Universitätsklinikum Hamburg- Eppendorf

Zentrum für Experimentelle Medizine Institut für Neurophysiologie und Pathophysiologie Direktor: Prof. Dr. Andreas Engel

MECHANISTIC AND NEURAL BASIS OF CHOICE-INDUCED BIASES IN DECISION-MAKING

Dissertation

zur Erlangung des Doktorgrades PhD

an der Medizinischen Fakultät der Universität Hamburg

vorgelegt von

Bharath Chandra Talluri

aus Indien

Hamburg- May, 2020

Angenommen von der Medizinischen Fakultät am: 17.03.2021 (Datum der mündl. Prüfung eintragen)

Veröffentlicht mit Genehmigung der Medizinischen Fakultät der Universität Hamburg.

Prüfungsausschuss, der/die Vorsitzende	Prof. Dr. Tobias Donner
Prüfungsausschuss, zweite/r Gutachter/in	Prof. Dr. Philipp Sterzer
Prüfungsausschuss, dritte/r Gutachter/in	Prof. Dr. Redmond O'Connell

We think, each of us, that we're much more rational than we are. And we think that we make our decisions because we have good reasons to make them. Even when it's the other way around. We believe in the reasons, because we've already made the decision

- Daniel Kahneman

CONTENTS

1	INTR	ODUCTION	3
	1.1	Choice-induced biases in decision-making	3
	1.2	Decision-making under experimental control	4
	1.3	Neural basis of perceptual decisions	7
	1.4	Confirmation bias in decision-making	11
	1.5	Temporal weighting of decision-relevant evidence	13
	1.6	Sequential bias in perceptual decisions	15
	1.7	Outline of the thesis	16
2	CON	FIRMATION BIAS THROUGH SELECTIVE OVERWEIGHTING OF	
	сно	ICE-CONSISTENT EVIDENCE	18
	2.1	Results	19
	2.2	Discussion	26
	2.3	Methods	29
	2.4	Data and software availability	49
	2.5	Acknowledgments	49
	2.6	Supplementary Figures	50
3	сно	ICES CHANGE THE TEMPORAL WEIGHTING OF DECISION EVI-	
	DEN	CE	54
	3.1	Introduction	55
	3.2	Results	57
	3.3	Discussion	63
	3.4	Methods	65
	3.5	Data availability	70
	3.6	Acknowledgments	70
4	NEUI	RAL BASIS OF SEQUENTIAL BIAS IN PERCEPTUAL DECISION-	
	MAK	ING	71
	4.1	Introduction	71
	4.2	Results	73
	4.3	Discussion	81
	4.4	Methods	84
	4.5	Acknowledgments	96

	4.6	Supplementary Figures	97
5	DISCUSSION		100
	5.1	Choice-induced biases in decision-making	102
	5.2	Confirmation bias in perceptual decision-making	103
	5.3	Temporal weighting of perceptual evidence	104
	5.4	Concluding remarks	105
	GENERAL SUMMARY		106
	BIBLIOGRAPHY		111
	ACKNOWLEDGMENTS		124
	CURI	RICULUM VITAE	128
	EIDE	SSTATTLICHE ERKLÄRUNG	129

ACRONYMS

ACC	Anterior Cingulate Cortex
AIP	Anterior Intraparietal Cortex
ANOVA	Analysis of Variance
BIC	Bayes Information Criterion
BOLD	Blood oxygenation level dependent
DDM	Drift Diffusion Model
dIPFC	dorsolateral Pre-frontal Cortex
EEG	Electroencephalography
FEF	Frontal Eye Fields
fMRI	functional Magnetic Resonance Imaging
FWHM	Full-width half maximum
GLM	Generalised Linear Models
HRF	Hemodynamic Response Function
IP	Intraparietal Cortex
IPS	Intraparietal Sulcus
ITI	Inter-trial Interval
LCA	Leaky Competing Accumulator model
LIP	Lateral Intra-parietal Area
LIPv/d	ventral/dorsal Lateral Intra-parietal Area
M1	Primary Motor cortex
MIP	Middle Intraparietal cortex
MEG	Magnetoencephalography
MOCS	Method of Constant Stimuli
MT	Middle Temporal Area
MVPA	Multivoxel Pattern Analysis
OFC	Orbito-frontal Cortex
PFC	Pre-frontal Cortex
PMd	dorsal Pre-Motor area
PPC	Posterior Parietal Cortex
ROC	Receiver-Operating Characteristic
ROI	Region of Interest
RT	Reaction Time
SC	Superior Colliculus
SEM	Standard Error of Means
SMA	Supplemental Motor Area
SVM	Support Vector Machines

INTRODUCTION

The survival of an organism depends on its ability to make inferences about the state of its world, a behavior exhibited by even single-celled non-neural organisms (Eisenstein 1975). The human brain has since come a long way in its ability to make decisions, ranging from simple sensory decisions such as detecting an apple among a bunch of oranges, to more complex decisions such as whom to vote for in an upcoming election. Decision-making is the deliberate process of evaluating information about various alternatives and committing to one. Imagine you are an emergency room doctor on the first day of Covid19 outbreak in your city. You have since diagnosed a few patients who tested positive for the disease. A new patient approached you suffering with cough, fever, and fatigue (a set of symptoms common to both seasonal flu and Covid 19), and you have to decide on a set of follow-up diagnostic tests. You choose to test for Covid19, since the last few patients with a similar set of symptoms you encountered tested positive for the disease. Would your decision about the follow-up tests change if you encounter the same patient but before the start of the outbreak? Common introspection suggests that you are more likely to test for seasonal flu first due to its prevalence in winter. Now, imagine you are a voter in an upcoming election. Throughout the electoral campaign, you follow the candidates, analyze their policy proposals at various time points, and make up your mind to vote for a candidate. Should a political scandal concerning your chosen candidate arise, will you repeat the entire deliberation process, now taking this new information into account? Again, introspection suggests that evaluation depends on the timing of information- whether it arrives before or after you make up your mind.

1.1 CHOICE-INDUCED BIASES IN DECISION-MAKING

People's behavior often exhibits strong biases i.e., systematic tendencies towards a particular choice, irrespective of the information available to them.

Sometimes, such biases arise from their own past choices. In the example above, the new patient may not be related to the other patients who tested positive for Covid19. Still, the doctor's decision about the follow-up tests was biased since the last few patients she encountered had the disease. This phenomenon where new information is evaluated in a biased manner depending on recent choice history is called **Sequential bias**.

People tend to treat information favorably if it agrees with their beliefs, a tendency we all experience at times. In some cases, beliefs arise from past choices themselves. For example, after committing in a poll that you will vote for a political candidate, you feel confirmed by the candidate's good policy proposals, but ignore the bad ones. This phenomenon where new information is treated preferentially depending on whether or not it agrees with your previously-held beliefs is called **Confirmation bias**.

In this thesis, I investigated the mechanistic and neural basis of these choice-induced biases. Using carefully controlled laboratory studies of decisionmaking, I addressed the following questions: What is the mechanistic basis of confirmation bias? How do intermittent choices impact the evaluation of decision-relevant evidence? What are the neural mechanisms underlying sequential bias?

1.2 DECISION-MAKING UNDER EXPERIMENTAL CON-

Scientific investigations of decision-making require careful quantification of behavior under controlled experimental settings. This is made possible by psychophysics, the study of the relationship between subjective sensations and the physical stimuli that produce them. A class of decisions that lend themselves well to such investigations are perceptual decisions. Perceptual decision-making is the process of deciding the identity of sensory information from a fixed set of alternatives. These decisions provide the experimenter with precise control over the sensory information available to the subject, enabling precise quantification of the effect of changes in sensory information over behavior. A typical perceptual decision-making experiment consists of tens to hundreds of individual *trials*, each up to a few seconds long. Each trial contains a presentation of sensory stimulus (also referred to as sensory evidence) and the subjects' response about that stimulus, in that order. All the experiments in this thesis showed the stimulus for a fixed duration, following which subjects were prompted for a response (so called interrogation design). Another type of commonly used trial design does not decouple the two events, instead, subjects respond immediately after they make up their mind during the stimulus presentation window (so called reaction-time design). Figure 1A shows an example sequence of events from an experiment in the thesis. Each trial starts with a baseline period alerting the subject about the upcoming sensory signal, followed by a short duration of sensory signal about which the decision has to be made. A cue then instructs subject to report their decision, followed by an inter-trial interval. By carefully monitoring responses given by subjects to repeated presentations of different intensities of the sensory stimulus, the experimenter can quantify subjects' choice patterns as a function of sensory input, using a **psychometric func**tion. A psychometric function models the relationship between a particular feature of the sensory evidence, and choices made by the subjects (Figure 1B). It allows the experimenter to make inferences about the sensitivity of the subject to a given sensory input, subject's bias towards one alternative, and the shape of the stimulus-response function.

All perceptual decision-making experiments described in this thesis used random dot motion as the sensory stimulus (Figure 1C). Random dot motion stimulus, as the name indicates contains a set of dots moving across different directions in a circular aperture. A proportion of dots, called coherent dots, move along a pre-defined direction, while the rest of the dots, called noise dots, move along random directions. Consequently, the directions of noise dots cancel each other, resulting in a net motion direction of the stimulus pattern along the coherent dot direction. The stochastic, and temporally extended nature of the signal makes random dot stimuli well suited to investigate certain aspects of perceptual decisions: subjects need to integrate the direction information across multiple frames of the stimulus to average away the directions of noise dots and improve their estimate of the signal (so-called signal-to-noise ratio). The proportion of coherent dots in every trial can be fixed by the experimenter to pre-determined levels.



Figure 1: **A**.Example sequence of events in a perceptual decision-making experiment with interrogation design. Different events during the course of a trial are cued by the fixation color. **B**. Illustration of a psychometric function. The slope of the function quantifies the subject's perceptual sensitivity, horizontal offset of the function quantifies subject's bias towards one choice. **C**.Illustration of random dot motion stimulus, showing pure random motion (Upper panel), and upwards coherent motion (lower panel).

People often make decisions based on information that unfolds over time. The computational mechanisms underlying such decisions have been formalized by a class of models called sequential sampling models. These models hypothesize that sensory evidence is gradually accumulated over time, and is transformed into a choice when this accumulated evidence (often referred to as decision-variable) reaches a threshold or a bound (Gold and Shadlen 2007). Well-known examples of this class of models include the drift diffusion model (DDM; Ratcliff 1978), and the leaky competing accumulator model (LCA; Usher and McClelland 2001). In two-alternative choice tasks for example, the DDM models an accumulator that integrates the difference in evidence supporting each of the alternatives linearly over time, without any loss. On the other hand, the LCA models separate *leaky* (leak refers to decay of accumulated information over time) accumulators for each alternative that compete with each other through mutual inhibition. These class of models have been widely accepted among decision-making researchers because of their ability to account for behavior in a variety of decision-making tasks (Bogacz et al. 2006; Ratcliff and McKoon 2008), and since, neural correlates of temporal accumulation have been identified in a variety of brain regions (see next section). However, finding neural correlates does not necessarily answer the question about how neural circuits implement such a temporal accumulation process. An important clue to address this question comes from the observation that some of the areas where correlates for the accumulation process hava been identified also show elevated persistent activity during working memory maintenance (Wang 2008), suggesting that similar neural mechanisms could underlie both the processes. Neural circuit models

developed with anatomically plausible architecture (commonly referred to as biophysically based models) showed that recurrent synaptic excitation, balanced by fast inhibition generates self-sustained **attractor states**, and slow transient dynamics in such networks could produce the dynamics observed during the temporal accumulation of information (Wang 2002; Wong and Wang 2006). Thus, as evidence is accumulated during the course of a trial, the neural circuits gradually settle into an attractor state that corresponds to one of the alternatives. The strength of excitation in such circuits determines the timescale of temporal accumulation, with stronger recurrent excitation facilitating the integration of information over longer timescales.

1.3 NEURAL BASIS OF PERCEPTUAL DECISIONS

Extensive research over the last few decades into perceptual decisions has provided a rich understanding of the neurophysiological mechanisms underlying perceptual decision-making (Gold and Shadlen 2007; Shadlen and Kiani 2013). The process by which sensory information is translated into a perceptual choice involves three stages: (i) sensory encoding stage, where the incoming sensory information is processed (ii) evidence accumulation stage, where the encoded sensory information is accumulated over time towards one of the possible alternatives, and (iii) choice commitment stage, where the accumulated information is transformed into an action plan once it reaches a threshold (Figure 2). The next paragraph summarizes investigations into the neural correlates of each of these stages in non-human primates.

Neural responses that correlate with perceptual choice of monkeys about random dot stimuli have been observed in visual area MT (middle temporal area or V5) (Britten et al. 1992, 1996; Newsome et al. 1989), an area previously shown to respond optimally to visual stimuli of particular direction and speed (Maunsell and van Essen 1983); and that about vibrotactile stimulus have been observed in the somatosensory cortex (Romo et al. 2002), previously shown to encode the frequency of vibration of a tactile stimulus (Mountcastle et al. 1969). The causal role of these sensory neurons in the decision-making process has been established via lesion, and microstimulation studies (Romo et al. 2000; Salzman et al. 1992). These findings suggest



Figure 2: Neural basis of perceptual decisions in the human brain. Sensory regions encode momentary sensory evidence which is fed forward to parietal, and prefrontal regions where it is accumulated across time. The accumulated evidence is transformed into an action plan once it reaches a threshold. Black arrows denote information flow.

that decision-making process in the brain starts with early sensory regions, with the sensory neurons encoding moment-by-moment sensory evidence, which is then fed to the next stage for accumulation across time. The sensory areas are thus referred to as **low-level areas**, owing to their presence in the decision-making hierarchy. Neural correlates of evidence accumulation have been identified in LIP (lateral intraparietal area) (Kiani et al. 2008; Roitman and Shadlen 2002; Shadlen and Newsome 1996, 2001), in FEF (frontal eye fields) (Gold and Shadlen 2000), and in PFC (prefrontal cortex) (Kiani et al. 2014; Kim and Shadlen 1999). The causal role of LIP neurons was established by microstimulation (Hanks et al. 2006; but see Katz et al. 2016), and by observing neuronal responses to brief perturbations of stimuli in a trial (Huk and Shadlen 2005). Microstimulation of neurons in FEF, a region that responds selectively for anticipatory eye-movements (Bruce and Goldberg 1985), showed that the correlates of the monkey's decision in this region corresponded to the development of motor commands (Ding and Gold 2012; Gold and Shadlen 2000, 2003), rather than the decision itself. However, it is difficult to disentangle the two when the mapping between the decision and the motor responses used to report the decision (eg., left saccade to indicate leftward choice, and right saccade to indicate rightward choice) is fixed (which is usually the case), in RT versions of the task. An important clue to disentangle these two processes came from the observation that the firing rates of choice-selective neurons in LIP reached a stereotypical level of

excitation independent of the strength of the stimulus just before execution of the motor response corresponding to the choice, analogous to a preset bound for the evidence accumulation process in the RT version of the task. In the interrogation design this stereotypical excitation was preceded by a persistent lower level of excitation, even in the absence of sensory stimulus (Roitman and Shadlen 2002). The stereotypical excitation could thus reflect the neural circuits reaching a motor activation threshold due to a transient gain modulation in the neural circuits at the time of choice commitment (Niyogi and Wong-Lin 2013). Thus, once the accumulated evidence in regions like LIP reaches this stereotypical bound, an action plan is generated in regions primarily responsible for execution of the action plan (like FEF) to signal choice-commitment. These regions are thus referred to as **high-level areas**, owing to their role in the decision formation, and execution.

The flow of decision-relevant information in the brain is not strictly **feedforward**. Recordings from visual cortex of monkeys making perceptual decisions showed that the activity of sensory neurons covaried with the monkey's choices (so-called choice probability). These correlations increased as a function of elapsed time in a trial, even when the influence of sensory stimulus on monkey's choice decreased (so-called **temporal weighting of evidence**) (Nienborg and Cumming 2010; Nienborg and Cumming 2009). Computational models using hierarchical networks of spiking neurons, suggested that decision-related information is fed back from high-level areas to sensory areas during the course of a trial, resulting in increased choice probability over time (Wimmer et al. 2015). These findings further provide strength to the idea that as sensory information is being integrated towards a choice, the **feedback** connections from high-level areas to low-level areas push the decision-making network into a stable attractor state, with new sensory information having little effect.

