and control traits did not significantly qualify the reported results. We therefore collapsed analyses across this factor.

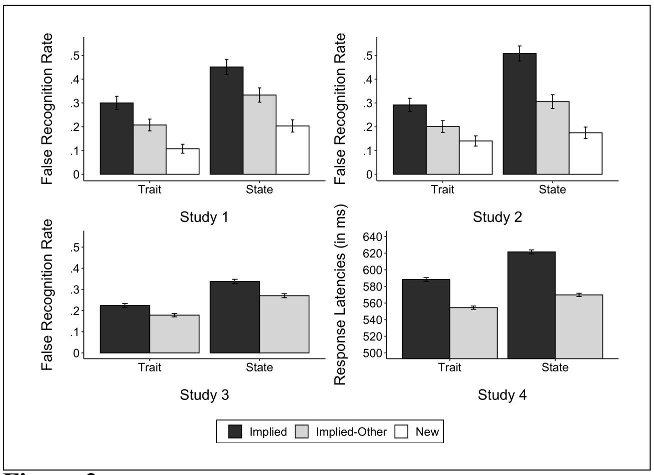

Figure 2

Mean proportion of false recognition responses (Studies 1-3) and mean response latencies for correct responses (Study 4) as a function of type of inference and experimental condition. Error bars represent SE.

significant albeit the *p*-value of $p = .03$ did not fulfill the preregistered significance criterion of $p \leq .018$. We thus prematurely stopped data collection. However, given that the pattern of results replicated typical STI findings, we opted against resuming data collection when noticing our mistake, and decided instead to invest our resources into a replication study conducted in the laboratory.

Study 2

Study 2 served as close replication of Study 1 with the only difference being that data collection was conducted in the laboratory instead of online.

Method

Sample size determination

Based on the effect sizes resulting from Study 1, we aimed at providing enough statistical power ($1 - \beta = .80$) to detect effect sizes of $d_z = 0.30$ with $\alpha = .05$ (one-tailed) for both, the state and trait conditions of this study. This would require a sample size of $n = 142$ per condition, calculated using G*Power 3.1 (Faul et al., 2007). Again, we relied on a sequential testing procedure (Lakens, 2014) with one interim analysis planned at 68 valid data sets or time = .38. We pre-registered to stop data collection if the observed effects were significant at the interim analysis at $\alpha_1 = .025$. If they were not, we planned to continue data collection until 180 valid data sets would have been collected and perform the final analysis with $\alpha_2 = .025$. In order to obtain the required numbers of valid data sets after applying the pre-registered exclusion criteria, we overpowered both studies by 12.5%, resulting in $n = 71$ and $n = 189$ for the interim and final analyses, respectively, for each condition. During data collection it became apparent that an unexpectedly high number of participants appeared to be non-native speakers of German—whose data would eventually need to be excluded from analyses (see sample description). Additionally, the randomized assignment of participants lead to a high imbalance of participants in the trait and state condition (with only $n = 28$ in the state and $n = 40$ in the trait condition at interim analyses). We therefore deviated from the pre-registered sample size and collected data from 117 participants in order to achieve a more balanced assignment of participants with sufficient language proficiency to both conditions. Note that interim analyses at the pre-registered $n = 68$ already fulfilled the aforementioned decision criteria to warrant applying the stopping rule (see online supplemental materials).

Participants

The current study relies on valid data from a total of $N = 91$ participants (34 female, average age $M = 29.9$, $SD = 11.2$, ranging from 18 to 71 years). The majority of participants indicated being native speakers of German (65%), 21% indicated to speak German as one of their native languages, and 14% indicated that they spoke German very well, albeit it was not their native language. Participants were mainly students from various faculties of a university in northern Germany, recruited via a university online job platform and were compensated 2.50 EUR (approx. 3.13 USD) for the duration of 15 minutes. Our study was the first to be conducted in an one-hour lab session followed by an unrelated study on face recognition. Participants were randomly assigned to either the trait condition ($n = 42$) or the state condition ($n = 44$). Following our pre-registered exclusion criteria, we excluded data from further participants: Six because they did not pass the initial attention test, twelve because they aborted the experiment before debriefing and thus did not give informed consent for data analysis, and 18 because they self-reported insufficient language proficiency.

Procedure

Study 2 used the same materials and followed the same procedure for the false recognition task as Study 1, with the exception that it was conducted in the laboratory. After completion of the false recognition paradigm, participants additionally completed a Navon task (Navon, 1977) and a measure of implicit personality theories (Implicit Personality Theories Questionnaire [8 items], translated into German; Dweck et al., 1995). A description and results of the exploratory analyses are reported in the online supplemental materials.

Results

Individual false recognition rates were submitted to a 2 (inference: trait vs. state) x 3 (condition: implied vs. implied-other vs. new) ANOVA with the factor inference varying between participants and condition varying within participants. Results show significant main effects for condition, $F(2, 168) = 40.75$, $p < .001$, $\eta^2_G = .18$, 90% CI [0.10, 0.27] and for inference, $F(1, 84) = 12.56$, $p < .001$, $\eta^2_G = .07$, 90% CI [.01, .18]. We also observed a significant interaction between inference and condition, $F(2, 168) = 5.84$, $p = .004$, $\eta^2_G = .03$, 90% CI [.00, .08], see upper right panel of Figure 2. Separate analyses confirmed expected effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied condition ($M = .29$, $SD = .24$) than in the implied-other trait condition ($M = .20$, $SD = .21$), $t(43) = 2.779$, $p = .008$ (one-tailed), $d_z = 0.398$, 95% CI [0.1, 0.69], and in the new trait condition ($M = .14$, $SD = .18$), $t(43) = 4.224$, $p < .001$ (one-tailed), $d_z = 0.697$, 95% CI [0.33, 1.06]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition ($M = .51$, $SD = .22$) than in the implied-other state condition ($M = .31$, $SD = .21$), $t(41) = 4.336$, $p < .001$ (one-tailed), $d_z = 0.937$, 95% CI [0.42, 1.45], and in the new state condition ($M = .18$, $SD = .20$), $t(41) = 8.061$, $p < .001$ (one-tailed), $d_z = 1.601$, 95% CI [1, 2.2].

Further analyses exploring the significant interaction effect of inference type by condition indicated that the mean difference of implied state and implied-other state ($M = .202$, $SD = .303$) was larger than the mean difference of implied trait and implied-other trait, albeit not significantly ($M = .091$, $SD = .217$), $t(74.123) = 1.996$, $p = .054$, $d_s = 0.425$, 95% CI [-0.01, 0.86].

Discussion

In Study 2 we observed a highly similar pattern of results as in Study 1, thus strengthening our previous conclusion that participants can spontaneously draw state inferences as well as trait inferences.

While Studies 1 and 2 serve as important proofs of concept and starting point for our research, there are, however, two characteristics of these first two studies that limit the interpretability of the results. On the one hand, we had implemented the state and trait conditions as a between-participants factor, which may have affected participants' general mode of information processing during the learning phase, increasing the general likelihood of inferring states or trait when only presented with strongly state-implying vs. trait-implying behaviors.

Furthermore, we employed only single-implication behavioral descriptions that exclusively and unambiguously implied *either* trait *or* state inferences. Past research on spontaneous trait inferences has frequently employed stimulus materials created such that they strongly afford trait inferences by describing rather extreme behaviors (e.g., "The farmer paints a swastika on the synagogue wall"; Winter & Uleman, 1984), by including temporal markers indicating repeated behaviors and thus temporal consistency ("I attend my church *twice a week* ..."; Carlston & Skowronski, 1994 [emphases added]), or by including expressions of inner dispositions such as attitudes (e.g.,"I *hate* animals. Today ... I saw this puppy. So I kicked it out of my way."; Carlston & Skowronski, 1994 [emphases added]). While the use of materials that strongly imply a specific inference is required to investigate process characteristics of any given inference, it provides strong stimulus-constraints and thus limits generalizability of results. Results of Studies 1 and 2 suggest that perceivers spontaneously draw trait inferences even from less extreme, relatively mundane behaviors, in the absence of strong trait-affording wording or temporal markers (see also Levordashka & Utz, 2017). Because we developed single-implication statements specifically for the investigation of spontaneous state versus trait inferences, these statements may possess certain characteristics that afford one type of inference only, while inhibiting other inferences. For example, if someone states that *"There is nothing going on this weekend, I am wasting my time channel-surfing"*, this person is most obviously bored by the specific situation which is hardly attributable to a personal disposition - at least when no further information is provided. We thus decided to develop a new set of materials with dual-implication statements, that – in principle – allow for the simultaneous occurrence of state and trait inferences. For example, in our introductory example of Daniel laughing at the joke, one may assume that he is a *jolly* person who generally laughs a lot or that he is so *amused* by this one joke that he bursts out laughing. If participants draw state inferences even if a trait inference would be just as warranted, we can conclude that these inferences need not be mutually exclusive. We use a false recognition paradigm (Study 3) and a probe recognition paradigm (Study 4). While Studies 1 and 2 were conducted in German language (with native German speakers as participants), Studies 3 and 4 were conducted in English language (with native English speakers as participants).