Functional imaging studies in humans corroborated and extended the findings above. Neural correlates of sensory encoding have been identified in area MT for random dot motion stimulus (Hebart et al. 2012; Siegel et al. 2007), and in somatosensory cortex for vibrotactile stimulus (Preuschhof et al. 2006; Tegenthoff et al. 2005). Choice selective signals have been identified in inferior parietal, posterior parietal and prefrontal regions when choice-response mappings were decoupled (Hebart et al. 2012, 2016), and

in dorsolateral prefrontal region (Heekeren et al. 2006; Philiastides et al. 2011) independent of response modality. When the choice-response mappings were fixed, choice predictive activity was also observed in motor cortex (Donner et al. 2009; de Gee et al. 2017; Kelly and O'Connell 2013; Wilming et al. 2020). Together, these results suggest that several parallels exist in the neural correlates of perceptual decisions between humans and non-human primates, establishing the generality of the decision-making hierarchy across species. Indeed, the simplistic nature of perceptual decision-making tasks makes them well suited to investigate and generalize the neural and computational mechanisms of perceptual decisions across species (Akrami et al. 2018; Brunton et al. 2013; de Gee et al. 2019; Lak et al. 2020).

In this thesis, I have used functional magnetic resonance imaging (fMRI) technique to characterise the neural mechanisms of perceptual decisions in humans. fMRI is a technique to measure the haemodynamic (blood flow) changes in brain regions using blood oxygenation level dependent (BOLD) contrast (Huettel et al. 2009). The BOLD contrast is used as a proxy to infer the (de)activation of a brain region in response to intrinsic, or experimentally controlled extrinsic variables (Logothetis et al. 2001). fMRI is an extremely popular technique among cognitive neuroscientists as it allows researchers to reliably identify brain areas underlying various cognitive processes with a high degree of spatial resolution, in a non-invasive fashion. A typical fMRI dataset contains a set of voxels, three dimensional cuboids with pre-set dimensions, measuring the BOLD signal of the brain tissue contained in the voxel. Each voxel has an associated time-course of BOLD signal measured at regular time intervals in a scanning session. Conventional statistical analyses characterise the relationship between cognitive variables and the activity of individual voxels: for example which voxels have a greater activation in response to the onset of sensory evidence compared to the pre-stimulus baseline interval. This approach to identify individual voxels showing differences in activity between two or more mental states (eg. trial engagement vs rest) is called a Univariate approach. In practice, univariate approaches are less sensitive to weak fluctuations in BOLD activity, making it difficult to find individual voxels where the differences between conditions is small, as is often the case. A more advanced approach combines activity from multiple voxels using pattern-classification techniques to identify fine-grained spatial patterns selective to different mental states, with increased sensitivity (Norman et

al. 2006). This approach is thus referred to as a **Multivariate** approach, or **multivoxel pattern analysis (MVPA)**. MVPA has been very successful as a method to identify distributed fine-grained patterns of representation sensitive to covert, and overt mental states in the brain (Haynes and Rees 2006). For example, using MVPA approaches, I've identified brain regions containing fine grained choice-specific representations in human subjects performing a perceptual decision-making task, as detailed in **chapter 4**.

1.4 CONFIRMATION BIAS IN DECISION-MAKING

Confirmation bias is the tendency to gather, recall or interpret new information such that it agrees with existing beliefs or hypothesis. This phenomenon was recognized by philosophers for a long time (usually through introspection and observation), with early descriptions dating back to 5th century BC, and is indeed one of the most widely documented biases of human reasoning (Nickerson 1998). Confirmation bias is pervasive, and can have disastrous consequences in certain real-world scenarios, like policy rationalization where policy makers ignore expert advice to justify their policies (Tuchman 1984), or when scientists resist discoveries that challenge their theoretical position (Barber 1961).

Empirical research showed that confirmation bias is an implicit process (Kahneman 2011), that people show this bias even for information that doesn't challenge beliefs they hold personal (Nickerson 1998; Oswald and Grosjean 2004). Early studies on confirmation bias used hypothesis testing, where human subjects initially formed a hypothesis (either through explicit instruction or inferred from observations), and were then provided additional evidence that supported or refuted it (Popper 1959; Wason 1960). These studies identified many forms in which humans exhibit confirmation bias: preferential sampling (Fischhoff and Beyth-Marom 1983), seeking (Zuckerman et al. 1995), biased recall (Kuhn 1989), or overweighting (Pyszczynski and Greenberg 1987) of information that supports their hypotheses, overconfidence in their judgments (Kahneman and Tversky 1973), and persistence of beliefs (Ross and Anderson 1982). Confirmation bias could arise as a result of suppression of post-decisional dissonance (Festinger 1957), when encountering evidence inconsistent with the initial hypothesis. More recent

studies investigated confirmation bias in the context of economic choices, and reinforcement learning paradigms. In free-choice paradigms, people assigned higher value to an item they have chosen from several equally desirable items (Chen and Risen 2010; Izuma and Murayama 2013). Confirmation bias modulated learning rates in standard reinforcement learning tasks, with higher learning rates for positive prediction errors of the chosen outcomes, and for negative prediction errors of the foregone outcomes (Chambon et al. 2019; Palminteri et al. 2017). Despite its prevalence, the mechanistic and neural basis of this bias has largely remained elusive, primarily because it is hard to measure under experimentally controlled conditions, and a bulk of empirical studies investigating this phenomenon used higher order cognitive paradigms whose neural mechanisms haven't been well established. Investigating confirmation bias in sensory decision-making tasks, by tightly controlling the decision-relevant evidence could illuminate the underlying mechanisms.

Conventional psychophysics experiments of perceptual decision-making adhere to a rigid trial structure, with each trial containing brief and discrete sequence of events, independent of all other trials (see Figure 1A). Such a structure stands in stark contrast to real life scenarios where decisions are embedded in sequence of other decisions. This discrepancy makes it harder to investigate confirmation bias in perceptual decisions. Decision-making researchers have started to address the gap between laboratory decisions and real-life decisions by developing novel psychophysics task protocols that go beyond the traditional design. Recent studies have probed the interaction between categorical choices and continuous estimations by combining discrimination and estimation judgments based on the same evidence (Jazayeri and Movshon 2006; Luu and Stocker 2018; Stocker and Simoncelli 2007). The biases in estimation judgments observed in such tasks were hypothesized to be a by-product of the sensory decoding being optimized to the discrimination judgment (Jazayeri and Movshon 2006, 2010) or due to the tendency of subjects to impose self-consistency between their discrimination and estimation judgments (Luu and Stocker 2018; Stocker and Simoncelli 2007). In numerical integration judgments, when additional evidence was presented between the discrimination and estimation judgments, subjects' showed reduced sensitivity to the new evidence (Bronfman et al. 2015). Taken

together, these studies provide a framework to develop novel task protocols to investigate confirmation bias in perceptual decisions.

1.5 TEMPORAL WEIGHTING OF DECISION-RELEVANT EVIDENCE

The source of sensory information in real life scenarios often changes abruptly and unpredictably. For example, while riding a bike in rain, the rider needs to constantly gather information *and* update their estimate about the source of a vehicle horn in order to avoid collision. Such scenarios require a flexible time scale of evidence accumulation, i.e., a delicate balance between gathering more information about the source of sensory evidence, and to discard old information. Experimental studies revealed that humans adjust to such changes in the source of signal by flexibly adapting their evidence accumulation time scale according to task demands (Glaze et al. 2015; Murphy et al. 2020; Ossmy et al. 2013).

However, conventional perceptual decision-making experiments assume stationarity of the signal in every trial (e.g. the direction, and proportion of coherent dots does not change during the trial). Subjects can maximize their performance in such tasks (correctly judge the stimulus direction), by giving equal weight to evidence in each sample during the course of a trial (Bogacz et al. 2006) i.e., a temporal weighting profile of evidence that is flat. The temporal weighting profiles can be constructed using **psychophysical kernels**, which quantify how sensory evidence at various time points in the trial affects the subjects' decision. Studies with human, and non-human subjects found conflicting results with the weighting of decision-relevant evidence. While some studies did find flat weighting profiles (Brunton et al. 2013; Raposo et al. 2014; Wyart et al. 2012), other studies found primacy, where evidence early in the trial had a greater weight in the decision (Kiani et al. 2008; Nienborg and Cumming 2009; Odoemene et al. 2018; Zylberberg et al. 2012); few studies found **recency**, where evidence late in the trial had a greater weight in the decision (Cheadle et al. 2014; Drugowitsch et al. 2016; Tsetsos et al. 2012); yet others found non-monotonic weighting profiles (Bronfman et al.





Figure 3: Illustration of different weighting profiles for decision-relevant sensory evidence. A multitude of factors influence whether subjects give more weight to evidence early in the trial (primacy), later in the trial (recency), or uniformly throughout the trial (flat weighting).

2016). Figure 3 illustrates these weighting profiles.

The discrepancy in the weighting profiles across studies has been attributed to various factors, some of which include the design of the experiment: interrogation design resulting in primacy vs reaction time design resulting in recency; stimulus presentation time: short duration resulting in primacy vs long duration resulting in recency; inter-subject differences for the same experimental design; and differences in sensory modalities (Bogacz et al. 2006; Bronfman et al. 2016; Cisek et al. 2009; Gold and Shadlen 2001; Thura et al. 2012; Tsetsos et al. 2012; Usher and McClelland 2001). Recent computational investigations accounted for these conflicting empirical findings by proposing variants to the standard accumulation to bound models: for example, changing the nature of the bound in DDM, or the relative strength of leak and inhibition in LCA. However, to fully understand the mechanisms underlying the relative weighting of decision-relevant evidence, we need a task design that can produce flexible temporal weighting profiles within each individual with minimal changes to the experimental variables. Such a task allows us to disentangle whether the changes in weighting occur due to task demands, or due to intrinsic changes in the state of the decision-making networks.



Figure 4: The psychometric function quantifying choice behavior across tens of trials (**A**) can be split into two separate psychometric functions conditioned on previous choice category (**B**), to visualise sequential bias as an offset between the two functions. This example shows a tendency to repeat the previous choice.

1.6 SEQUENTIAL BIAS IN PERCEPTUAL DECISIONS

Most perceptual decision-making studies treat trials as independent events, by randomizing the trial design (i.e., randomizing the identity of stimuli across successive trials) so as to study the response to each stimulus separately. However, it has been reported for a long time that past trials influence perceptual judgments in the current trials (Bertelson 1965; Cross 1973; Fernberger 1920). This phenomenon that perception is influenced not only by the current sensory information, but also by the sensory information and choices made in the recent past is called Sequential bias. Also called serial dependencies, choice hysteresis, or decision inertia, these biases recently received renewed interest among neuroscientists and psychologists, and have been observed in decision-making studies with humans (Fischer and Whitney 2014; Fritsche et al. 2017; Liberman et al. 2014), non-human primates (Gold et al. 2008), and rodents (Busse et al. 2011; Kiyonaga et al. 2017; Lak et al. 2020; Odoemene et al. 2018). Sequential biases are idiosyncratic (Braun et al. 2018; Fründ et al. 2014; Urai et al. 2019), adapt to the trial statistics (Abrahamyan et al. 2016; Braun et al. 2018), are driven by previous choices but not the motor commands used to report those choices (Akaishi et al. 2014; Braun et al. 2018), are influenced by decision confidence (Braun et al. 2018; Desender et al. 2018; Drugowitsch et al. 2019; Lak et al. 2020), and phasic arousal (de Gee et al. 2017, 2019; Urai et al. 2017), and are unaffected by task irrelevant stimulus presented between trials (Akaishi et al. 2014).

Sequential biases often go unnoticed in randomized trial designs when quantifying the average choice behavior across tens of trials (which is usually the case in perceptual decision-making investigations). Figure 4A shows a psychometric function of an example subject doing a two-alternative random dot motion direction discrimination task. Constructing separate psychometric functions conditioned on previous choice categories shows sequential bias as a horizontal offset between the two functions (Figure 4B). A more principled approach to quantify sequential bias is to model the contributions of individual components of previous trials namely previous choice, stimulus, or reward separately (Abrahamyan et al. 2016; Fründ et al. 2014). This approach made it easy to quantify sequential biases in perceptual decisions and offer a means to probe the role of biases and expectations in decision-making while retaining the control offered by perceptual decisions.

Recent studies have begun to characterise the neural mechanisms underlying sequential bias across different species. Neural correlates of sequential bias have been identified in early sensory (St John-Saaltink et al. 2016), parietal (Marcos et al. 2013; Scott et al. 2017), prefrontal (Akaishi et al. 2014), and motor (de Lange et al. 2013; Pape and Siegel 2016) regions. Optogenetic inactivation of the posterior parietal cortex in mice eliminated sequential bias in choices suggesting a causal role for this region (Akrami et al. 2018; Hwang et al. 2017), while neural correlates of this bias in early sensory regions could reflect choice-related feedback signals from higher level decision areas (Macke and Nienborg 2019; Wimmer et al. 2015). These results highlight the heterogenous nature of the format and locus of sequential bias signals in the brain. A comprehensive characterisation of bias signals across the cortex is yet to be addressed.

1.7 OUTLINE OF THE THESIS

It is now well-established that choices are not endpoints in perceptual decisionmaking tasks, but continue to influence subsequent evidence. Such biases possibly reflect the assumptions that people hold in their brains about the statistical regularities of the world around them. Thus, a comprehensive characterization of the mechanistic and neural basis of these biases in laboratory tasks could provide insights into the fundamental principles of decisionmaking, which is the central theme of this thesis.

Chapter 2 provides insights into the mechanistic basis of confirmation bias. Using novel task protocols, and extensive computational model-based and model-free analyses, this chapter shows that human subjects exhibit confirmation bias by assigning greater weight to new perceptual information that was consistent with their intermittent binary choice, compared to an inconsistent stimulus.

Chapter 3 characterizes the impact of intermittent binary choices on the accumulation of protracted decision-relevant evidence. We find that intermittent overt choices dynamically alter the processing of the evidence, consistent with an active internal state change triggered by choice-commitment. The generality of the findings in **chapters 2 & 3** were established by analyzing data from two different judgments- perceptual choices about the direction of random dot motion stimulus, and numerical integration judgments about the mean of a protracted stream of numbers.

Chapter 4 investigates the neural basis of sequential biases in perceptual decision-making, using a combination of computational modelling, and multivoxel pattern analysis of neuroimaging data collected using fMRI. We identified distinct networks of cortical regions containing transient choice representations at a single trial level, and sustained choice representations in the inter-trial intervals. We mapped behaviorally quantified sequential biases onto the neural activity patterns in these regions, and isolated behaviorally relevant effects from these patterns.

2

CONFIRMATION BIAS THROUGH SELECTIVE OVERWEIGHTING OF CHOICE-CONSISTENT EVIDENCE

PEOPLE'S ASSESSMENTS OF THE STATE OF THE WORLD OFTEN DEVIATE SYSTEMAT-ICALLY FROM THE INFORMATION AVAILABLE TO THEM (TVERSKY AND KAHNEMAN 1974). SUCH BIASES CAN ORIGINATE FROM PEOPLE'S OWN DECISIONS: COMMIT-TING TO A CATEGORICAL PROPOSITION, OR A COURSE OF ACTION, BIASES SUB-SEQUENT JUDGMENT AND DECISION-MAKING. THIS PHENOMENON, CALLED CON-FIRMATION BIAS (NICKERSON 1998), HAS BEEN EXPLAINED AS SUPPRESSION OF POST-DECISIONAL DISSONANCE (BREHM 1956; FESTINGER 1957). HERE, WE PRO-VIDE INSIGHTS INTO THE UNDERLYING MECHANISM. IT IS COMMONLY HELD THAT DECISIONS RESULT FROM THE ACCUMULATION OF SAMPLES OF EVIDENCE INFORM-ING ABOUT THE STATE OF THE WORLD (BOGACZ ET AL. 2006; GOLD AND SHADLEN 2007; LIU AND WANG 2008; WANG 2008). WE HYPOTHESIZED THAT CHOICES BIAS THE ACCUMULATION PROCESS BY SELECTIVELY ALTERING THE WEIGHTING (GAIN) OF SUBSEQUENT EVIDENCE, AKIN TO SELECTIVE ATTENTION. WE DEVELOPED A NOVEL

* shared first authors.

doi: 10.1016/j.cub.2018.07.052. This chapter was reprinted under a CC-BY-NC-ND 4.0 license. AUTHOR CONTRIBUTIONS: T.H.D., K.T., and M.U. formulated the idea for the study and designed the experiment. A.E.U., T.H.D., and M.U. specified the design for the perceptual tasks. A.E.U. programmed the task and collected data for the perceptual task. B.C.T., A.E.U., and K.T. analyzed data. T.H.D., M.U., and K.T. supervised the study. B.C.T., A.E.U., and T.H.D. wrote the paper. All authors revised and approved the manuscript.

PSYCHOPHYSICAL TASK TO TEST THIS IDEA. PARTICIPANTS VIEWED TWO SUCCESSIVE RANDOM DOT MOTION STIMULI AND MADE TWO MOTION-DIRECTION JUDGMENTS: A CATEGORICAL DISCRIMINATION AFTER THE FIRST STIMULUS AND A CONTINUOUS ESTIMATION OF THE OVERALL DIRECTION ACROSS BOTH STIMULUS AND A CONTINUOUS STIMULUS. PARTICIPANTS' SENSITIVITY FOR THE SECOND STIMULUS WAS SELECTIVELY ENHANCED WHEN THAT STIMULUS WAS CONSISTENT WITH THE INITIAL CHOICE (COM-PARED TO BOTH, FIRST STIMULI AND CHOICE-INCONSISTENT SECOND STIMULI). A MODEL ENTAILING CHOICE-DEPENDENT SELECTIVE GAIN MODULATION EXPLAINED THIS EFFECT BETTER THAN SEVERAL ALTERNATIVE MECHANISMS. CHOICE-DEPENDENT GAIN MODULATION WAS ALSO ESTABLISHED IN ANOTHER TASK ENTAILING AVERAG-ING OF NUMERICAL VALUES INSTEAD OF MOTION DIRECTIONS. WE CONCLUDE THAT INTERMITTENT CHOICES DIRECT SELECTIVE ATTENTION DURING THE EVALUATION OF SUBSEQUENT EVIDENCE, POSSIBLY DUE TO DECISION-RELATED FEEDBACK IN THE BRAIN (WIMMER ET AL. 2015). OUR RESULTS POINT TO A RECURRENT INTERPLAY BETWEEN DECISION-MAKING AND SELECTIVE ATTENTION.

2.1 RESULTS

Brain regions implicated in evidence accumulation, decision-making, and attentional control maintain their activity states over long timescales and send feedback to regions encoding the incoming evidence (Nienborg and Cumming 2009; Siegel et al. 2015; Wimmer et al. 2015). We thus reasoned that the consistency of new evidence with a previous choice might affect the decision-makers' sensitivity to the new evidence. Specifically, we hypothesized that a categorical choice induces a multiplicative gain modulation of new evidence, selectively boosting the sensitivity to consistent evidence. Such a selective gain modulation is commonly observed when explicit cues direct feature-based attention (Herrmann et al. 2012; Maunsell and Treue 2006; Reynolds and Heeger 2009).

Previous studies have identified gain modulations in evidence accumulation by presenting multiple samples of evidence in succession and asking participants to report a binary choice based on the mean evidence at the end of the sequence (Drugowitsch et al. 2016; Tsetsos et al. 2012, 2016; Wyart et al. 2012). Those studies did not assess the effect of intermittent choices in biasing the accumulation process. Other work has probed the interaction between categorical choices and continuous estimations by combining discrimination and estimation judgments based on the same evidence presented before (Jazayeri and Movshon 2007; Luu and Stocker 2018; Stocker and Simoncelli 2007; Zamboni et al. 2016). Here, choice-related estimation biases may be a by-product of the bottom-up sensory decoding (i.e., weighting of sensory neurons) being tailored to the discrimination judgment (Jazayeri and Movshon 2007;but see Luu and Stocker 2018; Stocker and Simoncelli 2007). Whether a categorical choice occurring during a protracted stream of decision-relevant evidence selectively modulate the gain of evidence subsequent to that choice, has remained unknown. We addressed this question by combining the above two approaches.

Our task required participants to report a continuous estimate of the overall motion direction across two successively presented random dot motion stimuli. In the majority of trials, participants were also prompted to report a binary categorical judgment after the first stimulus (see Figure 1A and Section 2.3): discriminating whether its direction was clockwise (CW) or counter-clockwise (CCW) with respect to a reference line. Importantly, the stimulus following the intermittent choice contributed only to the final estimation, but not to the discrimination judgment. This psychophysical protocol enabled us to isolate the impact of an intermittent categorical choice on decision-makers' sensitivity to subsequent evidence for continuous estimation.

Participants made use of the stimulus information for both judgments: The fraction of CW choices increased as a function of the direction of the first stimulus from the reference (Figure 1B; Figure S1A), and continuous estimations scaled with the mean stimulus direction across both intervals (Figure 1C, top; Figure S1 A, B). The estimations were generally attracted towards the reference (Figure 1C, top, compare black and grey dashed line), in line with the non-uniform distribution of the mean stimulus directions (Figure 1C, bottom; Section 2.3).

Post-decisional selective gain modulation predicts that evidence subsequent to a choice produces larger (smaller) deviations in the overall estimations when these new directions are consistent (inconsistent) with that choice. We used two complementary approaches to test this prediction. The first approach modeled the overall estimations as a noisy weighted average of the



Figure 1: A. Schematic sequence of events within a trial. A first dot motion stimulus was shown on all trials for 750 ms and then paused. On two-thirds of trials, an auditory prompt instructed a direction discrimination judgment (CW or CCW with respect to reference line, at 45° in this example trial) as shown here. A third of trials, not analyzed here, did not require a choice. After half of the discrimination judgments, feedback was given and the trial terminated. After the other half, a second motion stimulus was presented (equal coherence as first, but independent direction) and participants were asked to estimate the mean direction of both stimuli. **B.** Proportion of CW choices as a function of stimulus direction, along with psychometric function fit. C. Top, continuous estimations as function of mean direction across both stimuli. Bottom, distribution of mean directions across trials. Black, data; blue, predictions generated from best fitting parameters of Choice-based Selective Gain model. Data points, group mean; error bars, s.e.m. Gray, predictions by Extended Conditioned Perception model under several levels of output noise for average subject. Stimulus directions and estimations were always expressed as the angular distance from the reference, the position of which varied from trial to trial (0° equals reference). See also main text, Section 2.3, and Figure S1.

directional evidence in both stimulus intervals (see methods). The weight for each stimulus quantified its gain in the estimation process. Trials with second stimulus directions consistent or inconsistent with the choice were modeled separately. This model, referred to as the 'Choice-based Selective Gain' model in the following (Section 2.3), provided a good account of observers' estimation reports (Figure 1C, Figure 2). Smaller values of Bayes Information Criterion (BIC) within the majority of individual participants indicated that Choice-based Selective Gain explained the data better than a 'Baseline' model without choice-dependent change in evidence weighting (Figure 2A



Figure 2: A. Comparison between Choice-based Selective Gain and alternative models. Negative values are evidence for Choice-based Selective Gain. Gray shade, $|\Delta BIC| > 10$, indicating very strong evidence for model with smaller BIC (Section 2.3). B. Model weights for second stimulus in Consistent and Inconsistent conditions. Error bars, 66% bootstrap confidence intervals. Black cross, mean and s.e.m. Dashed line, Consistent=Inconsistent, and data points above dashed line, Consistent>Inconsistent. C. Mean model weights for both stimulus intervals in Consistent and Inconsistent. Error bars, s.e.m, F-statistic, interaction between interval and condition (2-way ANOVA). D. Difference between effect strength (difference: Consistent-Inconsistent) for second stimulus, in weights obtained from Choice-based Selective Gain and Stimulusbased Selective Gain models. E. ROC indices for second Consistent and Inconsistent stimulus, predicted by simulations of alternative models as indicated above (individual trial distributions and best fitting model parameters). F. As D, but for measured data. Data points in all panels but **C** are participants, with identical color scheme. Pvalues, permutation tests (100.000 permutations) comparing weights or ROC-indices between Consistent and Inconsistent across participants (N=10). See also Figure S2.

and Section 2.3). Further, choice-consistent second stimuli received larger weight than choice-inconsistent second stimuli (Figure 2B; see Figure S2A for noise estimates). This weight difference was not evident for the first stimuli (Figure 2C). Indeed, weights were increased compared to the first stimulus for choice-consistent second stimuli, and reduced for choice-inconsistent stimuli (Figure 2C). In sum, observers prioritized choice-consistent evidence after the categorical choice, in a way resembling feature-based attention.

The second, complementary approach corroborated this conclusion (Figure 2E, F). We developed a model-free measure based on the receiveroperating characteristic (ROC) that quantified the sensitivity to the second stimulus. ROC-indices measured the extent, to which single-trial estimations separated between second stimuli of nearby directions (i.e., 10° vs. 20°, or -10° vs. -20°; see Section 2.3 for details). Simulations confirmed that the difference between these ROC-indices, computed separately for choice-consistent and choice-inconsistent stimuli, captured the choice-dependent gain modulation described by the Choice-based Selective Gain model (Figure 2E, left panel; Figure S2B). Critically, for the actual data, ROC-indices were larger for the Consistent than Inconsistent condition (Figure 2F). In sum, also the model-free analysis revealed a selective modulation of sensitivity to additional evidence, in line with feature-based attention.

This consistency-dependent change in sensitivity for subsequent evidence, as quantified by the ROC-indices, could not be explained by other mechanisms lacking multiplicative gain modulation. In a first alternative model, biases shared among choice and subsequent estimations resulted from slow fluctuations in noise corrupting both judgments, without any genuine effect of the choice. This so-called 'Correlated Noise' model (Section 2.3) provided a worse account of estimation reports (in 9/10) than Choice-based Selective Gain (Figure 2A) and could not produce the consistency-dependent ROC effect: neither for the individually fitted parameters (Figure 2E, middle panel), nor for any combination of parameters that we simulated (Figure S2B).

In a second alternative model, the initial choice shifted the internal representation of the evidence towards the chosen category in an additive fashion. This 'Shift' model (Section 2.3) also produced systematic estimation biases and accounted well for the overall estimation behavior (Figure 2A). The shift parameter was larger than zero (p = 0.038, 2-sided permutation test), indicating that participants may have shifted their decision variable in the direction of the chosen category. The shift parameter was even significant (p = 0.05, 2-sided permutation test) for an Extended Choice-based Selective Gain model, which contained an extra free parameter for the shift (all other parameters constrained from the Choice-based Selective Gain model fits, Section 2.3; Figure S2F). But critically, the Shift model also could not capture the specific behavioral feature that was diagnostic of selective gain modulation: the consistency-dependent sensitivity change (Figure 2E, right panel) as was evident in the data (Figure S2B). It is possible that an additive shift and multiplicative gain modulation jointly governed choice-induced biases in the overall estimation behavior (see Section 2.2).

Taken together, the analyses presented so far indicate that consistencydependent gain modulation was necessary to account for certain features of participants' behavior. Further analyses indicated that this gain modulation was, in fact, induced by the intermittent choice (i.e., participants' categorization of the first stimulus), rather than by the first stimulus itself (Figure S2C, D) or by the disparity between first and second stimulus (Figure S2E). We fitted a variant of the Selective Gain model, in which the consistency of the second stimulus was defined based on the first physical stimulus direction, rather the participants' choice (Section 2.3). This so-called Stimulus-based Selective Gain model provided a worse account of the data than the Choice-based Selective Gain (Figure 2A). Critically, the selective gain effect was larger for the parameters estimated by Choice-based Selective Gain model (Figure 2D). In sum, the selective modulation in sensitivity was linked to the participants' categorical choice.

A recent Bayesian account of post-decision biases has proposed that perceptual inference is 'conditioned' on choice in order to ensure consistency between binary discrimination and continuous estimation judgments of the same stimulus (Luu and Stocker 2018; Stocker and Simoncelli 2007). This account is framed at a different level of description (Bayesian inference), but the notion of a choice-dependent prior for estimation is similar to our idea of a choice-induced top-down modulation. Could choice-based conditioning of internal representations explain the present results? Our task and analyses isolated the impact of binary choice on the processing of subsequent evidence for continuous estimation, requiring additional assumptions about the conditioning operation. If only the representation of the first stimulus was conditioned, this would yield an offset of the representation of the second stimulus - equivalent to the Shift model considered above, which did not account consistency-effect on ROC indices observed in the data (Figure 2E, right panel; Figure S2B). If also the representation of the second stimulus was conditioned on the choice (referred to as 'Extended Conditioned Perception', see Section 2.3), this reproduced the ROC-effect (Figure S2B, rightmost panel). However, the later model did not account well for the relationship between overall estimations and mean stimulus direction (gray lines in Figure 1C; for further comparison between Extended Conditioned Perception and Choice-based Selective Gain, see Figure S2G, H). Future work should develop biologically plausible and dynamic approximations of choice-based



Figure 3: **A.** Schematic sequence of events within a trial entailing intermittent binary choice. After the first sequence of eight numbers, participants discriminated the mean as larger or smaller than 50 (a quarter of trials, not analyzed here, did not require a choice (see Section 2.3 and (Bronfman et al. 2015)). Following discrimination report, the trial terminated with feedback (two-thirds of trials) or a second sequence of eight numbers was presented (mean independent from that of first interval). Participants were then asked to report the mean of the whole number sequence. **B**, **C.** Model-based (**B**) and model-free (**C**) measures of consistency-dependent sensitivity modulation (as Figure 2B, F). **D.** Correlation between consistency effect in model-based and model-free analyses across participants from both tasks. Effect strength is Consistent - Inconsistent difference in model weights or ROC. Data points, participants. P-values in (**B**), (**C**) from permutation tests across participants (100.000 permutations; N = 21). See also Figure S3.

conditioning operation, in order to unravel possible links to choice-dependent gain modulation.

The post-decisional biasing effect in the visual perceptual task resembled well-documented effects in reasoning (Nickerson 1998) and preference reports (Brehm 1956; Chen and Risen 2010). It is unknown, however, whether the latter high-level post-decision biases are mediated by selective gain modulations akin to attention. To test for this, we re-analyzed and modeled previously published (Bronfman et al. 2015) data from a numerical averaging task that also required the combination of evidence presented before and after a choice into an overall estimation (Figure 3A; see Section 2.3 for task and analysis details). Again, the weights were larger on Consistent than Inconsistent conditions, specifically for evidence after choice (Figure 3B) again with an interaction between interval and consistency (Figure S3D). Likewise, the ROC indices were also larger for Consistent than Inconsistent conditions (Figure 3C). In sum, the choice-induced biasing mechanism we uncovered for perceptual decision-making, including the selective gain modulation, also accounts for post-decision biases in higher-level decisions based on numerical evidence.

2.2 DISCUSSION

Decision-makers are often systematically influenced by their own choices: Committing to a categorical hypothesis, or choosing a course of action biases the subsequent evaluation of the decision-relevant evidence (Brehm 1956; Nickerson 1998). The mechanisms underlying such post-decisional confirmation biases have so far remained unknown. Here, we have shown that choices selectively increased the gain of subsequent evidence that was consistent with that choice, for perceptual as well as numerical decisions. A selective modulation of the gain of sensory responses is commonly observed during attention to certain features of the evidence (Herrmann et al. 2012; Maunsell and Treue 2006; Reynolds and Heeger 2009). In sum, our results illuminate the linkage between decision-making and attention – two capacities commonly studied in isolation, but interacting in real-life behavior. Our findings indicate that an agent's decision acts like a cue for selective attention, biasing subsequent decision processing.