Study 3

With Study 3, we aimed at providing an extended replication of our previous results by using novel behavioral descriptions during the learning phase of the false recognition paradigm that allowed for both, trait and state inferences to occur. Furthermore, all participants were probed for both, trait and state inferences during the recognition phase, thus avoiding potential systematic differences in processing mode that may have affected results based on the between-participants design of Studies 1 and 2.

Method

Sample size determination

Our primary interest in this study was the main effect of Condition (implied vs. implied-other) for both, trait and state inferences. In order to provide sufficient statistical power (1 - β = .80) to detect a main effect size of $\eta_p^2 = .091$ (estimation based on the smallest effect size observed in our previous studies) with α = .05 in a 2 x 2 repeated-measures ANOVA, valid data of *N* = 82 participants were required. Because we aimed to additionally test for a potential interaction effect – thus directly comparing state and trait inference effects – we included it into our power analyses. For this interaction, we considered a small effect size of $\eta_p^2 = .022$ as the smallest effect size of interest. In order to provide sufficient statistical power (1 - β = .80) with α = .05 for this effect size in a 2 x 2 interaction in a repeated-measures ANOVA, we planned to collect 352 valid data sets. Given an estimated exclusion rate of 5% with online data collection, we collected data from *N* = 376 participants. Power analyses were conducted using MorePower 6.0.4 (Campbell & Thompson, 2012).

Participants

The current study relies on valid data from a total of *N* = 365 participants (109 female, average age *M* = 35.6, *SD* = 12.3, ranging from 18 to 75 years). The majority of participants indicated being native speakers of English (96%), 3% indicated to speak English as one of their native languages, and 1 participant indicated that they spoke English very well, albeit it was not their native language. Participants were recruited using Prolific (www.prolific.ac) and received monetary compensation of 1.50 GBP (approx. 1.88 USD) for the average study duration of nine minutes. Following our pre-registered exclusion criteria, we excluded data of six further participants because they had failed an attention test and two further participants because they self-reported insufficient language proficiency. Data of 90 further participants were excluded because they aborted the experiment before debriefing and thus did not provide informed consent for data analyses.

Materials

We selected 24 dual-implication statements as target stimuli. These were intermixed with 12 filler statements of a similar structure to the target stimuli, with the exception that the fillers explicitly mentioned a trait or state word. Contrary to the previous studies, stimuli were not designed to look like messages from an instant messaging service application but formulated in third person and paired with a portrait and a name.

Design

We employed a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject design with rate of false recognitions (answers 'yes' in test phase) as dependent variable. In Study 3, we omitted the new trait and state conditions, because the implied-other condition represents the more conservative test for the occurrence of trait and state inferences. We implemented an additional between-subjects factor based on stimulus-set assignment: Stimuli were randomly assigned to separate sets of six to be presented equally often in the implied and implied-other condition across participants (using a Latin square design), in individual random order.

Procedure

Study 3 followed the same procedure as Study 1, with the exception that participants self-reported language proficiency using a 6-point scale (1: "English is my first (native) language" to 6: "It is very hard for me to speak and understand English").

Results

False recognition rates were submitted to a 2 (Inference: trait vs. state) x 2 (Condition: implied vs. implied-other) within-subject ANOVA. Results show significant main effects for inference, $F(1, 364) = 135.79$, $p < .001$, $\eta_G^2 = .05$, 90% CI [.02, .09], and for condition, $F(1, 364) = 43.89$, $p < .001$, $\eta_G^2 = .02$, 90% CI [0, .04], with no significant interaction, $F(1, 364) = 1.64$, $p = .202$, $\eta_p^G < .01$, 90% CI [0, .01], see lower left panel of Figure 2. Separate analyses confirmed spontaneous person inference effects in both inference conditions: In the trait inference condition, participants showed higher false recognition rates in the implied trait condition (*M* = .22, *SD* = .21) than in the implied-other trait condition (*M* = .18, *SD* = .20), $t(364) = 4.009$, $p < .001$ (one-tailed), $d_z = 0.22$, 95% CI [0.11, 0.33]. In the state inference condition, participants also showed higher false recognition rates in the implied state condition (*M* = .34, *SD* = .24) than in the implied-other state condition (*M* = .27, *SD* = .22), $t(364) = 5.292$, $p < .001$ (one-tailed), $d_z = 0.29$, 95% CI [0.18, 0.4]. Albeit effect sizes were somewhat larger in the state condition as compared to the trait condition, the inference effects (calculated as difference scores of implied – implied-other) did not differ significantly between conditions, $t(728) = 1.28$, $p = .201$, $d = 0.094$, 95% CI [-0.05, 0.24]. We report further auxiliary analyses in the Supplemental Materials.

Discussion

Results of Study 3 replicate and extend the findings of Studies 1 and 2 by providing first evidence that participants can draw both trait and state inferences from one and the same ambiguous behavior description. These results can, however, not yet be interpreted as indicators that trait and state inferences are *simultaneously* drawn when processing behavioral information (similar to the assumption of simultaneous trait and situation inferences in Krull & Ericson's (1995) model). Indeed, given our specific implementation of the false recognition paradigm, the same results would be observed if some behavioral descriptions exclusively or dominantly triggered trait inferences and other stimuli exclusively or dominantly triggered state inferences. Alternatively, some participants may have systematically only drawn trait inferences and no state inferences and other participants may have systematically drawn state inferences and no trait inferences. We conducted auxiliary correlational analyses that indicate that trait and state inferences seem not to have occurred at each others' expense, because we did not observe any negative correlation of state and trait inference effects, neither on a by-participant, nor on a by-item level (see online supplemental materials). Nevertheless, it is desirable to provide a more direct empirical test of the simultaneity and mutual non-exclusiveness of trait and state inferences.

The current results do also not fully refute the assumption that trait inferences may be the dominant inference from behavior. For example, one may assume that participants spontaneously only draw trait inferences at encoding of the behavioral information but later (re)consider state inferences once probed with a fitting state word during the recognition test. However, note that the same argument applies to the opposing assumption that participants spontaneously only draw state inferences at encoding and later (re)consider trait inferences when probed with a fitting trait word (but see Todorov and Uleman, 2002). We conducted auxiliary analyses of response times to explore whether response times of false recognition responses differed between the implied state and trait conditions, presuming that false recognition response based on inferences drawn at encoding may be faster as compared to responses after retrospect reconsiderations. However, there were no significant differences in false recognition response latencies between the implied-state as compared to the implied-trait condition (see online supplemental materials).

In Study 4, we used a more direct approach to investigate if and to what extent trait and state inferences are drawn *simultaneously* when participants process behavioral information about actors.

Study 4

In Study 4, we implemented a probe recognition paradigm (McKoon & Ratcliff, 1986) – another indirect paradigm frequently used in research on spontaneous trait and situation inferences (e.g., Ham & Vonk, 2003; Newman, 1991; Ramos, Garcia-Marques, Hamilton, Ferreira, & Van Acker, 2012). In this paradigm, participants read individual behavioral statements, each *immediately* followed by several probe words. Participants decide as quickly and accurately as possible whether the probes were part of the statement or not. In typical studies on STIs, these probe words are trait adjectives implied by the behavior or non-implied control probes, thus both requiring a negative response. If encoding of the behavioral statement automatically triggers a trait inference, participants should have more difficulty rejecting the corresponding trait probe and thus demonstrate slower response latencies and/or higher error rates as compared to non-implied control trait probes. One advantage of this paradigm is that each behavioral statement can be followed by several probes, thus allowing to test for multiple inferences referring to the same statement (e.g., Todd et al., 2011). In our adaptation of the paradigm, behavioral statements were followed by both, implied state and implied trait probes, which allows investigating the simultaneous occurrence of both types of inferences.

Method

Sample size determination

We followed the same pre-registered sample size rationale as in Study 3, thus planning to collect valid data from 352 participants.

Participants

The current study relies on valid data from a total of N = 341 participants (211 female, four other, one unspecified, average age M = 33.2, SD = 11.5, ranging from 18 to 71 years). The majority of participants indicated being native speakers of English (96%), 3% indicated to speak English as one of their native languages, and one participant indicated that they spoke English very well, albeit it was not their native language. Following our pre-registered exclusion criteria, we excluded the data of six further participants because they aborted the experiment before debriefing and thus did not give informed consent for data analysis, five because they self-reported insufficient language proficiency, and two because they responded accurately in less than 60% of trials. Participants were recruited via Prolific (www.prolific.ac) and received 2.20 GBP (approximately 2.75 USD) for the duration of 15 minutes.