Evidence inconsistent with an initial choice may induce post-decisional dissonance, possibly related to conflict between competing cognitive states or motor responses (Festinger 1957; van Veen et al. 2009). Previous work has shown that such conflict boosts top-down control, increasing task per-

formance and response caution on subsequent trials (Botvinick et al. 2001; Miller and Cohen 2001). But this line of work has not associated conflict with subsequent decision biases. In particular, it has not shown that conflict induces selective modulations of new information that is consistent with respect to a previous choice.

We have recently established that sensitivity for new information is generally reduced after an overt choice, compared to no overt choice (Bronfman et al. 2015). To this end, we assessed a non-selective reduction in sensitivity for any post-decision evidence. Our current work goes beyond this by uncovering a selective mechanism of confirmation bias: preferentially sampling the evidence that confirms one's prior belief. This effect indicates a more refined mechanism than the non-selective reduction in overall sensitivity due to an overt choice. Identifying this effect was afforded by an improved modeling approach (see Section 2.3) combined with a model-free behavioral readout of selective gain modulation (ROC-analysis), both yielding consistent results (Figure 3D). It is conceivable that a non-selective gain reduction due to overt choice (possibly reflecting reduced arousal and/or cortical attractor dynamics (Bronfman et al. 20155)) and selective attention towards choice-consistent evidence conspire to shape overall estimation behavior.

Our analysis of the perceptual task revealed, in some of the participants, also an additive shift in the direction of the chosen category, on top of the gain modulation. This additive shift may reflect previously identified choice-induced biases (Jazayeri and Movshon 2007). This additive shift could not, however, account for the consistency-dependent change in sensitivity (Figure 2E), which we found in the data (Figure 2F). The co-existence of additive and multiplicative effects may relate to the observation that common manipulations of selective attention produce effects on both, sensitivity and decision criteria, which are dissociable at behavioral and neural levels (Luo and Maunsell 2015, 2018). Our present experimental manipulation does not allow for distinguishing between an additive baseline shift in the sensory response and a shift of the starting point of the decision variable accumulating the sensory response. Future experiments could manipulate the duration of the second evidence to dissociate these two scenarios.

Our work contributes to recent progress in the understanding of historydependent biases in perceptual choice (Abrahamyan et al. 2016; Akaishi et al. 2014; Akrami et al. 2018; Braun et al. 2018; Fernberger 1920; Fritsche et al. 2017; Fründ et al. 2014; Kim et al. 2017; de Lange et al. 2013; Pape and Siegel 2016; Urai et al. 2017). One class of mechanism contributing to such biases is stimulus-selective adaptation, which can cause repulsion (Fritsche et al. 2017) or attraction (Brascamp et al. 2008; Cheadle et al. 2014; Kanai and Verstraten 2005; Pearson and Brascamp 2008), possibly owing to adaptation dynamics at different processing levels. Low-level adaptation after prolonged stimulus exposure as in our task (Figure 1A) commonly produces repulsive effects (Fritsche et al. 2017), due to suppressing sensory cortical responses (Kohn 2007; Krekelberg et al. 2006). This is inconsistent with our results because it predicts stronger sensitivity loss for congruent than incongruent stimuli. Higher-level adaptation mechanisms can cause attraction, especially in the face of ambiguous evidence (Brascamp et al. 2008; Kanai and Verstraten 2005; Pearson and Brascamp 2008), and has been linked to gain modulation induced by the stimulus sequence (Cheadle et al. 2014). We found that the consistency-dependent gain modulation was more strongly tied to observers' choices than the physical stimuli, implying a higher-level source. This, combined with the multiplicative nature of the effect, naturally links it to feature-based attention. Whether a local adaptation mechanism with such functional characteristics exists remains to be tested.

History biases in perceptual choice tasks requiring categorical judgments have specifically been linked to the history of previous choices (Akaishi et al. 2014; Bonaiuto et al. 2016; Braun et al. 2018; Fritsche et al. 2017; Urai et al. 2017) or choice outcomes (Abrahamyan et al. 2016; Busse et al. 2011). While these across-trial biases are idiosyncratic (Abrahamyan et al. 2016; Urai et al. 2017), the predominant tendency is to repeat choices more often than expected by chance (Akaishi et al. 2014; Braun et al. 2018; Fritsche et al. 2017; Urai et al. 2017), in line with the current choice-consistency bias established here within a single protracted decision process. Recent work on across-trial history biases in categorical choice points to a similar attentional mechanism giving rise to choice-repetition biases across trials (Urai et al. 2019).

The accumulation of fluctuating sensory evidence towards binary choices is well characterized at a neurophysiological level (Bogacz et al. 2006; Brunton et al. 2013; Ossmy et al. 2013; Wang 2008). Theoretical work points to an analogous mechanism underlying continuous decisions (Liu and Wang 2008). While less is known about continuous decisions based on two successive evidence streams, it is tempting to speculate that the selective re-weighting effect results from top-down feedback from cortical accumulator regions to regions that encode the evidence (Goris et al. 2017; Nienborg and Cumming 2009; Siegel et al. 2015; Wimmer et al. 2015). Such feedback interactions might alter the decision-maker's interpretation of incoming information by the evolving belief state (Haefner et al. 2016; Nienborg and Roelfsema 2015).

Our results have broader implications. First, insight into the computational mechanisms producing confirmation biases has considerable ecological value because these biases are pervasive in daily life, shaping human judgment in cases of critical significance (e.g. scientific hypothesis testing) (Nickerson 1998). Second, our work sets the stage for probing into the neural mechanisms of confirmation biases, in humans and animal models. Previous work into confirmation bias has focused on high-level judgment and reasoning (Nickerson 1998), the neural bases of which remain elusive. By contrast, neuroscience has accumulated substantial knowledge about the neural signals that encode the sensory evidence and evolving decision about visual motion (Shadlen and Kiani 2013; Wang 2008). The modulation of visual motion signals by attention is also well characterized (Maunsell and Treue 2006). Our current findings establish an analogous biasing mechanism in both domains high-level judgment and perceptual decisions - along with an effective behavioral readout and computational signature that constrains for the candidate neural mechanisms.

2.3 METHODS

2.3.1 Experimental model and subject details

Data from sixteen participants (six men and ten women) between the ages of 18 and 29 were collected for this study. Two participants did not complete the full experiment and were discarded from all analyses. The estimations of some subjects did not increase monotonically, quantified by the slope of best-fitting line, as a function of mean direction (red boxes in Figure S1A). We excluded four participants for whom the slopes were <0.3 (Figure S1B), and the results in Figures 1, 2 are based on the remaining 10 participants. All gave written informed consent prior to participation, and were naive to the aim of the experiment. The University of Amsterdam ethics review board approved the project. Each participant performed a total of 12 sessions, distributed across six days: One session to determine the motion coherence of the stimuli that corresponded to the individual psychophysical threshold and 11 sessions of the main experimental task. Each session of the main task consisted of 345 trials, divided into five experimental blocks of 69 trials. We used the first two sessions (690 trials) as training sessions to get participants acquainted to the task. We also re-analyzed previously collected data (Bronfman et al. 2015) from an additional 21 participants (age range: 21 to 29). In this data, after a short block of 20 practice trials, each participant completed 300 trials (5 blocks of 60 trials each).

2.3.2 Method details: Perceptual task

Stimuli

Stimuli were presented using PsychToolbox-3 (Kleiner et al. 2007) in MATLAB and were viewed in a dark, quiet room on a CRT monitor with a resolution of 1024 pixels x 768 pixels and a refresh rate of 60 Hz. Participants placed their heads on a chin-rest with a viewing distance of 50 cm from the screen. Dynamical random dot stimuli were presented in a central circle (outer radius 12°, inner radius 2°) around fixation. A field of dots with a density of 1.7 dots/degrees2 defined the annulus. Dots were 0.2° in diameter and were white, at 100% contrast from the black screen background (see Figure 1A). Signal dots were randomly selected on each frame and moved with 11.5°/second in the signal direction. Signal dots that left the annulus wrapped around and reappeared on the other side. Moreover, signal dots had a limited "lifetime", and were re-plotted in a random location after being on the screen for four consecutive frames. Noise dots were assigned a random location within the annulus on each frame, resulting in 'random position' (white) noise with a 'different' rule (Scase et al. 1996). Additionally, to avoid participants tracking individual signal dots as they move through the annulus, three independent

motion sequences were interleaved on subsequent frames (Roitman and Shadlen 2002), making the effective speed of dots 3.8° /second.

Procedure: Determining individual motion coherence thresholds

On the first day, participants were provided initial instructions about the task and performed a separate session in order to determine the individual motion coherence level for the main experiment. Individual participants' motion coherence thresholds were determined on a coarse (up vs. down) direction discrimination task. 600 trials of different motion strengths (0, 2.5, 5, 10, 20 and 40% coherence) were randomly interleaved (duration: 750 ms). For each participant, we fit a cumulative Weibull function to the proportion of correct choices as a function of motion coherence *c*:

$$P(correct|c) = \delta + (1 - \delta - \gamma)(1 - e^{(-c/\alpha)^{\beta}})$$
(2.1)

where δ was the guess rate (chance performance), γ was the lapse rate, and α and β were the threshold and slope of the psychometric Weibull function, respectively. While keeping the guess rate δ fixed at 50% correct, we fit the parameters γ , α and β maximizing the likelihood function (Wichmann and Hill 2001) using a Nelder-Mead simplex optimization algorithm. The individual threshold was taken as the stimulus difficulty corresponding to an 80% correct fit of the cumulative Weibull. Across participants, motion coherence thresholds ranged from 11% to 28% (mean 18%).

Procedure: Main experiment

Each trial had lasted for about 5 s, throughout which a red fixation mark was presented, followed by a black screen in the inter-trial interval. Participants self-initiated the next trial by pressing a mouse button. Within each trial, two random dot motion stimuli were presented in succession, each with independently chosen direction (Figure 1A) and an individually titrated near-threshold coherence levels (see previous section). In addition, auditory signals were presented prompting the participants' responses or providing feedback (see below). Each trial began with a blank fixation period (600-800 ms, uniform distribution), followed by the first motion stimulus (750 ms) during which the signal dots moved in one of five directions relative to a reference mark (see below). The reference mark was a white line in the circle,
with randomly changing position from trial to trial. Following the offset of the first dot motion stimulus, one of two tones prompted participants to either click the central mouse wheel (No-Choice trials) or the left and right mouse button to report a CW vs. CCW choice (Choice trials). After half of the Choice trials, participants received auditory feedback about the correctness of their choice (assigned randomly for 0° stimuli) and the trial ended. In the remaining trials, a second dot motion stimulus was presented for 750 ms. The delay between the first and second dot motion stimulus was always 2 s, regardless of reaction time. After the offset of the second stimulus the reference mark turned red, prompting participants to estimate the average motion direction across both dot motion stimuli. They reported their estimate by dragging the red line around the circle, starting from the position of the reference, and by then clicking the mouse at the endpoint.

For each participant, the reference position was constrained to be either within the top (0°-180°) or the bottom half (180°-360°) of the stimulus unit circle (balanced across participants) in order to keep the mapping between CW/CCW choices and left/right button presses constant within each participant. There were five possible directions (-20° , -10° , 0° , 10° , 20°) of each dot motion stimulus, yielding 25 possible combinations of directions across both subsequent stimuli. Of those, only 23 were used, excluding the two most obviously conflicting combinations ($-20^{\circ}/20^{\circ}$ and $20^{\circ}/-20^{\circ}$). The resulting distribution of mean directions was non-uniform and bi-modal (Figure 1C). Feedback about their estimation performance was given at the end of each block as the mean deviation across trials of their estimation reports from the physical stimulus directions ¹.

In total, 90 trials for each combination of first and second directions were presented per participant (45 in Choice and 45 in No-Choice trials). Trials were excluded from analysis according to the following criteria: (i) participants did not comply with the instructions (i.e., pressing the mouse wheel on Choice trials or a choice key on No-Choice trials); (ii) binary choice reaction time was below 200 ms (i.e., shorter than regular reaction times on two-choice tasks); and (iii) estimation outliers (defined as estimations beyond 1.5 times

¹ A video demonstration of the task can be found online with the article at https://www.cell.com/cms/10.1016/j.cub.2018.07.052/attachment/ 1194906f-486c-47fd-8668-5f39d21998da/mmc2.

the interquartile range, above upper or below lower quartile). In total \approx 7% of the total trials across the 10 participants were excluded. In addition, we excluded all No-Choice trials from our analyses as we focus only on Choice trials here. The distributions of the remaining trials used for analysis are shown in Figure S1C, D.

Task instructions

Instructions for the perceptual task were provided to participants before the start of the experiment as a written numbered list with graphics. Below, we provide an abbreviated version of all points from that list:

- Every block consists of several trials of the same visual motion task. Always keep your gaze on the red fixation point in the center of the screen.
- 2. Blank screen: Each trial will begin with a black screen. The red fixation appearance indicates that the trial is about to start.
- 3. Interval 1: You will see a cloud of dots moving, with some of the dots moving together in a particular direction. Your task is to determine whether the dots are moving to the left or to the right of the reference mark.
- 4. Binary Response: Once the dots stop moving, you will hear an auditory prompt to report your decision about the direction by clicking the corresponding mouse key (left or right). Try your best to make this decision as quickly and accurately as possible.
- 5. After your response, trials will continue with either:
 - a) Feedback: Once you've pressed a mouse key, you will hear feedback about your response in some trials. A correct choice will be followed by a high beep, and an incorrect choice will be followed by a low beep. Following feedback, you will move on to the next trial.

OR

b) Interval 2: After your response, you will see a second cloud of dots moving, with some of the dots moving together in a particular direction. These dots may have a different angle of motion from the first stimulus. Your task is to determine and estimate what the average overall angle of motion is from this cloud and the first one combined.

- 6. Estimation Response: When the dots stop moving and the reference mark will turn red, you must complete an estimation task. Move the mouse to align the cursor to the average angle of motion you saw in the two trials. Once you are satisfied with your estimate, click the mouse to confirm your response.
- 7. When you see a blank screen, the trial is over and you will have the opportunity for a break. Rest your eyes for a moment to help you keep them open and fixated during the experimental trials. When you want to continue the experiment, click the mouse to continue.
- 8. After each block of around 10 minutes, you will see a screen indicating your performance on the last block and telling you to take a break.

2.3.3 Method details: Numerical task

The task was identical to the perceptual task described above, with the following exceptions. Participants viewed two sequences of eight two-digit numbers each and reported the mean of all 16 samples as a continuous measure. Each sequence lasted for 4 s, and each numerical sample was displayed for 500 ms. In 75% of all trials, prompted by a visual cue, subjects made a binary choice about the mean of the first sequence of numbers- whether the mean was greater than or less than 50 by pressing the corresponding key. In a proportion of these trials (25% of all trials), the binary response was followed by a second sequence of eight numbers after which subjects made the estimation judgment by vertically sliding a mouse-controlled bar set on a number ruler between 0 and 100. Numbers were generated from pre-defined distributions ranged between 10 and 90. The data we analyzed in this paper constituted 25% of trials from each subject (\approx 75 trials). Please see the original report on this data set (Bronfman et al. 2015) for a more detailed description of the task, and https://datadryad.org/resource/doi: 10.5061/dryad.40f6v for the behavioral data.

2.3.4 Modeling discrimination judgments

Performance on the binary choice task in both datasets was quantified by fitting a sigmoidal probit psychometric function (Figure 1B) to each participant's proportion of CW choices (> 50 choices in numerical integration task), as a function of the stimulus direction (trial-wise mean of 8 samples in numerical integration task) in interval 1:

$$P(Choice = CW) = \Phi(\delta + \alpha \phi_1)$$
(2.2)

where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{(-t^2/2)} dt$ was the cumulative Gaussian function, α was the slope of the psychometric function (i.e., perceptual sensitivity), and δ was the horizontal shift of the psychometric function (i.e., systematic bias towards one of the two choice options; see Figure 1B). The inverse of the parameter α quantifies the internal sensory noise. The free parameters α and δ were estimated by maximum likelihood optimization (Wichmann and Hill 2001).