Materials

We used a different subset of 24 behavioral statements developed for Study 3 as target statements and 24 filler statements explicitly mentioning a state or a trait word (see online supplemental materials for the complete list of statements and pre-test results). Given the framing of the task as measuring automatic text comprehension, behavioral statements included the actor name but were presented on screen without images.

Design

We employed a 2 (inference: trait vs. state) x 2 (condition: implied vs. implied-other) within-subject design with response latencies of correct probe rejections of target trials as dependent variable.

Procedure

Participants completed the same introduction procedure and attention check as in the previous studies. The probe recognition paradigm was then introduced to participants as a study on language comprehension. The study was designed in PsychoPy 3.1.2 (Peirce et al., 2019) and conducted online on Pavlovia (www.pavlovia.org). Participants were presented with 48 behavioral statements (24 targets, 24 filler) in individual random order and were instructed to read them carefully. Each statement was presented for 3000ms. Immediately after each statement, participants completed eight probe recognition trials, in which they indicated for each probe word whether it had been part of the previous statement or not. Each probe was preceded by a blank screen (250ms) and a fixation cross (500ms) and remained on screen until a response was recorded. Participants were instructed to indicate via key press whether the word had appeared (press [D]) or not (press [K]) in the statement. Participants were instructed to respond as accurately and as fast as possible. Erroneous responses were signaled to the participants by a red cross displayed for 1000ms. After completion of all eight probe trials for each statement, the next statement was presented with an intertrial interval of 500ms.

Target probe trials always consisted of the trait and the state adjectives implied by the behavioral statement of the same person (implied condition), as well as a trait and a state adjective implied by the behavioral statement of another person (implied-other condition), thus all requiring a "no" response. In order to balance the ratio of correct "yes" and "no" responses for each trial, we additionally presented four filler probes consisting of words that had actually appeared in the statement (i.e., names, objects, verbs, and prepositions). In order to avoid that participants recognized that any type of probe consistently required a "yes" or "no" response, probes for the filler statements were chosen such that correct responses were "yes" for adjectives and "no" for names, objects, verbs, and prepositions. For each trial, the order of the eight probes was individually randomized, with the restriction that the first probe was never a target probe (Stewart, Weeks, & Lupfer, 2004). Responses and response latencies were recorded. After completion of the probe recognition task, participants provided demographic information (age, gender, language proficiency, education, profession). Additionally, participants were asked to indicate how seriously they complied with task instructions using a 10-point scale (0 = not at all and 10 = very much, $M = 9.5$, $SD = 0.8$) and asked to speculate about the hypothesis of the study. At the end of the study, participants were fully debriefed about the purpose of the study and once again asked for consent for data storage and analyses.

Results

We had pre-registered response latencies of correct rejections as main dependent variable for our analyses. Analyses of response latencies usually require corrections of outlying slow responses (Ratcliff, 1993). To our knowledge, there is no convention how to correct for outliers in the probe recognition paradigm. Therefore, we applied different trimming criteria for slow responses (2500, 2000, 1500ms, individual $M \pm 2.5SD$), and transformations (log- and inverse transformation) and compared their impact on analyses. Across the different trimming criteria and transformations, effect sizes differed by small to medium amounts (trait inference effect: $d_z = 0.580$ - 0.806; state inference effect: $d_z = 0.790$ - 1.117; Inference x Condition interaction: $\eta_p^2 = .03$ - $.07$; see online supplemental materials). Analyses reported in the text are based on log-transformed data with a cut-off of 1500 ms. For ease of understanding, descriptive statistics are based on untransformed but trimmed response latencies.

The pre-registered 2 (inference: trait vs. state) x 2 (condition: implied vs. implied-other) repeated-measures ANOVA of response latencies for correct rejections of target probes documented significant main effects of inference, $F(1, 340) = 224.09$, $p < .001$, $\eta_G^2 = .01$, 90% CI [0, .04], and condition, $F(1, 340) = 592.05$, $p < .001$, $\eta_G^2 = .03$, 90% CI [.01, .07], that were qualified by a significant interaction, $F(1, 340) = 23.55$, $p < .001$, $\eta_G^2 = .001$, 90% CI [0, .02], see lower right panel of Figure 2. Separate analyses confirmed expected effects in both inference conditions: Participants were slower to reject the implied trait probes ($M = 592$ ms, $SD = 122$) compared to the implied-other trait probes ($M = 557$ ms, $SD = 105$), $t(340) = 13.928$, $p < .001$ (one-tailed), $d_z = 0.286$, 95% CI [0.25, 0.33]. Similarly, the participants were slower to reject the implied state probes ($M = 625$ ms, $SD = 115$) compared to the implied-other state probes ($M = 573$ ms, $SD = 115$), $t(340) = 19.502$, $p < .001$ (one-tailed), $d_z = 0.425$, 95% CI [0.38, 0.47]. The mean difference of implied state and implied-other state probes ($M = 51$ ms, $SD = 49$) was significantly larger than the mean difference of implied trait and implied-other trait probes ($M = 34$ ms, $SD = 46$), $t(680) = 4.794$, $p < .001$, $d_z = 0.367$, 95% CI [0.22, 0.52].

Discussion

The results of Study 4 conceptually replicate and complement the results of our previous studies in a different experimental paradigm. Again, we observed significant state and trait inference effects, with state effects being significantly larger than trait effects. Most importantly, the use of the probe recognition paradigm in Study 4 allows the conclusion that observers drew state and trait inferences both spontaneously and simultaneously when encoding the behavioral information. As in Study 3, our auxiliary correlational analyses (see online supplemental materials) showed no significant interrelations of trait- and state-inference effects, neither on the participant level nor on the by-item level, supporting the conclusion that trait and state inferences did not occur at each others' expense.

Our research thus far strongly supports the notion that people draw spontaneous trait and state inferences when forming impressions from others' behaviors. However, the interpretability of these effects remains limited because the differentiation of trait and states in these studies entirely relies on our pilot study. In this study, we had asked participants to deliberately judge the single adjectives with regard to their perceived temporal and situational stability after explicitly instructing them about our theoretical conceptualization of states and traits. We cannot be entirely sure that participants draw the same conceptual distinction between traits and states when inferring them from specific behavioral statements. Albeit our pilot study clearly indicated that the employed adjectives by themselves were understood as either states or traits, several of these adjectives may still appear as possibly referring to both, states and traits - depending on the contexts in which they are used. For example, we may use the state adjective *sad* when considering a person to *feel sad* (state) or to be a *sad person* (trait). Because German and English have no clear linguistic markers that conceptually distinguish between traits and states (as in Spanish or Portuguese, for example), it is conceivable that the state inferences we observed in our studies are nothing but trait inferences in disguise. So far, we cannot conclude from our data that state inferences actually differ functionally from trait inferences in the eyes of perceivers. We thus conducted a fifth study in order to investigate whether participants actually represent our pre-defined state and trait inferences as functionally different from each other.

Study 5

From a theoretical standpoint, traits and states can be distinguished from one another quite precisely (e.g., Hamaker et al., 2007; but see Fleeson & Jayawickreme, 2015). The most obvious difference is their relative stability: Traits describe dispositions that are relatively stable over time and across situations; states describe dispositions that are considerably less stable over time and across situations. This stability advantage of traits over states renders trait attributions appealing for impression formation because by increasing the predictability of other peoples' future behavior they may reduce uncertainty to a higher degree than state attributions. Trait inferences should thus influence predictions about future behavior, and have been demonstrated to do so (McCarthy & Skowronski, 2011b; Nussbaum et al., 2003). State inferences, on the other hand, describe variable, more fleeting properties of a person and should therefore have less influence on future behavior predictions. If people mentally represent this functional differentiation of trait and state inferences, they should rely on them differently when using observed behavior to make predictions about actors' future behavior. In Study 5, we investigated whether the traits and states that participants spontaneously inferred from the dual-implication behavioral statements used in Studies 3 and 4 do indeed hold differential value for predicting behavior. We therefore presented participants with the same behavioral statements, either paired with the implied state or the implied trait word, and asked them to judge how likely the actors would show such behavior again in the future. We chose this procedure (instead of presenting only the behavioral statements or only the probe actives) based on the following rationale: On the one hand, it would make little sense to only present the state and trait adjectives without the behavioral statements, because the meaning of the adjectives may vary with different behavioral contexts (e.g., *curious* means something different when inferred from the statement "My neighbour Betty is peeping through the curtain slit" or from "My neighbour Betty wants to learn everything there is to know about the Mars mission"; see also Kunda, Sinclair, & Griffin, 1997). Thus, while our pretests already suggested that the selected state and trait words generally differ in their perceived temporal stability, we needed to verify that this also applied to the specific behavioral contexts of our studies. On the other hand, it would make no sense to present only the behavioral statements without adding the trait vs. state inferences, because results from Studies 3 and 4 demonstrate that participants draw both state and trait inferences spontaneously and simultaneously. Thus, presenting only the behavioral statements would make it impossible to judge if and to what extent individual behavioral predictions were influenced by state vs. trait inference. By explicitly mentioning the implied state vs. the implied trait, the respective inference is made more salient to the participant. If these inferences have a functional value for further information processes, the more salient inference should influence judgments of future behavior more than the non-salient inference. We thus expected to observe that for one and the same behavior, the salient trait inference would lead to a higher perceived probability of repeated behavior, than the salient state inference.