In the numerical task, the means of samples from the first interval exhibited substantial trial-to-trial variability. In order to compute psychometric functions, we binned trials by sample means into six bins, three on each on either side of the reference (50). We used the bin means as input to the psychometric function.

2.3.5 Modeling estimation reports

We used a statistical modeling approach to estimate the relative contributions of the evidence conveyed by both successive dot motion stimuli, or number streams, to participants' trial-by-trial estimation reports. All models described in this section were fit exclusively to the Choice trials. A comparison of Choice and No-choice trials was beyond the scope of this study; it can be found in (Bronfman et al. 2015) for the numerical task data and will be subject to a subsequent report for the perceptual task data.

Baseline model

As reference for assessing the importance of choice-related biases in the measured estimation data, we designed a Baseline model that did not entail any choice-related bias, but only participants' overall directional bias (estimated from the psychometric function, see below), as well as possible temporal biases in the combination of the two stimulus samples into the final estimation. The Baseline model was as follows:

$$y = w_1 X_1 + w_2 X_2 + N(0, \xi)$$
(2.3)

where y was the vector of single-trial estimations expressed as angular deviation from the reference mark, X_1 and X_2 were the noisy representation of stimulus direction 1 and 2 respectively (see below), w_1 and w_2 were the weights assigned to the corresponding evidence, and $N(0,\xi)$ was zero-mean Gaussian estimation noise with variance ξ . Because estimations y were expressed relative to the reference, so were the internal representations X_i . We used this format of internal representation and estimation reports because (i) in our design, the cursor movements used for estimation report were always initiated at the reference and (ii) recent work on post-decision biases has highlighted the importance of the reference (Zamboni et al. 2016). The reference-dependent format of internal representation assumed for X_i thus did not describe the sensory representation of motion direction, but rather a statistic extracted from that sensory representation in a task-dependent fashion (e.g. the direction with the largest posterior probability relative to reference (Jazayeri and Movshon 2006).

Here and below, X_1 and X_2 were computed by replacing the angular deviation of the physical stimulus from the reference, ϕ , with:

$$X_{i} = \phi_{i} + N(\delta, \sigma) \tag{2.4}$$

where $i \in (1,2)$, and δ and σ were each observer's individual overall bias and sensory noise parameters taken from the psychometric function fit to the binary choice data (see Equation 2.2 above). X_1 and X_2 thus approximated the noisy internal representation that governed observers' estimations. Specifying these two parameters in this fashion avoided adding additional free parameters to Equation 2.3. The approach was based on the assumption that a substantial portion of biases was shared between the choice and estimation judgments. We validated this assumption by confirming that biases were strongly correlated between binary choices (quantified as the horizontal shift of psychometric function estimations) and estimation reports (mean angular estimation error), across the 10 observers (Spearman's rho = 0.79; p = 0.0098). In sum, the Baseline model had three free parameters, w_1 , w_2 and ξ .

Correlated Noise model

This model assumed shared noise in the internal representations X_1 and X_2 . Specifically, the model assumed that additive noise corrupting the transformation of physical stimulus directions ϕ_i into X_i was correlated, to some degree, across the two intervals, inducing correlations between X_1 and X_2 , as follows:

$$X_1 = \phi_1 + \delta + \varepsilon$$

$$X_2 = \phi_2 + \delta + (1 - c)N(0, \sigma) + c\varepsilon$$
(2.5)

where ε was the noise in X_1 in a given trial, drawn from the distribution $N(0,\sigma)$, and $c \in [0,1]$. Parameter c governed the degree of correlation among the two internal representations. Thus, the noise in X_2 was made up of a portion correlated with the noise in X_1 (i.e., $c\varepsilon$) and another portion independent of X_1 (i.e., $(1-c)N(0,\sigma)$). The estimations were modeled by Equation 2.3 above.

Shift model

In this model, the choice induced a shift of the estimations into the direction of the chosen category, thus inducing an additive estimation bias consistent with the binary choice. Specifically, the estimations in this model were given by:

$$y = w_1 X_1 + w_2 X_2 + kD + N(0, \xi)$$
(2.6)

where κ was the additive shift parameter, and D was the vector of intermediate binary choices, taking the values [1, -1].

Selective Gain models

The Selective Gain model enabled testing for a selective change in sensitivity to Consistent vs. Inconsistent evidence conveyed by the direction of the second stimulus ϕ_2 . Consistency of that direction could be defined with respect to the initial choice or with respect to the first stimulus direction ϕ_1 . This led to two alternative versions of the Selective Gain model, specified next.

Choice-based Selective Gain. This model was as the Baseline model, except that the weights were allowed to vary depending on whether ϕ_2 was consistent or inconsistent with the initial choice:

$$y = w_{1c}X_1 + w_{2c}X_2 + N(0,\xi_c) \quad \text{if} \quad \operatorname{sign}(\phi_2) = D$$

$$y = w_{1i}X_1 + w_{2i}X_2 + N(0,\xi_i) \quad \text{if} \quad \operatorname{sign}(\phi_2) \neq D$$
(2.7)

where w_{1c} (w_{2c}) and w_{1i} (w_{2i}) were the weights for Consistent and Inconsistent trials, respectively, ϕ_1 and ϕ_2 were the physical stimulus directions from both intervals, and D was the vector of intermediate binary choice (values: 1 or -1 for CCW and CW reports, respectively). Since consistency could not be defined in trials where ϕ_2 was 0° , we excluded this subset of trials before fitting Equation 2.7.

Extended Choice-based Selective Gain. We also tested whether there was an additive shift, over and above the multiplicative gain modulation described by the Choice-based Selective Gain model. To this end, we extended the model from Equation 2.7 by means of the shift parameter from Equation 2.6, as follows:

$$y = w_{1c}X_1 + w_{2c}X_2 + kD + N(0,\xi_c) \quad \text{if} \quad \operatorname{sign}(\phi_2) = D$$
$$y = w_{1i}X_1 + w_{2i}X_2 + kD + N(0,\xi_i) \quad \text{if} \quad \operatorname{sign}(\phi_2) \neq D$$
(2.8)

This Extended Choice-based Selective Gain model was fit by constraining all parameters to take the values estimated by the basic version of the model (i.e. Equation 2.7), with the shift as the only free parameter. Parameter recovery indicated that leaving all parameter free to vary in the fit made the model too complex given the limited amount of data (see Parameter Recovery below).

Stimulus-based Selective Gain. This version of the model was as the previous one, except that consistency depended on the direction of the first stimulus

(specifically: the sign of its difference from the reference), not the initial choice:

$$y = w_{1c}X_1 + w_{2c}X_2 + N(0,\xi_c) \quad \text{if} \quad \operatorname{sign}(\phi_2) = \operatorname{sign}(\phi_1)$$
$$y = w_{1i}X_1 + w_{2i}X_2 + N(0,\xi_i) \quad \text{if} \quad \operatorname{sign}(\phi_2) \neq \operatorname{sign}(\phi_1)$$
(2.9)

Since consistency could not be defined in trials where ϕ_1 and ϕ_2 was 0°, we excluded this subset of trials before fitting Equation 2.9).

Choice-based Selective Gain model after matching evidence disparity. As control for the differences in disparity between motion directions in first and second interval, we randomly subsampled trials such that the absolute distance between ϕ_1 and ϕ_2 in Consistent and Inconsistent trials is matched. This was done before fitting the Choice-based selective gain model described above.

Choice-based Selective Gain model for numerical task. The data from the numerical task (numerical integration, Figure 3A) were also fit with the Choice-based Selective Gain model, but with the following differences due the nature of the task design and the smaller number of trials per individual than available for the perceptual task. As in the perceptual task, the mean evidence also exhibited a small bias after splitting by choice-consistency. Different from the perceptual task, the group average estimations exhibited a small opposite trend, i.e., there was an interaction between mean evidence vs. estimations and (Consistent, Inconsistent) condition. The model as in Equation 2.7 could not capture this interaction, and indeed we found that fitting the model without accounting for it yielded poor fits. (Please note that the model-free analysis of sensitivity described below was unaffected by this issue.) To account for this, we introduced two additional free parameters θ_c and θ_i , as follows:

$$y = w_{1c}X_1 + w_{2c}X_2 + N(\theta_c, \xi_c) \quad \text{if} \quad \text{sign}(\phi_2) = D$$

$$y = w_{1i}X_1 + w_{2i}X_2 + N(\theta_i, \xi_i) \quad \text{if} \quad \text{sign}(\phi_2) \neq D$$
(2.10)

where θ_c and θ_i accounted for the above interaction, ϕ_i was the mean of 8 samples (again relative to the reference, i.e., 50) in each interval (i = 1, 2),

 X_i was the noise-corrupted and biased internal representation of the mean value in each interval computed as in Equation 2.4. For the results shown in Figure 3, we constrained all weights to be positive based on the assumption that weights should not be negative. We also constrained the possible values of θ_c and θ_i to be within a reasonable range, [-10, 10], still far larger than the magnitude of the interaction observed in the group average estimation data. The above constraints were introduced in order to obtain reliable model fits in the face of limited data (trials). Results were qualitatively similar (especially, significant difference between Consistent and Inconsistent weights for the second interval) when fitting the model without those constraints.

Further, as an additional control, we also fitted the model without constraints and without θ (i.e. Equation 2.7), after randomly sampling trials from Consistent and Inconsistent conditions, so as to minimize the above interaction through matching the mean evidence between both conditions. Because this procedure substantially reduced the number of trials (23%), we only fitted the model on the remaining data after pooling trials across all participants. We repeated this 'mean-matching' procedure 500 times and re-fitted the model for each random trial selection. The median across iterations of the difference in weights for Consistent vs. Inconsistent was 0.065, with a 95% confidence range that excluded zero (0.00001, 0.133). In sum, also this second approach for fitting the data from the numerical task supported the re-allocation of the weights for the second evidence dependent on choiceconsistency as observed in the first model-based approach (Figure 3B) as well as in the model-free analysis (Figure 3C).

Likelihood computation

We used maximum likelihood estimates to estimate parameters and the goodness of fit of different models. For any unique combination of experimental variables (first and second stimuli, and choice), we numerically derived the estimation distribution of each model for a given parameter set and used this estimation distribution to assess the likelihood of the estimation reported by participants on a given trial with the corresponding experimental variables. All models described above assume that the stimuli on each trial are represented in the form of scalar values. Thus, the estimation distribution represents a distribution of estimations over trials. Specifically, the estimation distribution was the expected distribution of estimations for a given set of experimental variables, if the model was simulated several times using the same set of parameters. Using this numerical method avoided the need to rely on large number of stochastic simulations in order to compute the likelihoods and made the fitting procedure less prone to converging to local minima. For all models except for the Correlated Noise model (see below), we numerically derived each model's estimation distribution for each experimental condition, by first generating Gaussian distributions centered at w_1X_1 and w_2X_2 with standard deviation $|w_1|\sigma$ and $|w_2|\sigma$ for intervals 1 and 2 respectively. Then, we set the probability in the non-chosen side to zero in the interval 1 distribution (i.e. we truncated the distribution to only have density in the chosen side) and normalized it so as the integral of the resulting distribution is equal to 1. We combined the probability distributions corresponding to stimulus 1 (truncated distribution) and 2 using convolution and renormalized the resulting distribution. Note that different weights applied to stimulus 1 and stimulus 2 distributions (see Equation 2.3, Equation 2.6, Equation 2.7, Equation 2.8, Equation 2.9, Equation 2.10) in different models. We then generated a zero-mean Gaussian probability distribution with variance ε and convolved this distribution with the distribution from the previous step, renormalizing the resulting distribution to obtain the estimation probability distribution for that trial. We used this probability distribution to calculate the likelihood of the reported estimation in the trial. Finally, we summed the logarithm of likelihood values over all trials to obtain the final log-likelihood value for a given set of parameters.

For the Correlated Noise Model, we used Monte Carlo techniques to simulate a probability distribution of estimations over trials. For each combination of experimental variables (combination of first and second stimuli, and choice), we generated a set of 10,000 normally distributed noisy representations (N(0, σ)) or noisy samples for interval 1 (X₁, Equation 2.5, top row). From these noisy samples, we discarded those where the sign of X₁ did not match the binary choice of the subject. These maintained samples featured as variable ε in interval 2 (Equation 2.5, bottom row). Another set of noisy samples was generated afresh for interval 2 (N(0, σ) in Equation 2.5). Note that the number of new noisy samples and thus of the simulated representations X₂ was less than 10,000 because of the sub-selection described above. We combined X₁ and X₂ using Equation 2.3, to obtain a distribution of estimations for this trial. Smoothing kernels were obtained from the simulated estimations in each trial in order to identify the underlying distributions, which were then used to calculate the likelihood. The kernels were defined using non-parametric Epanechnikov function. Finally, we summed the logarithm of likelihood values of all trials to obtain the final log-likelihood value for a given set of parameters (w_1 , w_2 , ξ and c). In some trials, the likelihood of the estimations was zero regardless of the values of the parameters (possibly because these estimations were motor lapses), resulting in an optimization function that never converged. To address this, we added one simulated estimation trial in the response range (-180° to 180° in steps of 1) to the distribution of estimations obtained from Equation 2.3, before obtaining the estimation kernels. This did not influence the maximum likelihood fitting procedure in other trials, but just gave an estimation kernel with non-zero probability value for the whole range of estimations.

Comparison of fitting procedure to Bronfman et al. 2015

In our previous report on analyses of the data from the numerical task (Bronfman et al. 2015), we had also fitted a so-called Selective Gain model to the data from the numerical estimation task and compared that to a model without such selective gain modulation. Here, we applied a model fitting procedure that differed from the previous one in two important respects. First, in the previous study, models were fitted to the across participants aggregated data (i.e., a 'fixed effects' approach) (Bronfman et al. 2015), while we here fitted models to each participant's data individually. The second difference concerns the calculation of the likelihood in (Bronfman et al. 2015): In the previous study, model estimations were derived for each trial (using 1000 simulated trials) and then the likelihood of the reported estimation was computed under the simplifying assumption that the model's estimation distribution is Gaussian. This was done in order to avoid kernel-based smoothing of the simulated estimations, which could significantly slow down the fitting procedures. In the current study, however, we did not make this assumption since in all models the predicted distribution, albeit symmetric looking under most parameter sets, always had non-zero skew. We thus computed likelihoods from the actual non-Gaussian distributions that the models predict.

Fitting procedure and computation of confidence intervals

To obtain the best fitting parameters that maximize the likelihood function of each model, we used Subplex algorithm (Bogacz and Cohen 2004), a generalization of the Nelder-Mead simplex method, which is well suited to optimize high dimensional noisy objective functions. Subplex starts at a specified starting point of the objective function and works by dividing the parameter space into subspaces. It then performs a simplex search in each of these subspaces before converging on the set of parameters that maximize the function. The starting points were randomly chosen from the interval [0,20] for ε and [0 1] for w_1 and w_2 .

We used bootstrapping (Efron and Tibshirani 1986) to obtain confidence intervals for the fitted parameters for each individual. Specifically, we randomly selected trials with replacement and fit the selective gain model to these resampled datasets. We repeated this procedure 500 times, each time using Subplex optimization with starting points at the best-fitting parameters of the actual data. We then obtained confidence intervals from the distribution of estimated parameters.

Parameter recovery

We simulated data with different sets of parameters using the number of trials as in a typical dataset. We then fit the simulated data using the Choice-based Selective gain model (Equation 2.7). Overall, the model recovered parameters well: The Spearman correlations between actual and recovered parameters was 0.8 for Noise parameters ($p < 10^{-10}$), and it ranged between 0.91 to 0.94 for all weight parameters ($p < 10^{-10}$). Importantly, the model also did not introduce any spurious correlations between the recovered parameters. Inter-parameter correlations for the actual parameters ranged between -0.03 to 0.03 (p > 0.39), and between -0.05 to 0.03 (p > 0.23) for the recovered parameters. This allowed us to confirm that our fitting procedures were able to recover the parameters when the ground truth of the data was known.