Method

Sample size determination

Our primary interest in this study was the effect of Adjective (trait vs. state) in a *t*-test with prediction rating as the dependent variable. We aimed to provide sufficient statistical power ($1 - \beta = .80$) to detect the minimal effect size of interest of d_s = .50 with α = .05 for the between-subjects comparison of only the first response (see Analysis Plan). Thus, valid data of N = 102 participants was required. Power analyses were conducted using G*Power 3.1 (Faul et al., 2007).

Participants

The current analyses relied on valid data from a total of N = 97 participants (64 female, one other, average age M = 33.7, SD = 11.9, ranging from 18 to 66 years). Following our pre-registered exclusion criteria, we excluded the data of five further participants because they aborted the study before completion due to a programming error. Participants were recruited via Prolific (www.prolific.ac) and received 1.56 GBP (approximately 2 USD) for the duration of 10 minutes.

Materials

We used all 41 behavioral statements used in Studies 3 and 4 as target statements, paired with a statement describing the actor with the implied state or trait adjectives, respectively. Behavioral statements were presented on screen without images.

Procedure

Participants were introduced to a study on behavior prediction and memory performance. The study was designed and conducted in Qualtrics software. Participants were instructed to read the behavioral statements and judge how likely each actor would perform behavior such as the one described again in the future using a slider ranging from 0 to 100. Additionally, they were instructed to memorize actor name, behavior, and adjective for a later memory test. We implemented the additional memory test in order to ensure that participants actually process the trait and state adjectives and do not base their judgements on the relative frequency of the described behaviors only. Participants completed the same attention check procedure with the first instruction page as in the previous studies. They were then presented with 82 trials including each behavioral statement, once paired with the implied state, once paired with the implied trait. Participants submitted their slider responses by pressing the space bar, which allowed us to record responses and response latencies.

In order to verify that participants had actually processed behaviors and state and trait adjectives, they completed ten recognition trials consisting of an actor's name and one of the adjectives from the rating phase. In five of these recognition trials, the adjectives had been presented with that same actor, and in the other five, the adjectives had been presented with a different actor. Participants were asked to judge whether actor and adjective had been presented together before (correct decision rate: M = .67, SD = .18).

Finally, participants provided demographic information (age, gender, language proficiency, education, profession, ethnicity). Additionally, participants were asked to indicate how seriously they complied with task instructions using a 10-point scale (0 = not at all and 10 = very much, M = 8.9, SD = 1.4).

Analysis Plan

We had planned two ways of analysing the results of this Study. First, we planned a simple within-participants t-test in order to inspect whether participants attributed higher likelihood of behavior repetition when the behavioral statements were paired with the implied trait as compared to the implied state adjective. However, because this within-participants design asks participants to estimate the repetition likelihood for the same behavior twice, differences between implied state and trait adjectives may be under-estimated: Regardless of randomized sequence, participants may be inclined to base their second estimation on the response given to the same behavior earlier. To account for this possibility and rule out any form of cross-contamination, we planned a second analysis using only participants' first responses to each statement, discarding their second judgement.

Results

The within-participants analyses show that participants rated the probability for behavior repetition higher when the behavioral statement was paired with a trait adjective (M = 73.46 , SD = 11.4) than with a state adjective (M = 69.06 , SD = 10.69), $t(96)$ = 4.95, p < .001 (one-tailed), d_z = 0.398, 95% CI [0.233, 0.563]. The planned additional analysis of the first judgements on each behavioral statement revealed a similar effect with a slightly larger effect size: Participants who saw the statement paired with a trait (M = 73.21, SD = 11.68) rated the probability of behavior repetition significantly higher than participants who saw the statements paired with the respective states (M = 68.44, SD = 10.48), $t(96)$ = 4.958, p < .001 (one-tailed), d_z = 0.427, 95% CI [0.25, 0.605]. We observed the same pattern of results for the second presentation of statements with a slightly reduced effect size. Participants who saw the statement paired with a trait (M = 73.68, SD = 12.11) rated the probability of behavior repetition significantly higher than participants who saw the statements paired with the respective states (M = 69.64, SD = 12.17), $t(96)$ = 3.696, p < .001 (one-tailed), d_z = 0.332, 95% CI [0.15, 0.514].

Auxiliary Analyses

We further conducted exploratory analyses in order to explore whether the relative strength of state and trait inferences in Studies 3 and 4 were related to the aggregated repetition predictions in Study 5. The rationale behind this analysis is the following: By-item analyses of data from Studies 3 and 4 suggest that not all state and trait inferences were equally strong for each behavioral statement (see online supplemental materials). If trait and state inferences are represented as functionally different in the perceivers' minds and trait inferences are more important for future predictions than state inferences, we would expect differential predictive values of the inference effects: That is, the magnitude of the trait inference should be positively related to the repetition likelihood ratings, whereas the magnitude of state inference should not matter for these predictions. We thus computed state and trait inference scores for each statement aggregated across participants as a simple difference between implied and implied-other responses. Because both studies used different dependent variables, we z-transformed these inference scores and averaged the z-values for those statements that were used in both studies. Simple bivariate correlations indicate that the size of spontaneous trait inferences in Studies 3 and 4 was related to the perceived likelihood of behavior repetition in Study 5, $r = .439$, $t(46) = 3.316$, $p = .002$, 95% CI [.177, .643], whereas the size of the state inference effects in Studies 3 and 4 was not significantly related to the repetition predictions in Study 5, $r = .078$, $t(46) = 0.533$, $p = .596$, 95% CI [-.210, .355]. Of course, these auxiliary analyses have to be interpreted with caution, because of the limited test power given the sample size of only $n = 47$ statements.

Discussion

Results of Study 5 support the assumption that participants processed trait and state inferences as indicating differential predictive value. Participants judged the likelihood that actors would repeat their behaviors in the future as higher when the behaviors were paired with the implied trait inference than the implied state inference. Albeit limited by their relatively low power, the exploratory by-item analyses further support this interpretation: The stronger the trait inference from a behavior, the more likely it seems that this behavior may be repeated in the future. No such relation was observed for state inferences drawn from the same behaviors. Note that this should not be read as causal interpretation; one could similarly conclude that the more a behavior elicited an expectation to be repeated in the future, the more likely participants draw a trait inference, whereas state inferences appeared independent of the perceived repetition likelihood. Nevertheless, the Results of Study 5 support the assumption that the trait and state inferences measured in Studies 3 and 4 do indeed represent functionally distinct inferences.

General Discussion

Previous research on impression formation from behavior has provided a vast amount of evidence that when observing others' behaviors, people spontaneously draw trait inferences (e.g., Moskowitz & Olcaysoy Okten, 2016; Uleman et al., 2012). In a series of five pre-registered studies we provide consistent evidence that people spontaneously and simultaneously infer both traits and states from behavior.

In Studies 1 and 2, we employed behavioral statements with single-implications of either a trait or a state in a false recognition paradigm. We observed significant inference effects for both trait and state adjectives, in that participants more frequently falsely recognized implied trait and state probes as having been previously mentioned in an actor's behavioral statement, as compared to a trait or state implied by a different actor's behavior, or a new trait or state adjective. We thus replicated the established spontaneous trait inference effect and, more importantly, provided first evidence that people can spontaneously infer states as well.

In Studies 3 and 4, we used dual-implication stimuli, that is, behavioral statements that allowed for both trait and state inferences. Again, we observed significant inference effects for both trait and state adjectives in a false recognition paradigm (Study 3) and in a probe recognition paradigm (Study 4).

Results of Study 5 support the assumption that trait and state inferences are represented as functionally different: Participants rated actors as more likely to show similar behavior again in the future when the behavior was paired with the implied trait adjective, as compared to the implied state adjective.

On the spontaneity of spontaneous state and trait inferences

The results of our research are in line with the many findings from different domains of person perception showing that people have a strong tendency to spontaneously make sense of their social environment by going beyond the given information, adding self-generated information and thus drawing causal inferences and attributions.

Our studies show that person inferences from others' behaviors include both, considerations of the actors' enduring traits as well as considerations of their more transient states. Note that the semantic and pragmatic rules of the employed experimental paradigms and their instructions actually call for noninferences: Participants were never asked to form impressions of others – neither state nor trait impressions – but to merely process and memorize statements and images. Deliberate impression formation would have been a rather distracting activity during information encoding and also hindering during task performance, given that these inferences increase false recognition rates and slow down responding. Thus, the indirectness of the inference effects – the increase of false recognition responses to implied states and traits in the false recognition paradigm and the slowing of correct rejection responses in the probe recognition paradigm – support the assumption that both state and trait inference effects are *spontaneous* both in the sense of *unintentional* as well as in the sense of *uncontrolled* (Uleman et al., 1996). Future research may systematically investigate further process characteristics of state and trait inferences, such as the relative efficiency and resource (in)dependence, general controllability, and subjective awareness of these inference processes and/or their outcomes.

Using actors' portraits for the assessment of inference effects in the recognition test of the false recognition paradigm further allows for the conclusion that the observed trait and state inference effects should not be interpreted as mere side-effects of text-comprehension, but as actor-specific person inferences. Furthermore, we always compared responses to the state and trait adjectives implied by the actors' behavior to trait and state adjectives implied by behaviors of other actors during the same learning phase. It is thus not mere familiarity or traces of prior activation of text-based associations of behaviors with trait or state words that increase false recognition rates (e.g., reading the verb *laughing* activating associations like *jolly* or *amused*), but inferences directly tied to the specific actors of that behavior (e.g., Laughing *Dave is* jolly and/or amused; see also Orghian, Ramos, Reis, & Garcia-Marques, 2018).

These considerations have been discussed previously and at length with regard to spontaneous trait inferences (e.g., Todorov & Uleman, 2002, 2003) and our research demonstrates that they apply to state inferences as well. Future research may more systematically investigate to what extent state and trait inferences are based on associative versus inferential processing (e.g., Carlston & Skowronski, 2005).

However, the more important contribution of our findings is that they show that state and trait inferences occur *simultaneously* at encoding of the behavioral information and are *mutually non-exclusive*. Specifically, in Studies 3-5, we employed dual-implication stimuli, that were carefully pretested to support both trait and state inferences in deliberate impression formation - and lead to both trait and state inference effects in both the false recognition and the probe recognition paradigm assessing effects of spontaneous impression in Study 3 and 4. To better understand the results of the false recognition effects in Study 3, we had conducted auxiliary correlational analyses both on the by-participant and the by-item level (see online supplemental materials). Observing negative correlations on either level would have implied that trait and state inferences may be mutually exclusive; on the participant level, if participants who show strong trait inference effects show weak state inference effects or vice versa; on the item-level, if statements leading to strong trait-inference effects (aggregated across participants) lead to weak state inference effects or vice versa. The results indicate, however, that state and trait inferences were not drawn at the expense of each other: There was no indication that participants within our samples had individual response tendencies towards either spontaneous trait or spontaneous state inferences, nor that individual behavioral statements would only prompt either spontaneous state or spontaneous trait inferences. The strongest evidence for the mutual non-exclusiveness of simultaneously and spontaneously drawn state and trait inferences, however, stems from the probe recognition paradigm used in Study 4, in which we prompted participants with both implied trait and state adjectives immediately after encoding of the behavioral statements. Again, we observed both inferences effects and again we did not find any negative correlations between state and trait inferences neither on the participant nor on the item level. Thus, we feel confident to conclude that spontaneous person inferences simultaneously include considerations about stable person dispositions as well as transient mental states and that these are mutually non-exclusive at the stage of spontaneous information processing.

Multiple spontaneous social inferences in the impression formation process

Our findings complement empirical evidence that people can draw spontaneous situation inferences from behavior (Ham & Vonk, 2003; Todd et al., 2011). However, while states certainly have a situational component, they differ from situation inferences quite substantially. In the example *Pete tells the waiter the food tastes good* (adapted from Ham & Vonk, 2003), a situation inference would refer to a property of the situation: the food is *delicious*. In contrast, a state inference would refer to the person: Pete is *pleased*. Our results are also compatible with recent documentations of other spontaneous person inferences, such as goals (e.g., Hassin et al., 2005; Olcaysoy Okten & Moskowitz, 2019; Van Overwalle et al., 2012) or motivations (e.g., Reeder, 2009a, 2009b; Reeder et al., 2001). Without question, goals and motivations represent subcategories of mental states. However, our findings indicate that people infer mental states that go beyond these previously documented inferences. In the example above of Pete complementing the waiter on the food, he might have the goal of *wanting to be nice* which may be relatively independent of his state of being *pleased*. Our studies are the first to demonstrate that people spontaneously draw state inferences beyond behavioral goals.

In this respect, our findings are in line with theoretical accounts on social inferences that presume that behavior attribution includes inferences of subcategories of states. For example, in the MIM, Reeder (2009b) suggests that people might draw spontaneous inferences about both motives and traits simultaneously from behavior. Similarly, Malle (2001, 2007) proposed in their folk theory of mind that people deliberately make sense of other people's behavior mainly in terms of beliefs, desires, and intentions. Our research expands the empirical support of these approaches in two ways. On the one hand, and as previously discussed, we have provided empirical evidence regarding the spontaneity of these inferences. Future research will help to further understand the process characteristics of both trait and state inferences. On the other hand, we enlarge the range of inferred states to affective or emotional states (e.g., happy, surprised, anxious, ashamed), cognitive and attentional states (e.g., focused, interested, distracted, bored) and physiological states (e.g., hungry, full, sick, tired), thus supporting the idea that observers' inferences about actors' states are not limited to motives, beliefs, desires, and intentions, but encompass a great variety of mental states. Future research may address the role of perceived intentionality in inference-making (i.e., whether there might be any differences between inferring intentional and unintentional states from behaviors).

Further implications for theory development

The finding of spontaneous and simultaneous trait and state inferences has potential implications for our theorizing about other effects of person perception and impression formation. For example, they question previous assumptions regarding the underlying mechanisms of the correspondence bias in person perception: Given that we are able to simultaneously and spontaneously draw trait and state inferences, how come that in deliberate impression formation we still favor dispositional inferences above all others? For example, Krull & Erickson's (1995) model of dispositional inference presumes that behavioral observations are initially and automatically processed only in terms of trait inferences, thus causing a dispositional bias in impression formation when controlled processes do not operate to correct these inferences (e.g., under cognitive load or time pressure; Gawronski, 2004). However, this assumption is difficult to uphold given the multiplicity of spontaneous social inferences.

In this line, Krull (1993) argued that the processing sequence and thus relative dominance of different inferences is not fixed, but depends on whether the inferential goal of the perceiver focuses on the acting person or the situation as causes of behavior: Whatever is in the perceiver's focal interest is processed spontaneously and effortlessly, thus dominates automatic inferences and is only revised given sufficient motivation and cognitive resources (Krull & Erickson, 1995). Based on this account, we may conclude that participants in our studies drew simultaneous state and trait inferences because both equally satisfy a predominant goal of forming an impression about the acting person. This rationale, however, cannot explain why and how perceivers generalize from spontaneous person impressions that initially include at least state and trait inferences – if not even further inferences – to correspondence biases that attribute behavior predominantly to only traits. While previous theorizing and research in the STI domain has treated person inferences as equal to dispositional inferences, our research along with others (e.g., Olcaysoy Okten & Moskowitz, 2018) documents that a more fine-grained theorizing is needed that includes different types of spontaneous person inferences from behavior (e.g., traits, states, intentions, values, etc.; Malle, 2005) and their functional consequences for impression formation. Even the more recent theories of spontaneous social inference processes so far fail to integrate the multiplicity of spontaneous social inferences into a coherent process model (e.g., Read & Miller, 2005; Trope & Gaunt, 2000).

We are convinced that future theorizing will benefit from more targeted considerations of different types of person inferences, how they may be related to each other and what independent and interactive roles they may have in impression formation from behavior. For example, one could speculate that spontaneous state inferences are inherently related to both trait as well as situation inferences from behavior. This conceptualization follows from an application of whole trait theory (Fleeson & Jayawickreme, 2015; Jayawickreme et al., 2019) to person perception. Whole trait theory defines traits as dispositions to experience trait-related states frequently and across many situations, hence to frequently feel and act in trait-specific ways. Construing traits this way suggests that perceivers may spontaneously infer an actor's current state and attribute it to a stable cause within the person. A similar mechanism could be construed for situational attributions: Perceivers may need to infer an actor's current state and attribute it to a situational cause outside the person. These speculations are not necessarily incompatible with long-standing models of impression formation (e.g., Gilbert et al., 1988; Krull, 1993; Trope, 1986), which assume that the first step in impression formation from behavior is "behavior categorization" (Trope, 1986), or "behavior interpretation" (Krull, 1993), followed by the actual attributional inference. It is possible that this processing stage may include inferring the actor's current state(s), which may then be attributed to traits or situations as causes of the behavior. It remains unclear, however, if states, traits, and/or situations are related in this fashion in spontaneous social inference processes. Just because we can theoretically construe a logical sequence of inferential processes, the actual processes or processing stages in spontaneous impression formation may look very different. However, our speculations illustrate the potential impact that these first results may have on further development of process-oriented theories and research on impression formation from behavior.