Simulations of the most complex model assessed here, the Extended Choice-based Selective Gain (Equation 2.8) with all parameters left free to vary, also showed decent overall recovery of parameters (Spearman correlations ranged between 0.78 to 0.94, p < 10^{-10}). However, for special cases, the fits exhibited significant spurious correlations between parameter estimates. Specifically, we introduced a few iterations where the actual shift

parameter (k_{actual}) was 0, or the consistency parameter for weights of interval 2 ($\Delta w_{2,actual} = w_{2c,actual} - w_{2i,actual}$) was 0. When k_{actual} was 0, the mean recovered shift ($k_{recovered}$) in these iterations was -1.63 (p = 0.0043). This spurious shift was introduced by the model at the expense of Δw_2 i.e., the correlation between $k_{recovered}$ and $\Delta w_{2,actual} - \Delta w_{2,recovered}$ was -0.85 (p < 10⁻⁵). Likewise, in iterations where $\Delta w_{2,actual} = 0$, mean $\Delta w_{2,recovered} = 0.115$ (p = 0.046). This spurious consistency effect was introduced by the model at the expense of the shift parameter i.e., the correlation between $\Delta w_{2,recovered}$ and $k_{actual} - k_{recovered}$ was -0.65 (p < 10⁻³). Because these spurious correlations rendered fits of this complex model generally hard to interpret, we did not report any parameter estimates from this model.

In order to test whether there was evidence for a shift, over and above selective gain modulation, we constrained all parameters in Extended Selective Gain to take the fit values for basic Choice-based Selective Gain, allowing only the shift parameter free to vary. The thus estimated biasing effects (i.e., shift and weight difference for consistent and inconsistent second stimulus) did not exhibit any correlation across participants (r = 0.38, p = 0.279), ruling out spurious dependencies.

Model comparison

We used Bayesian Information Criterion (BIC) to quantitatively compare the ability of different models to explain the data. BIC is given by:

$$BIC = -2.\ln(L) + m.\ln(n)$$
(2.11)

where L was the likelihood value, m was the number of free parameters in the model and n was the number of observations that are used to fit the model (Schwarz 1978). BIC values were compared across models and the model with lowest BIC value was identified as the model that best explains the data among all candidate models. Specifically, a difference of 10 in BIC values suggests very strong evidence in favor of the model with the lower BIC value (Kass and Raftery 1995). Since BIC values depended on the number of observations used to fit the model, we fit all models under comparison on the same subset of trials to enable us identify the model that best explains the data. We calculated BIC values for all individual model fits to identify the model that better explained the data for that participant.

2.3.6 Model-free analysis of estimation reports

We also assessed the impact of the second stimuli to participants' estimations in a model-independent fashion. The rationale was to quantify the impact of small differences in the evidence values (stimulus directions or numerical means) on the estimation reports produced by the participant, depending on whether the two directions were consistent or inconsistent with the previous choice. Our analysis aimed to compare the separability of distributions of single-trial estimations from subsets of trials, between Consistent and Inconsistent conditions. We quantified the separability of estimation distributions by means of the receiver operating characteristic (ROC) from signal-detection theory (Green and Swets 1966). The area under the ROC-curve, referred to as ROC-index, could range from 0 to 1. An index of 0.5 implied perfectly overlapping distributions (i.e. no sensitivity to the 10° difference) and any deviation from 0.5 implies some sensitivity to the evidence. An ROC-index of 1 (or 0) implied that the two distributions were completely separable.

We intended to use the ROC measure for specifically quantifying the sensitivity to the smallest presented difference (10°) in the direction of the second stimulus (ϕ_2), while eliminating the impact of the direction of the first stimulus (ϕ_1) on the final estimation. To this end, we used the following procedure for the perceptual task. All trials (except for those with $\phi_2 = 0^\circ$ where choiceconsistency was not defined) were first sorted by whether the direction of second interval was consistent or inconsistent with the initial choice. For each thus-defined condition (Consistent, Inconsistent), we further sorted trials by ϕ_1 (i.e., -20°, -10°, 0°, 10°, 20°). For each ϕ_1 we compared estimation distributions from trials with ϕ_2 differing by 10° (i.e. -20° vs. -10° and 10° vs. 20°). The resulting ROC-indices were then pooled across the different ϕ_1 directions, so as to yield a single pooled ROC index, separately for Consistent and Inconsistent conditions. We pooled the ROC indices by means of weighted averaging, whereby the weight of each ROC index was determined by the number of trials that went into the calculation of that ROC-index. That number differed substantially between ROC indices due to the uneven distribution of pairs of directions of the first and second stimulus (Figure S1D). The resulting ROC-indices were compared between both conditions by means of permutation tests (see next section). We obtained qualitatively identical results when simply discarding the trial pairs with small trial numbers (< 15) and averaging the other ROC-indices without weighting (mean difference in ROC-index between Consistent and Inconsistent trials across subjects = 0.04).

ROC indices for the numerical task (numerical integration, Figure 3C) were computed as for the perceptual task, with the following exceptions. We binarized the mean of the stimulus presented in the first interval into two bins (ϕ_1 > 50 and ϕ_1 < 50), and split the mean of the second stimulus into four bins with means at 40, 47, 53 and 60, two each on either side of the reference number 50. Those four bins were treated equivalently to the different ϕ_2 values in the description for the perceptual task above.

2.3.7 Simulated estimations from the models

We used two different methods to assess if the individual models could explain the effect captured by the model-free analysis described above. Specifically, we simulated estimations both using the best fitting parameters from each model, and by sampling from a wider range of parameters. We then calculated the ROC indices on the simulated estimations in consistent trials and inconsistent trials.

To simulate estimations in each trial, we first calculated the internal representations x_1 and x_2 using Equation 2.4 (Equation 2.5 for the Correlated Noise model). In addition, we ensured that the sign of x_1 matches the binary decision made by the subject in the trial. We then combined these internal representations to obtain a simulated estimation for the trial, using the corresponding parameters and equations for each model.

Simulated estimations from the best fitting parameters

We used the best fitting parameters for each individual from each model, simulated the estimations, and calculated the ROC-indices using the procedure described above. This process was repeated 500 times for each subject and each model, to obtain the confidence intervals. We then compared the median ROC-indices between consistent trials and inconsistent trials across subjects.

Simulated estimations using a range of parameters

For each model, we simulated estimations across a range of parameters in order to identify the dependence of the consistency effect in the model-free analysis on the relevant parameters. We simulated a single fixed-effects subject, whose trial distribution was obtained by combining the trial distributions of all subjects. For each combination of parameters in each model, we calculated the estimations using the corresponding equations described above. We then performed the model-free analysis to obtain the ROC-indices for consistent trials and inconsistent trials. We defined "Consistency" as the difference between these ROC-indices. A positive value of Consistency suggests that this set of parameters replicates the model-free findings observed in the behavioral data. This procedure was repeated 100 times and the mean of the difference in ROC-index between Consistent and Inconsistent conditions was shown in color code for each parameter combination in Figure S2B. We showed the heatmaps as a function of the two parameters that may give rise to a difference non-zero Consistency, by marginalizing this value across all other parameters.

Simulations of Extended Conditioned Perception model

For simulations, we extended the Conditioned Perception model described in (Luu and Stocker 2018; Stocker and Simoncelli 2007) for discrimination and estimation judgments on a single stimulus to our task with two successive stimuli, intermittent choice, and a total estimation judgment at the end. We simulated a version of this 'Extended Conditioned Perception" model, in which the posterior distribution over stimulus directions after both stimuli were conditioned on the intermittent choice. The resulting procedure was as described in the section Likelihood computation for Choice-based Selective Gain model above, with the following differences. In the Conditioned Perception model, contrary to all other models presented, external stimuli on each trial are not represented by scalar variables but as posterior probability distributions over stimulus features (i.e., directions), given the sensory stimulus and the choice. Before conditioning, the mean of this posterior distribution for each stimulus was given by each stimulus' direction and its standard deviation was given by the individual psychometric noise (parameter σ from Equation 2.4, so-called 'input noise' capturing imprecise encoding of direction). Per trial two such distributions were obtained, one per stimulus. Both of these distributions

were then conditioned on the choice: the probability was set to zero on all stimulus directions that were inconsistent with the choice, for the first and the second stimulus. The two resulting distributions were finally combined with equal weight, producing an overall estimation distribution for each trial. We extracted the mean from the resulting distribution as the model's estimate of direction. In different runs of the simulations we added different amounts of independent Gaussian (zero-mean) noise to these estimates ('output noise' capturing both imperfect memory of stimulus identity as well as motor noise). The resulting value was taken as the estimation report on a given trial. We simulated this model in order to assess if it would produce similar behavioral features as Choice-based Selective Gain. Please note that a more elaborate version of the Conditioned Perception model has been used to fit estimation data in (Luu and Stocker 2018).

For each version of the Extended Conditioned Perception model, we systematically varied the input and output noise parameters, applied ROC analysis to the resulting estimation reports for Consistent and Inconsistent trials, and plotted the difference between ROC indices for both conditions as a function of the parameter combination (Figure S2B, right panel).

In further simulations of this model, we used an average subject (pooling the trial distributions across all subjects), with input noise as mean parameter σ across all participants, and varied only the output noise. We computed the mean estimations as a function of average stimulus direction predicted by the model for several levels of output noise (Figure 1C). We performed an additional analysis to uncover subtler differences in the behavior of Extended Conditioned Perception and Choice-based Selective Gain (Figure S2G, H). The Extended Conditioned Perception model was essentially insensitive to new evidence inconsistent with the choice. Thus, we reasoned that the fraction of inconsistent estimations (i.e., estimations falling on the side of the reference opposite from the choice) predicted by this model should be lower than the fraction predicted by the Choice-based Selective Gain model. Specifically, the increase in this fraction as a function of inconsistent second stimulus should be higher for the Choice-based Selective Gain model. To test this prediction, we simulated estimations for the fixed-effects subject using the Conditioned Perception model for different levels of output noise, and the Choice-based Selective Gain model with the mean of the best fitting parameters across

subjects. We then calculated the fraction of inconsistent estimations for correct and error trials, separately for positive and negative direction of the first stimulus ($x_1 = -10^\circ$, 10°), as well as for ambiguous first direction ($x_1 = 0^\circ$), in the simulated estimations. We repeated this procedure for 100 iterations, and compared the mean fraction across the iterations in both models to that of the behavioral data (Figure S2G, H).

2.3.8 Quantification and statistical analysis

Non-parametric permutation tests (Efron and Tibshirani 1986) were used to test for group-level significance of individual measures, unless otherwise specified. This was done by randomly switching the labels of individual observations either between two paired sets of values, or between one set of values and zero. After repeating this procedure 100,000 times, we computed the difference between the two group means on each permutation and obtained the p value as the fraction of permutations that exceeded the observed difference between the means. All p values reported were computed using two-sided tests.

2.4 DATA AND SOFTWARE AVAILABILITY

Data and analysis scripts are available on https://github.com/BharathTalluri/ postchoicebias.

2.5 ACKNOWLEDGMENTS

We thank Ana Vojvodic for help with data collection, and Zohar Bronfman and Noam Brezis for continued discussion and helpful feedback on the manuscript. This research was supported by grants from the German Research Foundation (DO 1240/2-1 and DO 1240/3-1, to THD), the German Academic Exchange Service (to AEU), a Marie Curie Individual Fellowship (to KT). We acknowledge computing resources provided by NWO Physical Sciences.



2.6 SUPPLEMENTARY FIGURES

Figure S1: **A.** As in Figure 1B, Figure 1C top panel, but for individual subjects. Subjects 1 to 7 from left to right in top row, subjects 8 to 14 from left to right in bottom row. β values show the slope of best-fitting line for bottom panel in each subject. Green boxes show subjects used for the analysis, and red boxes show subjects discarded from the analysis since their estimations did not monotonically increase with mean stimulus direction (see Section 2.3). **B.** β values from panel A for individual subjects, showing the 4 subjects with low β values discarded from the analyses. C. Distribution of analyzed motion directions. Each cell in the matrix represents the average number of trials presented to the participants for the corresponding values of $\phi_1 \& \phi_2$ directions where the participant made a valid response (see Section 2.3). Trial numbers are printed and color coded. **D.** Same as **C**, but for Consistent trials (left panel), and Inconsistent trials (right panel). Related to Figure 1.



Figure S2: A. Estimation noise parameters of the Choice-based Selective Gain model for Consistent and Inconsistent trials. Format same as Figure 2B. B. The difference in ROC-indices between Consistent and Inconsistent conditions in simulated estimations for various models, as a function of different model parameter combinations. Estimations were simulated using a wide range of parameters for each model (see Section 2.3). Only the Choice-based Selective Gain and Conditioned Perception models showed a positive difference depending on parameters. For Choice-based Selective Gain, this difference was positive when $w_{2c} > w_{2i}$ and negative difference when $w_{2c} < w_{2i}$. **C**, **D**. Same as figure 2B, but when the model was fit only on error trials (i.e., stimulus category opposite to choice) and trials where $\phi_1 = 0^\circ$ (i.e., no evidence). We observed a trend towards the selective gain effect (i.e., Consistent vs. Inconsistent weight difference) in both cases. E. The Choice-based Selective Gain model was fitted after first equating (through random sub-selection, see Section 2.3) Consistent and Inconsistent trials for the difference between motion directions from first and second stimuli. Any form of sensory adaptation should depend on this similarity and thus disappear in this control analysis. Instead, we found the same choice consistency dependent modulation in sensitivity (Figure S2E). The resulting weights for Consistent and Inconsistent conditions are shown in the same format as in Figure 2B. Also in this control analysis, weights were larger for second evidence that was consistent compared to inconsistent with the initial choice, ruling out any effect of evidence disparity across the two intervals. F. Residual shift parameter from the extended Choice-based Selective gain model. G. Fraction of inconsistent estimations as function of second stimulus direction, for the data, and as predicted by the Selective Gain and Extended Conditioned Perception models. A multiplicative dependence of estimations on interval 2, as predicted by the Choice-based Selective Gain model, suggests an increase in the fraction of inconsistent estimations with increasing inconsistent stimulus in interval 2. Thus, in correct trials, this fraction will increase as we progress along the X-axis from left to right when $\phi_1 = -10^\circ$, and vice-versa when $\phi_1 = 10^\circ$. This trend reverses in error trials. The increase in the fraction of inconsistent estimations as a function of inconsistent stimulus in interval 2 in the data suggests a multiplicative gain reduction for choice-inconsistent stimulus. The extended Conditioned Perception model predicts the fraction to be largely independent of second stimulus. **H.** Same as **G**, but for trials in which ϕ_1 = 10°. The Selective Gain model predicts that in inconsistent trials (where binary choice and sign of ϕ_2 are opposite), the fraction of inconsistent estimations increases as the value of inconsistent stimulus in interval 2 increases (V-shaped curve), and in consistent trials (where binary choice and sign of ϕ_2 are the same) this fraction decreases (inverted v-shape). Again, fractions predicted by the Selective Gain model agree with the data, while that of extended Conditioned Perception model are largely independent of second stimulus. Related to Figure 2.



Figure S3: **A-C.** Parameters of the Choice-based Selective Gain model in the numerical task. **A.** Format same as Figure 3B, but for weights for first interval. **B.** Estimation noise. **C.** θ parameter. **D.** as Figure 2C, but for numerical task. Related to Figure 3.

CHOICES CHANGE THE TEMPORAL WEIGHTING OF DECISION EVIDENCE

DECISIONS DO NOT OCCUR IN ISOLATION, BUT ARE EMBEDDED IN SEQUENCES OF OTHER DECISIONS, OFTEN PERTAINING TO THE SAME SOURCE OF EVIDENCE. HERE, WE CHARACTERIZED THE IMPACT OF INTERMITTENT CHOICES ON THE ACCUMULATION OF A PROTRACTED STREAM OF DECISION-RELEVANT EVIDENCE TOWARDS A FINAL DECISION. HUMAN PARTICIPANTS PERFORMED TWO VERSIONS OF A DECISION TASK BASED ON EITHER PERCEPTUAL EVIDENCE (NET DIRECTION OF RANDOM DOT MOTION PATTERNS) OR SYMBOLIC EVIDENCE (RAPID SEQUENCES OF NUMBERS). BOTH TASK VERSIONS REQUIRED TWO SUCCESSIVE JUDGMENTS OF THE EVIDENCE AT DIFFERENT TIMES DURING THE EVIDENCE STREAM: AN INTERMITTENT CATEGORICAL CHOICE HALF-WAY THROUGH THE EVIDENCE STREAM, AND A CONTINUOUS ESTIMATION OF THE MEAN EVIDENCE AT THE END OF THE EVIDENCE STREAM. IN A CONTROL CONDITION, SUBJECTS EXECUTED A SIMPLE MOTOR RESPONSE (NO CHOICE) INSTEAD OF REPORT-ING BINARY CHOICE, USING THE SAME MOTOR RESPONSE AS USED FOR CHOICE, TO CONTINUE THE TRIAL. THE INTERMITTENT RESPONSE PROMPT WAS FOLLOWED BY LARGER DILATIONS OF PARTICIPANTS' PUPILS ON TRIALS REQUIRING A CHOICE THAN ON TRIALS ONLY REQUIRING A SIMPLE MOTOR RESPONSE, INDICATING A LARGER TRANSIENT ELEVATION OF CENTRAL AROUSAL LEVELS IN THE CHOICE CONDITION. THE

Preprint posted to bioRxiv.

doi: 10.1101/2020.03.06.979690. This chapter was reprinted under a CC-BY 4.0 license. AUTHOR CONTRIBUTIONS: Conceptualization, T.H.D. and M.U.; Methodology, B.C.T., A.E.U., Z.Z.B., N.B. and K.T.; Software, B.C.T., A.E.U., and K.T.; Formal analysis, B.C.T., and A.E.U.; Investigation; B.C.T. and A.E.U.; Writing- original draft, B.C.T. and T.H.D.; Writing- review & editing, B.C.T., A.E.U., Z.Z.B., N.B., K.T., M.U., T.H.D.; Visualization, B.C.T.; Supervision, T.H.D., M.U. and K.T.; Project administration, T.H.D. and M.U.; Funding acquisition, T.H.D.

INTERMITTENT CHOICE ALSO REDUCED THE SENSITIVITY TO SUBSEQUENT EVIDENCE, AND FLIPPED THE RELATIVE TEMPORAL WEIGHTING OF EARLY AND LATE EVIDENCE IN THE FINAL ESTIMATION JUDGMENT. THE INDIVIDUAL EXTENT OF THE CHOICE-INDUCED OVERALL (NON-SELECTIVE) SENSITIVITY REDUCTION PREDICTED THE EXTENT OF THE SELECTIVE DOWN-WEIGHTING OF SUBSEQUENT EVIDENCE INCONSISTENT WITH THE INITIAL CHOICE, A FORM OF CONFIRMATION BIAS. ALL CHOICE-INDUCED CHANGES IN EVIDENCE SENSITIVITY WERE PRESENT IN BOTH, PERCEPTUAL AND NUMERICAL TASKS. WE CONCLUDE THAT CATEGORICAL CHOICES ABOUT A PROTRACTED EVIDENCE STREAM DYNAMICALLY ALTER THE PROCESSING OF THE EVIDENCE, CONSISTENT WITH AN ACTIVE INTERNAL STATE CHANGE TRIGGERED BY THE CHOICE.

3.1 INTRODUCTION

Many decisions are made under uncertainty, on the basis of noisy, incomplete, or ambiguous decision-relevant 'evidence'. An extensive body of research on perceptual decisions under uncertainty has converged on the idea that evidence about the state of the sensory environment is continuously accumulated across time (Bogacz et al. 2006; Gold and Shadlen 2007). In the perceptual choice tasks commonly used in the laboratory (but see (Glaze et al. 2015; Murphy et al. 2020; Ossmy et al. 2013)), performance is maximized by weighing evidence equally across time (Bogacz et al. 2006). Yet, the evidence weighting applied by human and non-human decision-makers often deviates substantially from such flat weighting profiles (but see (Brunton et al. 2013; Wyart et al. 2012)): some studies found stronger weighting of early evidence ('primacy'; (Kiani et al. 2008; Nienborg and Cumming 2009; Odoemene et al. 2018; Zylberberg et al. 2014; Drugowitsch et al. 2016; Tsetsos et al. 2012)), and yet others even non-monotonic weighting profiles (Bronfman et al. 2016).

Most of these studies of perceptual choice have focused on the within-trial factors governing decision-making, ignoring interactions between consecutive decisions or stimuli. However, real-life decisions are not isolated events, but embedded in a sequence of judgments based on continuous streams of information. Indeed, a growing body of evidence has shown that perceptual choices are biased by the choices made on previous trials (Abrahamyan et al. 2016; Akaishi et al. 2014; Braun et al. 2018; Busse et al. 2011; Fernberger 1920; Fischer and Whitney 2014; Fritsche et al. 2017; Fründ et al. 2014; Kim

et al. 2017; Liberman et al. 2014; Pape and Siegel 2016; Raviv et al. 2012; St John-Saaltink et al. 2016; Tsunada et al. 2019; Urai et al. 2017, 2019; Yu and Cohen 2009). Recently developed task protocols provide new tools for assessing the impact of choices on the accumulation of subsequent decision evidence. These tasks prompt two successive judgments within the same trial: commonly a binary choice followed up by a continuous estimation (Bronfman et al. 2015; Jazayeri and Movshon 2007; Luu and Stocker 2018; Stocker and Simoncelli 2007; Talluri et al. 2018; Zamboni et al. 2016) or a confidence (Fleming et al. 2018; Rollwage et al. 2018) judgment. Specifically, some tasks prompt binary choice and estimation judgments sequentially, separated by a second evidence stream presented in between (Bronfman et al. 2015; Fleming et al. 2018; Rollwage et al. 2018; Talluri et al. 2018). These task designs have led us to two insights. First, the overall sensitivity to evidence following the intermittent choice is reduced in a non-selective ('global') fashion, a finding obtained in the domain of numerical decisions (Bronfman et al. 2015). Second, sensitivity for information consistent with the binary choice is selectively enhanced, at the expense of less sensitivity for choice-inconsistent evidence, a choice-induced evidence re-weighting that produced a bias to confirm the initial choice and that was found for both perceptual and numerical decisions (Talluri et al. 2018). In the latter study, we did not examine the non-selective impact of choice in reducing sensitivity for subsequent evidence.

Here, we re-analyzed the datasets from both the previous studies (Bronfman et al. 2015; Talluri et al. 2018) to develop a more comprehensive understanding of these two choice-dependent effects (selective evidence reweighting and non-selective sensitivity reduction), as well as their relationship. We tested for the following three outstanding issues: (i) if the choice-induced sensitivity reduction observed in the domain of numerical decisions generalizes to the domain of perceptual decisions; (ii) how an intermittent overt choice affects the temporal weighting of evidence, and (iii) if and how the overt choice-induced, overall reduction of sensitivity relates to the choiceinduced confirmation bias towards consistent evidence.



Figure 1: A. Perceptual task. After a first random dot motion stimulus was shown for 750 ms, participants received an auditory cue about whether to report a binary choice about the net motion direction (Choice trials) or to continue the trial (No-Choice trials). The choice entailed discriminating the motion direction as CW or CCW with respect to the reference (white line shown at about 45° in this example). On half the Choice-trials, auditory feedback was then given and the trial terminated. In the other half, and in all No-Choice trials, a second motion stimulus was presented (with equal coherence as the first, but independent direction), and participants were asked to report a continuous estimate of the mean direction of both stimuli by dragging a line along the screen with the mouse¹. **B.** Numerical task. After the first sequence of eight numerical samples, participants were instructed to either press the space bar (a quarter of all trials; No-Choice trials), or to give their binary choice about the average of the eight samples (mean > or < 50; Choice trials). On two-thirds of Choice trials (constituting a half of all trials), auditory feedback was presented and the trial terminated. In the rest, a second sequence of eight numerical samples was presented and participants were instructed to give a continuous estimate of the mean across the two sequences. Adapted from (Talluri et al. 2018) under a CC-BY license.

3.2 RESULTS

Participants reported a continuous estimate of the mean of fluctuating sensory (perceptual task, Figure 1A) or symbolic (numerical task, Figure 1B) evidence across two successive intervals. This estimate needed to be based on accumulating some internal representation of the fluctuating evidence – motion direction or numerical value in the perceptual or numerical tasks, re-



Figure 2: **A,B.** Perceptual task. **C.** Numerical task. **A.** Top: Continuous estimations as a function of mean direction across both stimuli. Bottom: Distribution of mean directions across trials. Data points, group mean; error bars, SEM. Stimulus directions and estimations were always expressed as the angular distance from the reference, the position of which varied from trial to trial (0° equals reference). **B.** Time courses of average pupil diameter aligned to trial onset for Choice and No-Choice conditions in Perceptual task. Left, average time course across whole trial. Right, close-up of time course during second stimulus interval (following intermediate motor response). Dashed vertical lines, mean response times across participants; grey vertical lines, different events during the trial. **C.** Same as **A** but for Numerical task. Mean evidence across intervals 1 & 2 in C split into discrete bins for illustration. All panels: solid lines, mean across participants; shaded region, SEM; black horizontal bars, p<0.05, cluster-based permutation test Choice vs. No-Choice.

spectively – across the two stimulus intervals. On a subset of trials (so-called Choice trials), participants were also asked to report an intermediate categorical choice after the first stimulus: a fine direction discrimination judgment relative to a visually presented reference line (Perceptual Task) or comparison of the numerical mean with 50 (Numerical Task). On the remaining set of trials (No-Choice trials), participants were asked to simply press a button for continuing the trial, without reporting a categorical judgment of the first evidence. The cue informing participants whether to report the discrimination judgment or to press a choice-independent button press came after the first stimulus interval. This design enabled us to quantify the degree to which evidence in each interval contributed to the final estimation and whether this depended on the overt report of a categorical choice (Section 3.4).

Estimation responses in both tasks increased with mean directional evidence across the two intervals (Figure 2A, Figure 2C), and did not differ between Choice and No-Choice trials, with negligible and statistically nonsignificant differences in the regression slopes for evidence against estimations (perceptual task: 0.0256, p = 0.8449; numerical task: 0.0125, p = 0.9083). Participants' pupils constricted after trial onset in the same way for Choice and No-Choice trials, an expected response to the onset of the random dot stimulus (pupil light reflex, Figure 2B; pupil diameter was only measured during the perceptual task). This constriction was followed by a dilation about 1 s after the intermittent response (see red/blue dashed vertical lines), indicating a phasic activation of central arousal systems (Aston-Jones and Cohen 2005; de Gee et al. 2017; Joshi et al. 2016; Reimer et al. 2016). Critically, this dilation was bigger for Choice than No-Choice trials (Figure 2B), reflecting the internal decision process associated with Choice (de Gee et al. 2014, 2017). Indeed, the bigger pupil dilation during Choice was not due to longer response times in that condition (and the associated longer accumulation of central inputs in the peripheral pupil apparatus; (de Gee et al. 2014, 2017)): response times were, in fact, shorter in Choice than No-Choice trials (see blue and red vertical lines in Figure 2B; permutation test, p = 0.0112).

3.2.1 Global down-weighting of subsequent evidence following intermittent choice

We first replicated our finding, previously reported for the Numerical Task (Bronfman et al. 2015), of lower sensitivity to subsequent evidence in the Choice condition, for the Perceptual Task (Figure 3). A statistical model-based as well as a model-free approach (Section 3.4) both showed a choice-dependent sensitivity reduction for subsequent evidence (Figure 3). Model weights for the second stimulus were significantly smaller in Choice trials compared to No-Choice trials (most individual participants, and the mean, above identity line in Figure 3A, Figure 3C). Likewise, a model-free measure of sensitivity to subsequent evidence (area under the ROC curve) was smaller on Choice trials compared to No-choice trials (most individual participants, and the mean, above identity line in Figure 3B, Figure 3D). In sum, overt choices reduce the sensitivity to subsequent evidence not only for numerical, but also for perceptual decisions.



Figure 3: **A.** Model weights for sensitivity to second interval in Choice and No-Choice conditions in Perceptual task. Dashed line, identity of Choice and No-Choice; points above diagonal indicate larger weights to No-Choice. **B.** As **A**, but for ROC indices quantifying the sensitivity to second interval in a model-free way in Perceptual task. Dashed line, identity of Choice and No-Choice; points above diagonal indicate greater sensitivity to No-Choice. Data points, individual participants, with identical color scheme from (**A**, **B**). (**C**, **D**.) As (**A**, **B**), but for Numerical task. Perceptual task, n = 10 participants; Numerical task, n = 20 participants; p values, permutation tests across participants (100,000 permutations).

3.2.2 Intermittent choice alters temporal weighting of sensory evidence

Having generalized the choice-induced sensitivity reduction across both domains of decision-making, we assessed if and how the intermittent choice affected the relative weighting of early vs. late evidence in the decision process underlying the final estimation judgments. For both tasks, the weights in Choice trials were higher for the first interval, and lower for the second interval compared to No-Choice trials, with a significant interaction between trial type (Choice vs. No-Choice) and interval (Figure 4A). In other words, the evidence weighting across the two intervals flipped from recency to primacy between No-Choice and Choice conditions (Figure 4B). This choice-induced flip in temporal weighting was also evident in the individual data: The sums of weights from both intervals were highly similar for Choice and No-Choice



Figure 4: **A.** Mean model weights for both stimulus intervals in Choice and No-Choice conditions in Perceptual task (left, n = 10 participants) and Numerical task (right, n = 20 participants). Error bars, SEM; F-statistic, interaction between interval and condition (repeated measures 2-way ANOVA). **B.** Direction of temporal weighting quantified as difference in model weights between interval 2 and interval 1, separately for each task; * p < 0.05, ** p < 0.005, permutation tests across participants (100,000 permutations). **C.** Sum of weights across both intervals in Choice and No-Choice condition, across participants from both tasks. **D.** Difference between weights in Choice condition and No-Choice condition, in both intervals across participants from both tasks. Data points, individual participants; solid lines, best fitting lines; r, Pearson's correlation coefficients; * p < 0.05, ** p < 0.005, ** p < 0.005, ** p < 0.005.

trials in each subject (Figure 4C), but the difference in Choice and No-Choice weights was negatively correlated between intervals (Figure 4D). Please note that no such constraint was imposed in the statistical models used to estimate the weights for both intervals (Section 3.4).

These results are in line with a 'push-pull' mechanism, in which a limited resource was distributed across sensitivity to evidence in both intervals: Reporting an intermittent choice after the initial evidence boosted sensitivity to that early evidence, but at the cost of reducing sensitivity to subsequent evidence. This effect could also explain the similarity in overall estimation accuracy between Choice and No-Choice conditions (Figure 2A, Figure 2C).



Figure 5: **A**. Relationship between global (quantified as the difference in weights to second interval between No-Choice and Choice conditions) and selective sensitivity modulation giving rise to Confirmation bias (quantified as the difference in weights to second interval between Consistent and Inconsistent conditions, from (Talluri et al. 2018)), across participants from both tasks. **B**. Relationship between the weighted mean of Consistent and Inconsistent conditions (from (Talluri et al. 2018)), weighted by the number of trials in each condition) and weights for the Choice condition for the second interval across participants from both tasks. **C**, **D**. Relationship between No-Choice weights for the 2nd interval; and Consistent weights (**C**), and Inconsistent weights (**D**). Data points, individual participants; solid lines, best fitting lines; dashed lines, identity lines; p-values comparing the two conditions in **B**, **C** & **D** are obtained from 2-way ANOVA with the conditions and tasks as factors; * p < 0.05, ** p < 0.005.