Besides the focus on understanding the psychological mechanisms of inference processes in spontaneous and controlled impression formation, we also think that considerations regarding the functional value of first impressions warrants further thought and development: How do spontaneous inferences that people draw from behavior affect their further judgement, decision making, and interaction behavior? Many scholars have offered a functional perspective of spontaneous impression formation processes, presuming that people draw inferences because they can make use of them for forming expectations and adapting their own (interaction) behavior. This served as an explanation for a frequently presumed dominance of trait inferences - because inferring a stable personal disposition as cause of behavior allows predictions of future behavior and thus adaptation for future interactions (e.g., Newman, 1996). However, such functional consequences of first impression formation have only rarely been studied (but see Costabile & Madon, 2019; McCarthy & Skowronski, 2011a, for exceptions). The results of our research also question this simplified approach: While we show that participants simultaneously draw state and trait inferences from behavior, results of Study 5 demonstrate that these same inferences are related to different expectations regarding the future and the likelihood that the same kind of behavior will be repeated. Thus, people may not only draw spontaneous inferences in order to predict future behavior, but also to infer temporal characteristics whose predictive value is limited. Further research is needed to assess the functional properties of spontaneous state inferences in more detail. Also, because trait and state inferences may potentially differ with regard to other functional characteristics (e.g., internal vs. external person attributions), which may differentially affect impression formation and their functional consequences in social information processing.

Boundary conditions of spontaneous social inferences

Similarly, future theorizing and research could and should address whether there are specific conditions that impact whether one type of inference would outweigh the other in spontaneous impression formation. Are there any characteristics of situations, perceivers, targets, and/or behaviors that moderate the occurrence of different types of inferences? In order to measure spontaneous inferences in the present research, we selected only those stimuli for which a majority of participants had generated the same deliberate inferences in the pretests. The selected stimuli represent only a small fraction of all tested materials, let alone the theoretical space of all behaviors, which suggests that strong stimulus constraints may apply. It makes sense to expect that inferences drawn from behavioral statements are strongly regulated by the stimuli themselves. One the one hand, there may be (socially shared) classification rules about presumed causes of certain types of behavior (causal schemata, implicit personality theories, e.g., Anderson, 1981; Kelley, 1973). On the other hand, the additional information provided in a behavioral statement may also impact the likelihood of certain inferences. For example, linguistic markers of temporal consistency have been demonstrated to increase trait inferences in deliberate impression formation (Olcaysoy Okten & Moskowitz, 2018). Thus, behaviors and behavioral statements themselves present strong constraints with regard to the inferences that can be drawn from them. Similarly, target and/or perceiver effects are likely to influence inferences drawn from behavior (Kenny, 2019). For example, research on the linguistic intergroup bias (LIB; Maass, Salvi, Arcuri, & Semin, 1989) has provided evidence that people represent others' behavior differently depending on its evaluation and whether it is performed by ingroup vs. outgroup members: Whereas desired ingroup behavior and undesired outgroup behavior is represented in more abstract terms (e.g., by using trait adjectives), undesired behavior of ingroup members and desired behavior of outgroup members is represented in less abstract forms (e.g., by using state verbs). Although we are not aware of any research investigating if the LIB applies to spontaneous impression formation as well, it is conceivable that observers apply state and trait inferences differently to the same behaviors based on the group memberships of actor and perceiver. A further moderating condition may be cultural differences in person construal. There is abundant research showing that people with an individualistic cultural mindset tend to construe individuals as rather separate entities with higher individualistic agency, whereas people with a collectivistic cultural mindset tend to construe individuals and their agency as fundamentally connected to their social and non-social environment (e.g., Markus, Uchida, Omoregie, Townsend, & Kitayama, 2006). Some cross-cultural research has indicated lower levels of trait inferences in collectivistic than individualistic cultures (Shimizu, 2017; but see Bott et al., 2021). For state inferences, there may be a reversed cultural difference. Last but not least, the temporary mindset and goals of an observer may constitute important moderating conditions, as has already been shown with respect to trait versus situation inferences. In our own exploratory attempts at identifying individual differences, we did not find any effects relating spontaneous state (and trait) inferences with implicit personality theories (Dweck et al., 1995) or construal level (Navon, 1977; Trope & Liberman, 2010, see online supplemental materials).

The process of person perception from behavior is very complex, and it is likely that a multitude of factors influence which inferences a specific perceiver draws from a specific behavior shown by an specific actor. Caution is advised to not over-generalize our current findings to person perception in general. Future research is needed to investigate which processes contribute to social inferences from behavior and to better understand the systematic behind aforementioned stimulus constraints and how target and perceiver effects may impact spontaneous state and trait inferences in deliberate and spontaneous impression formation from behavior.

In light of the multitude of spontaneous social inferences, we need to further develop our theorizing of impression formation from behavior to understand why and how we tend to overly attribute observed behaviors to peoples' dispositions. We are confident that the presented research constitutes an important first step in this direction.

References

Ames, D. R. (2004). Inside the mind reader's tool kit: Projection and stereotyping in mental state inference. *Journal of Personality and Social Psychology*, *87*(3), 340–353. doi:10.1037/0022-3514.87.3.340

Anderson, N. H. (1981). *Foundations of information theory*. New York: Academic Press.

Apperly, I. A., Riggs, K. J., Simpson, A., Chiavarino, C., & Samson, D. (2006). Is belief reasoning automatic? *Psychological Science*, *17*(10), 841–844. doi:10.1111/j.1467-9280.2006.01791.x

Bainbridge, W. A., Isola, P., & Oliva, A. (2013). The intrinsic memorability of face photographs. *Journal of Experimental Psychology: General*, *142*(4), 1323–1334. doi:10.1037/a0033872

Biner, P. M., Angle, S. T., Park, J. H., Mellinger, A. E., & Barber, B. C. (1995). Need state and the illusion of control. *Personality and Social Psychology Bulletin*, *21*(9), 899–907. doi:10.1177/0146167295219004

Bott, A., Brockmann, L., Denneberg, I., Henken, E., Kuper, N., Kruse, F., & Degner, J. (2021). *Spontaneous trait inferences from behavior: A systematic meta-analysis*. Manuscript submitted for publication.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, *64*(2), 123–152.

Campbell, J. I. D., & Thompson, V. A. (2012). MorePower 6.0 for ANOVA with relational confidence intervals and Bayesian analysis. *Behavior Research Methods*, *44*(4), 1255–1265. doi:10.3758/s13428-012-0186-0

Cantor, N., & Mischel, W. (1977). Traits as prototypes: Effects on recognition memory. *Journal of Personality and Social Psychology*, *35*(1), 38–48. doi:10.1037/0022-3514.35.1.38

Cantor, N., & Mischel, W. (1979). Prototypes in person perception. *Advances in Experimental Social Psychology*, *12*(100), 3–52.

Carlston, D. E., Skowronski, J. J., & Sparks, C. (1995). Savings in relearning: II. On the formation of behavior-based trait associations and inferences. *Journal of Personality and Social Psychology*, *69*(3), 420–35. doi:10.1037/0022-3514.69.3.429

Carlston, D. E., & Skowronski, J. J. (1994). Savings in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology*, *66*(5), 840–856. doi:10.1037/0022-3514.66.5.840

Carlston, D. E., & Skowronski, J. J. (2005). Linking versus thinking: Evidence for the different associative and attributional bases of spontaneous trait transference and spontaneous trait inference. *Journal of Personality and Social Psychology*, *89*(6), 884–98. doi:10.1037/0022-3514.89.6.884

Cook, A. E., & O'Brien, E. J. (2017). Fundamentals of inferencing during reading. *Language and Linguistics Compass*, *11*(7), 1–16. doi:10.1111/lnc3.12246

Costabile, K. A., & Madon, S. (2019). Downstream effects of dispositional inferences on confirmation biases. *Personality and Social Psychology Bulletin*, *45*(4), 557–570. doi:10.1177/0146167218789624

Diergarten, A. K., & Nieding, G. (2016). Online emotional inferences in written and auditory texts: a study with children and adults. *Reading and Writing*, *29*(7), 1383–1407. doi:10.1007/s11145-016-9642-z

Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*(4), 568–581.

Dweck, C. S., Chiu, C., & Hong, Y. (1995). Implicit theories and their role in judgments and reactions : A word from two perspectives. *Psychological Inquiry*, *6*(4), 267–285. doi:10.1207/s15327965pli0604

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, *39*(2), 175–191. doi:10.3758/BF03193146

Ferreira, M. B., Garcia-Marques, L., Hamilton, D., Ramos, T., Uleman, J. S., & Jeronimo, R. (2012). On the relation between spontaneous trait inferences and intentional inferences: An inference monitoring hypothesis. *Journal of Experimental Social Psychology*, *48*(1), 1–12. doi:10.1016/j.jesp.2011.06.013

Fiedler, K., & Schenck, W. (2001). Spontaneous inferences from pictorially presented behaviors. *Personality and Social Psychology Bulletin*, *27*(11), 1533–1546. doi:10.1177/01461672012711013

Fiedler, K., Watling, M., Menges, J. I., & Schenck, W. (2005). Priming trait inferences through pictures and moving pictures: The impact of open and closed mindsets. *Journal of Personality and Social Psychology*, *88*(2), 229–244. doi:10.1037/0022-3514.88.2.229

Fleeson, W., & Jayawickreme, E. (2015). Whole trait theory. *Journal of Research in Personality*, *56*, 82–92. doi:10.1016/j.jrp.2014.10.009

Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European Review of Social Psychology*, *15*(1), 183–217. doi:10.1080/10463280440000026

Gernsbacher, M. A., Goldsmith, H. H., & Robertson, R. R. (1992). Do readers mentally represent characters' emotional states? *Cognition and Emotion*, *6*(2), 89–111. doi:10.1080/02699939208411061

Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, *54*(5), 733–740. doi:10.1037/0022-3514.54.5.733

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, *101*(3), 371–395. doi:10.1037/0033-295x.101.3.371

Ham, J., & Vonk, R. (2003). Smart and easy: Co-occurring activation of spontaneous trait inferences and spontaneous situational inferences. *Journal of Experimental Social Psychology*, *39*(5), 434–447. doi:10.1016/S0022-1031(03)00033-7

Hamaker, E. L., Nesselroade, J. R., & Molenaar, P. C. M. (2007). The integrated trait-state model. *Journal of Research in Personality*, *41*(2), 295–315. doi:10.1016/j.jrp.2006.04.003

Hassin, R. R., Aarts, H., & Ferguson, M. J. (2005). Automatic goal inferences. *Journal of Experimental Social Psychology*, *41*(2), 129–140. doi:10.1016/j.jesp.2004.06.008

Heider, F. (1958). *The psychology of interpersonal relations*. New York: John Wiley & Sons, Inc.

Hewstone, M., Stroebe, W., & Jonas, K. (2016). *An introduction to social psychology*. Hoboken, NJ: John Wiley & Sons.

Hoffman, C., Mischel, W., & Mazze, K. (1981). The role of purpose in the organization of information about behavior: Trait-based versus goal-based categories in person cognition. *Journal of Personality and Social Psychology*, *40*(2), 211–225. doi:10.1037/0022-3514.40.2.211

Jayawickreme, E., Zachry, C. E., & Fleeson, W. (2019). Whole trait theory: An integrative approach to examining personality structure and process. *Personality and Individual Differences*, *136*, 2–11. doi:10.1016/j.paid.2018.06.045

Jones, E. E., & Davis, K. (1965). From acts to dispositions: The attribution process in social psychology. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219–266). doi:10.1016/S0065-2601(08)60107-0

Kelley, H. H. (1967). Attribution theory in social psychology. In *Nebraska symposium on motivation*. University of Nebraska Press.

Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, *28*(2), 107–128. doi:10.1037/h0034225

Kenny, D. A. (2019). *Interpersonal Perception* (Second Edition). New York: The Guilford Press.

Kiley Hamlin, J., Ullman, T., Tenenbaum, J., Goodman, N., & Baker, C. (2013). The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Developmental Science*, *16*(2), 209–226. doi:10.1111/desc.12017

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge University Press.

Korman, J., & Malle, B. F. (2016). Grasping for traits or reasons? How people grapple with puzzling social behaviors. *Personality and Social Psychology Bulletin*, *42*(11), 1451–1465. doi:10.1177/0146167216663704

Krull, D. S. (1993). Does the grist change the mill? The effect of the perceiver's inferential goal on the process of social inference. *Personality and Social Psychology Bulletin*, *19*(3), 340–348.

Krull, D. S., & Erickson, D. J. (1995). Inferential hopscotch: How people draw social inferences from behavior. *Current Directions in Psychological Science*, *4*(2), 35–38. doi:10.1111/1467-8721.ep10770986

Kunda, Z., Sinclair, L., & Griffin, D. (1997). Equal ratings but separate meanings: Stereotypes and the construal of traits. *Journal of Personality and Social Psychology*, *72*(4), 720–734. doi:10.1037/0022-3514.72.4.720

Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, *44*(7), 701–710. doi:10.1002/ejsp.2023

Levordashka, A., & Utz, S. (2017). Spontaneous trait inferences on social media. *Social Psychological and Personality Science*, *8*(1), 93–101. doi:10 . 1177 / 1948550616663803

Lillard, A. S., & Skibbe, L. (2006). Theory of mind: Conscious attribution and spontaneous trait inference. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 277–305). New York: Oxford University Press.

Lupfer, M. B., Clark, L. F., & Hutcherson, H. W. (1990). Impact of context on spontaneous trait and situational attributions. *Journal of Personality and Social Psychology*, *58*(2), 239–249. doi:10.1037/0022-3514.58.2.239

Maass, A., Salvi, D., Arcuri, L., & Semin, G. (1989). Language use in intergroup contexts: the linguistic intergroup bias. *Journal of Personality and Social Psychology*, *57*(6), 981–993. doi:10.1037/0022-3514.57.6.981

Malle, B. F. (2005). Folk theory of mind: Conceptual foundations of human social cognition. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 225–255). New York: Oxford University Press.

Malle, B. F., Moses, L. J., & Baldwin, D. A. (Eds.). (2001). *Intentions and intentionality: Foundations of social cognition*. Cambridge, MA: MIT Press.

Malle, B. F. (2004). *How the mind explains behavior*. Cambridge, MA: MIT Press.

Malle, B. F. (2008). Fritz Heider's legacy: Celebrated insights, many of them misunderstood. *Social Psychology*, *39*(3), 163–173. doi:10.1027/1864-9335.39.3.163

Malle, B. F., & Holbrook, J. (2012). Is there a hierarchy of social inferences? The likelihood and speed of inferring intentionality, mind, and personality. *Journal of Personality and Social Psychology*, *102*(4), 661–684. doi:10.1037/a0026790

Malle, B. F., Knobe, J. M., & Nelson, S. E. (2007). Actor-observer asymmetries in explanations of behavior: New answers to an old question. *Journal of Personality and Social Psychology*, *93*(4), 491–514. doi:10.1037/0022-3514.93.4.491

Markus, H. R., Uchida, Y., Omoregie, H., Townsend, S. S., & Kitayama, S. (2006). Going for the gold models of agency in Japanese and American contexts. *Psychological Science*, *17*(2), 103–112. doi:10.1111/j.1467-9280.2006.01672.x

McCarthy, R. J., & Skowronski, J. J. (2011a). The interplay of controlled and automatic processing in the expression of spontaneously inferred traits: A PDP analysis. *Journal of Personality and Social Psychology*, *100*(2), 229–240. doi:10.1037/a0021991

McCarthy, R. J., & Skowronski, J. J. (2011b). What will Phil do next? Spontaneously inferred traits influence predictions of behavior. *Journal of Experimental Social Psychology*, *47*(2), 321–332. doi:10.1016/j.jesp.2010.10.015

McClung, J. S., & Reicher, S. D. (2018). Representing other minds: Mental state reference is moderated by group membership. *Journal of Experimental Social Psychology*, *76*, 385–392. doi:10.1016/j.jesp.2017.11.013

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(1), 82–91. doi:10.1037/0278-7393.12.1.82

Mischel, W. (1968). *Personality and assessment*. New York: John Wiley & Sons, Inc.

Moskowitz, G., & Olcaysoy Okten, I. (2016). Spontaneous goal inference (SGI). *Social and Personality Psychology Compass*, *10*(1), 64–80. doi:10.1111/spc3.12232

Moskowitz, G. B. (2005). Correspondence bias and spontaneous trait inferences. In G. B. Moskowitz (Ed.), *Social cognition: Understanding self and others* (pp. 267–309). New York, New York: Guilford Press.

Myers, D., & Twenge, J. (2018). *Social psychology*. McGraw-Hill Education.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9*(3), 353–383. doi:10.1016/0010-0285(77)90012-3

Newman, L. (1991). Why are traits inferred spontaneously? A developmental approach. *Social Cognition*, *9*(3), 221–253. doi:10.1521/soco.1991.9.3.221

Newman, L. S. (1996). Trait impressions as heuristics for predicting future behavior. *Personality and Social Psychology Bulletin*, *22*(4), 395–411. doi:10.1177/0146167296224006

Nussbaum, S., Trope, Y., & Liberman, N. (2003). Creeping Dispositionism: The Temporal Dynamics of Behavior Prediction. *Journal of Personality and Social Psychology*, *84*(3), 485–497. doi:10.1037/0022-3514.84.3.485

Olcaysoy Okten, I., & Moskowitz, G. B. (2018). Goal versus trait explanations: Causal attributions beyond the trait-situation dichotomy. *Journal of Personality and Social Psychology*, *114*(2), 53608. doi:10.1037/pspa0000104

Olcaysoy Okten, I., & Moskowitz, G. B. (2019). Spontaneous goal versus spontaneous trait inferences and explanations: How ideology shapes attributions. *European Journal of Social Psychology*, *117*(1), 1–12. doi:10.1002/ejsp.2611

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872. doi:10.1016/j.jesp.2009.03.009

Orghian, D., Ramos, T., Reis, J., & Garcia-Marques, L. (2018). Acknowledging the role of word-based activation in spontaneous trait inferences. *Analise Psicologica*, *36*(1), 115–131. doi:10.14417/ap.1319

Pahl, R. (2018). *Groupseq: A gui-based program to compute probabilities regarding group sequential designs*. R package version 1.3.5.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., … Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. doi:10.3758/s13428-018-01193-y

R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.

Ramos, T., Garcia-Marques, L., Hamilton, D. L., Ferreira, M., & Van Acker, K. (2012). What I infer depends on who you are: The influence of stereotypes on trait and situational spontaneous inferences. *Journal of Experimental Social Psychology*, *48*(6), 1247–1256. doi:10.1016/j.jesp.2012.05.009

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, *114*(3), 510–532. doi:10.1037/0033-2909.114.3.510

Read, S. J., & Miller, L. C. (2005). Explanatory coherence and goal-based knowledge structures in making dispositional inferences. In B. F. Malle & S. D. Hodges (Eds.), *Other minds: How humans bridge the divide between self and others* (pp. 124–139). Guilford Press.

Reeder, G. D. (2009a). Mindreading and dispositional inference: MIM revised and extended. *Psychological Inquiry*, *20*(1), 73–83. doi:10.1080/10478400902868565

Reeder, G. D. (2009b). Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychological Inquiry*, *20*(1), 1–18. doi:10 . 1080 / 10478400802615744

Reeder, G. D., Hesson-McInnis, M., Krohse, J. O., & Scialabba, E. A. (2001). Inferences about effort and ability. *Personality and Social Psychology Bulletin*, *27*(9), 1225–1235. doi:10.1177/0146167201279014

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173–220). doi:10.1016/S0065-2601(08)60357-3

Ruffman, T. (2014). To belief or not belief: Children's theory of mind. *Developmental Review*, *34*(3), 265–293. doi:10.1016/j.dr.2014.04.001

Schneid, E. D., Carlston, D. E., & Skowronski, J. J. (2015). Spontaneous evaluative inferences and their relationship to spontaneous trait inferences. *Journal of Personality and Social Psychology*, *108*(5), 681–696. doi:10.1037/a0039118

Schneid, E. D., Crawford, M. T., Skowronski, J. J., Irwin, L. M., & Carlston, D. E. (2015). Thinking about other people: Spontaneous trait inferences and spontaneous evaluations. *Social Psychology*, *46*(1), 24–35. doi:10.1027/1864-9335/a000218

Scott, R. M., & Baillargeon, R. (2017). Early false-belief understanding. *Trends in Cognitive Sciences*, *21*(4), 237–249. doi:10.1016/j.tics.2017.01.012

Shimizu, Y. (2017). Why are negative behaviours likely to be immediately invoked traits? The effects of valence and frequency on spontaneous trait inferences. *Asian Journal of Social Psychology*, *20*(3-4), 201–210. doi:10.1111/ajsp.12183

Stewart, T. L., Weeks, M., & Lupfer, M. B. (2004). Spontaneous stereotyping: A matter of prejudice? *Social Cognition*, *21*(4), 263–298. doi:10.1521/soco.21.4.263.27003

Thornton, M. A., & Tamir, D. I. (2020). People represent mental states in terms of rationality, social impact, and valence: Validating the 3d Mind Model. *Cortex*, *125*, 44–59. doi:10.1016/j.cortex.2019.12.012

Todd, A. R., Molden, D. C., Ham, J., & Vonk, R. (2011). The automatic and co-occurring activation of multiple social inferences. *Journal of Experimental Social Psychology*, *47*(1), 37–49. doi:10.1016/j.jesp.2010.08.006

Todorov, A., & Uleman, J. S. (2004). The person reference process in spontaneous trait inferences. *Journal of Personality and Social Psychology*, *87*(4), 482–493. doi:10.1037/0022-3514.87.4.482

Todorov, A., & Uleman, J. S. (2002). Spontaneous trait inferences are bound to actors' faces: Evidence from a false recognition paradigm. *Journal of Personality and Social Psychology*, *83*(5), 1051–1065. doi:10.1037/0022-3514.83.5.1051

Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology*, *39*(6), 549–562. doi:10.1016/S0022-1031(03)00059-3

Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, *93*(3), 239–257. doi:10.1037/0033-295X.93.3.239

Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology*, *79*(3), 344–354. doi:10.1037//0022-3514.79.3.344

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, *117*(2), 440–463. doi:10.1037/a0018963

Uleman, J. S. (2005). Introduction: Becoming Aware of the New Unconscious. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 3–18). New York, New York: Oxford University Press.

Uleman, J. S., Hon, A., Roman, R. J., & Moskowitz, G. B. (1996). On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin*, *22*(4), 377–394. doi:10 . 1177 / 0146167296224005

Uleman, J. S., Rim, S., Saribay, A. S., & Kressel, L. M. (2012). Controversies, questions, and prospects for spontaneous social inferences. *Social and Personality Psychology Compass*, *6*(9), 657–673. doi:10.1111/j.1751-9004.2012.00452.x

Uleman, J. S., Saribay, S. A., & Gonzalez, C. M. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual review of psychology*, *59*, 329–60. doi:10.1146/annurev.psych.59.103006.093707

Van Overwalle, F., Van Duynslaeger, M., Coomans, D., & Timmermans, B. (2012). Spontaneous goal inferences are often inferred faster than spontaneous trait inferences. *Journal of Experimental Social Psychology*, *48*(1), 13–18. doi:10.1016/j.jesp.2011.06.016

Wells, B. M., Skowronski, J. J., Crawford, M. T., Scherer, C. R., & Carlston, D. E. (2011). Inference making and linking both require thinking: Spontaneous trait inference and spontaneous trait transference both rely on working memory capacity. *Journal of Experimental Social Psychology*, *47*(6), 1116–1126. doi:10.1016/j.jesp.2011.05.013

Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of inferences. *Journal of Personality and Social Psychology*, *47*(2), 237–252. doi:10.1037/0022-3514.47.2.237

List of Figures

List of Figures

List of Tables

Colophon

This thesis was typeset with $\LaTeX\ 2_\varepsilon$. It uses an adaptation of the *Clean Thesis* style developed by Ricardo Langner.

Erklärung

Hiermit erkläre ich, Felix Kruse, an Eides statt, gemäß §5 (4d) und §9 (1c und 1d) der Promotionsordnung des Instituts für Psychologie der Universität Hamburg vom 20.08.2003

1. dass ich mich an einer anderen Universität oder Fakultät noch keiner Doktorprüfung unterzogen oder mich um Zulassung zu einer Doktorprüfung bemüht habe.

2. dass die von mir vorgelegte Dissertation nicht Gegenstand eines anderen Prüfungsverfahrens gewesen oder in einem solchen Verfahren als ungenügend beurteilt worden ist.

3. dass ich die von mir vorgelegte Dissertation selbst verfasst, keine anderen als die angegebenen Quellen und Hilfsmittel benutzt und keine kommerzielle Promotionsberatung in Anspruch genommen habe. Die wörtlich oder inhaltlich übernommenen Stellen habe ich als solche kenntlich gemacht.

4. dass die gedruckte und die digitale Version dieser Dissertation inhaltlich identisch sind. Die gedruckte Version ist mit Ausnahme der Abbildungen in Graustufen gedruckt, die digitale Version ist vollfarbig.

Hamburg, 31.03.2022

Felix Kruse