3.2.3 Choice-dependent non-selective, and selective sensitivity modulations are coupled

Finally, we found that the individual degree of the choice-dependent, overall (non-selective) reduction in sensitivity to subsequent evidence was closely related to the selective confirmation bias effect defined as a larger sensitivity to subsequent evidence consistent than inconsistent with the initial choice (Figure 5). We quantified the overall ('global') gain modulation as the difference in weights of interval 2 between Choice trials and No-Choice trials, and the selective choice-driven gain modulation as the difference in weights of interval 2 between trials with choice-consistent and -inconsistent evidence.

Participants with a stronger global gain reduction also showed a stronger selective gain modulation (Figure 5A). The weights for interval 2 in the Choice condition in Figure 3A, Figure 3C are indeed the weighted average of the corresponding weights in choice-consistent and -inconsistent evidence (Figure 5B). Furthermore, we found that the weights for interval 2 in the No-Choice condition were comparable in magnitude, and correlated across participants with the weights for choice-consistent evidence (Figure 5C), but not for choice-inconsistent evidence (Figure 5D). Thus, the reduction in sensitivity following an overt choice observed in Figure 3 was primarily driven by the reduction in sensitivity to evidence inconsistent with the initial choice.

3.3 DISCUSSION

Recent work has begun to expose the impact of choices on the accumulation of subsequent decision evidence, revealing an overall reduction in sensitivity to subsequent evidence (Bronfman et al. 2015) combined with a selective suppression of the gain of evidence inconsistent with the initial choice (confirmation bias; Talluri et al. 2018). Here, we extend this nascent line of work, by showing that the report of an intermittent choice about a protracted stream of perceptual or symbolic (numerical) evidence alters the weights assigned to pre- and post-choice evidence from recency to primacy. Further, we show that the above three effects are tightly related, consistent with generation by the same mechanism.

The overall reduction in sensitivity due to the initial choice observed here (Figure 3) corroborates earlier analyses of the Numerical Task data (Bronfman et al. 2015). The correspondence between these and our current findings from the Perceptual Task indicate that choice-induced decreases in sensitivity generalizes across different formats of decision evidence (from symbolic to low-level perceptual). We used the same methods to analyze data from the perceptual task presented here (but different from those in (Bronfman et al. 2015), see Materials and Methods), and data from the numerical integration task using a similar task protocol (Bronfman et al. 2015), and found strong correspondence between the two.

One important implication of our findings is that the temporal weighting profiles in evidence accumulation are neither fixed traits of decision-makers nor fixed properties of certain tasks, but are flexibly altered on the fly in a given task, depending on the presence or absence of an intermittent choice. Previous studies investigating temporal biases in decision-making found conflicting results, ranging from recency to flat profiles, to primacy. Differences in the task protocols and idiosyncratic tendencies of decision-makers are important confounds that complicate the comparison between these studies. Our current results show that the temporal weighting profile, within a given individual and a given task, can be effectively flipped, simply by asking the participant for an intermittent choice half-way through the evidence stream.

Importantly, the here-discovered, strong relationship between the individual strength of the choice-induced, global sensitivity reduction and choiceinduced, selective confirmation bias (Figure 5) is not a given, because both effects were operationalized in terms of two orthogonal comparisons: the sensitivity reduction by comparing sensitivity between trials with an intermittent choice and trials without such a choice; the confirmation bias by comparing trials with subsequent evidence that was consistent or inconsistent with the choice, within the trials that contain an intermittent choice. Thus, presence of a global sensitivity reduction effect does not imply presence of the confirmation bias, and vice versa. Even so, their correlations were tight, in line with a common underlying mechanism. In particular, the weights of new evidence in the No-Choice condition were comparable to, and correlated with the corresponding weights for the choice-consistent but not for choiceinconsistent evidence in the Choice-condition. This observation suggests a distinct state of the decision-maker when faced with information inconsistent with previous decisions, possibly reflecting the suppression of post-decisional dissonance (Festinger 1957).

It is tempting to interpret our findings as a signature of decision-related cortical evidence accumulation dynamics (Rolls and Deco 2010; Wang 2008; Wimmer et al. 2015), combined with neuromodulatory input (Eckhoff et al. 2009; de Gee et al. 2017; Murphy et al. 2020). Once the decision circuits have settled in an attractor (choice commitment), this will reduce the decision-maker's sensitivity to all subsequent evidence (see Bronfman et al. 2015, Supplement) – an effect that may hold regardless of whether that evidence is

consistent or inconsistent with the choice. Due to selective feedback from the accumulator circuit to early sensory regions encoding the evidence, the attractor state in accumulator networks may additionally cause selective gain modulation of subsequent incoming evidence (Wimmer et al. 2015) that can translate into the consistency-dependent, selective confirmation bias effect we found earlier (Talluri et al. 2018). These effects may have been amplified by choice-induced, phasic neuromodulatory input to cortex. Our task entailed an interrogation protocol, in which a categorical choice was prompted by the experimenter, when cortical decision circuits might not yet have reached a stable decision commitment; the choice prompt might then trigger the release of certain neuromodulators that pushes the decision circuits into an attractor state (Bogacz et al. 2006; Bronfman et al. 2015). Such a neuromodulatory signal might be reflected in pupil dilation (Breton-Provencher and Sur 2019; de Gee et al. 2017; Joshi et al. 2016), which we found to be larger during the choice, compared to the no-choice condition. In sum, by altering the dynamical properties of decision circuits in the brain, choices can have versatile and coupled effects on evidence accumulation.

3.4 METHODS

3.4.1 Experiment details

Perceptual task

The University of Amsterdam ethics review board approved the study (reference number 2014-BC- 3517). All participants gave their informed consent. Participants were presented with two random dot motion stimuli in succession, and were asked to estimate the average motion direction across the two intervals in each trial (Figure 1A). A white line plotted on top of the circular aperture served as the reference, whose position changed between trials. An auditory cue after the first interval instructed the participants to either (i) report a binary choice about the direction of dots in the first interval (clockwise or counter-clockwise w.r.t the reference; two-third proportion of all trials), or (ii) make a choice-independent button press (one-third proportion of all trials). This intermittent response allowed us to investigate if participants showed different sensitivity to the second stimulus depending on whether they reported a binary choice (so called "Choice trials"), or made a choice-independent motor response (so called "No-Choice trials"). The delay between the first and second stimuli was fixed (2 seconds), regardless of the reaction time of the subject. Half of all choice trials ended with an auditory feedback about the correctness of the binary choice to motivate participants to take the binary choice component seriously. The coherence of the stimuli was fixed at a pre-determined level for each subject, while the direction of coherent dots in the two intervals was sampled independently from five possible values (-20° , -10° , 0° , 10° , 20° relative to the reference line). 23 possible combinations of directions were used in the experiment (excluding the two most obviously conflicting directions: $-20^{\circ}/20^{\circ}$ and $20^{\circ}/-20^{\circ}$).

In all the analyses that follow, we used trials where participants made an estimation judgment (Choice trials and No-Choice trials). We excluded trials in which participants did not comply with the instructions i.e., when they pressed the mouse wheel on Choice trials or a choice key on No-Choice trials, trials in which the binary choice response time was below 200 ms, and trials where estimations were outliers. Outliers were defined as those trials whose estimations fall beyond 1.5 times the interquartile range above the upper-quartile or below the lower-quartile of estimations. Together, these excluded trials corresponded to \approx 7% of the total trials across participants.

We analyzed data from the same set of participants as in our earlier report (Talluri et al. 2018). Please refer to this report for a detailed description of the task, participants, and stimuli used in the experiments.

Numerical task

We reanalyzed data from the numerical integration task in (Bronfman et al. 2015) using the same analyses methods as the perceptual task. The task has a similar structure as the perceptual task above, with the exception that participants saw 16 numerical samples displayed in succession and reported their mean as a continuous estimate (Figure 1B). Like the perceptual task, participants received a cue midway through the trial (i.e., after the first 8 samples) to either report a binary choice about the mean of the 8 samples (greater, or less than 50), or make a choice-independent motor response. In 50% of all trials, the trial terminated with feedback after the binary choice. On

another 25% of the trials, participants saw the second stream of 8 numerical samples and made the continuous estimation judgment at the end (Choice trials). In the rest 25% of trials, participants made a choice-independent motor response (No-Choice trials) instead of the binary choice, and a continuous estimation judgment at the end. We analyzed data from all the trials where participants made the estimation judgment (50% of all trials).

We analyzed data from 20 out of 21 participants participated in the study. The remaining subject (subject 21) failed to do the task (Spearman's correlation between estimation and mean evidence in No-Choice trials: rho = 0.18, p = 0.117; and in Choice trials: rho = 0.17, p = 0.156). Please see the earlier reports (Bronfman et al. 2015; Talluri et al. 2018) for more detailed description of the task, stimuli, and participants.

3.4.2 Pupillometry

Horizontal and vertical gaze position as well as pupil diameter were recorded at 1000 Hz using an EyeLink 1000 (SR Research). The eye tracker was calibrated before each block. Blinks detected by the EyeLink software were linearly interpolated from -150 ms to 150 ms around the detected velocity change. All further data analysis was done using FieldTrip (Oostenveld et al. 2011) and custom Matlab scripts. We estimated the effect of blinks and saccades on the pupil response through deconvolution, and removed these responses from the data using linear regression. The pupil signal was bandpass filtered from 0.01 to 10 Hz using a second-order Butterworth filter, z-scored per block of trials, and down-sampled to 20 Hz. For both experimental conditions (Choice, No-Choice), we then averaged the time courses across trials, time-locked to either the onset of the evidence sequence.

3.4.3 Modelling estimation reports

General approach

We used a statistical modelling approach to estimate the relative contributions of the sensory evidence (i.e., physical stimulus corrupted by sensory noise) conveyed by both successive dot motion stimuli to participants' trial-
by-trial estimation reports, as in our previous report (Talluri et al. 2018). The noisy sensory evidence was described by:

$$X_i = \phi_i + \delta + N(0, \sigma) \tag{3.1}$$

where $i \in (1,2)$ denotes the interval, ϕ_i is the physical stimulus direction, N(0, σ) was zero mean Gaussian noise with variance σ , δ and σ were each observer's individual overall bias and sensory noise parameters taken from the psychometric function fit to the binary choice data (see STAR methods in Talluri et al. 2018 (Section 2.3 in this thesis)).

Global Gain model

We modelled a global, choice related reduction in sensitivity to evidence following an overt choice, by allowing the weights to vary separately in Choice trials and No-Choice trials. The estimations were modelled by:

$$y = w_{1c}X_{1c} + w_{2c}X_{2c} + N(0, \xi_c)$$

$$y = w_{1nc}X_{1nc} + w_{2nc}X_{2nc} + N(0, \xi_{nc})$$
(3.2)

where y was the vector of estimations, w_{1c} (w_{1nc}) and w_{2c} (w_{2nc}) were the weights for the noisy evidence encoded in intervals 1 and 2 in Choice (No-Choice) trials respectively. $N(0,\xi)$ was zero-mean Gaussian estimation noise with variance ξ that captured additional noise in the estimations, over and above the sensory noise corrupting binary choice.

Fitting procedure

We used maximum likelihood estimates to estimate parameters and the goodness of fit for different models. To obtain the best fitting parameters that maximize the likelihood function of each model, we used the Subplex algorithm (Bogacz and Cohen 2004; Rowan 1990), a generalization of the Nelder-Mead simplex method, which is well suited to optimize high dimensional noisy objective functions. Please refer to our earlier report for a detailed description of the fitting procedure (Talluri et al. 2018; Section 2.3 in this thesis).

3.4.4 ROC analysis for differences in sensitivity to evidence in interval 2

We assessed the impact of an overt choice on sensory evidence in interval 2 from participants' estimations in a model-free fashion, using the so-called ROC analysis. This analysis was based on the receiver operating characteristic (Green and Swets 1966), similar to the one used in our earlier report (see "model-free analysis of estimation reports" in Talluri et al. 2018; Section 2.3 in this thesis). By computing ROC indices between sets of trials that differed in their input, we could assess the sensitivity of the observer in using that input to guide their estimation reports.

For the perceptual task, in each condition (Choice and No-Choice), we first sorted trials based on the directional evidence in interval 1 (ϕ_1). For each ϕ_1 , we ran the ROC analysis on all pairs of estimation distributions, separated by 10° of directional evidence in interval 2 (ϕ_2): -20° vs. -10°, -10° vs. 0°, 0° vs. 10°, and 10° vs. 20°. This gave us 4 ROC-indices per ϕ_1 , one index for every pair of distributions compared. We then computed a weighted average ROC-index for each ϕ_1 , weighting the individual ROC-indices by the number of trials that went into the ROC analysis. The resulting ROC indices, which are robust to changes in ϕ_1 , are averaged to obtain one single ROC index per subject for each condition.

ROC indices for the numerical task were computed similar to the above procedure with the following exceptions: mean evidence in interval 1, and interval 2 were binarized (mean > 50 or mean < 50) resulting in two bins for interval 1, and interval 2 respectively. These binarized values were treated equivalent to ϕ_1 and ϕ_2 in the perceptual task above.

3.4.5 *Statistical tests*

Non-parametric permutation tests (Efron and Tibshirani 1986) were used to test for group-level significance of individual measures for each task, unless otherwise specified. This was done by randomly switching the labels of individual observations either between two paired sets of values, or between one set of values and zero. After repeating this procedure 100,000 times, we computed the difference between the two group means on each permutation and obtained the p value as the fraction of permutations that exceeded the observed difference between the means. All p values reported were computed using two-sided tests, unless otherwise specified.

To obtain the correlation values for data pooled from both the tasks (Figure 4, Figure 5), we first obtained Pearson's correlation coefficient for dataset from each task (also reported in the figures). We then applied Fisher-transformation on the correlation values, calculated their weighted average to obtain the pooled Fisher-transformed correlation coefficient. This quantity is used to obtain the pooled Pearson's correlation coefficient (using inverse Fisher transformation), and its corresponding p-value.

3.5 DATA AVAILABILITY

Behavioural data for the Perceptual Task is available at https://doi.org/ 10.6084/m9.figshare.7048430 (Talluri et al. 2018), and for the Numerical Task is available at https://datadryad.org/resource/doi: 10.5061/dryad.40f6v (Bronfman et al. 2015).

3.6 ACKNOWLEDGMENTS

We thank Ana Vojvodic for help with data collection. This research was supported by the German Academic Exchange Service (DAAD) (to AEU), a European Research Council Starting Grant under the European Union's Horizon 2020 research and innovation program (Grant No. 802905) (to KT), and the following grants from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): DO 1240/2-1 and DO 1240/3-1 (to THD). We acknowledge computing resources provided by NWO Physical Sciences.

NEURAL BASIS OF SEQUENTIAL BIAS IN PERCEPTUAL DECISION-MAKING

HUMAN DECISION-MAKING IS INFLUENCED BY THE HISTORY OF PRECEDING CHOICES, A WELL-KNOWN TENDENCY CALLED SEQUENTIAL BIAS. DESPITE ITS PREVALENCE IN VARIOUS TYPES OF DECISIONS, THE NEURAL BASIS OF SEQUENTIAL BIAS IN HU-MAN PERCEPTUAL DECISION-MAKING IS UNCLEAR. HERE, USING A COMBINATION OF BEHAVIORAL MODELLING, FUNCTIONAL MAGNETIC RESONANCE IMAGING, AND MULTIVOXEL PATTERN ANALYSIS TECHNIQUES, WE CHARACTERIZED CHOICE REPRE-SENTATIONS ACROSS THE WHOLE CORTEX DURING TASK-BASED ACTIVATIONS, AND DURING ONGOING ACTIVITY IN THE INTER-TRIAL INTERVALS IN DISTINCT NETWORKS OF REGIONS. WE FOUND THAT CHOICE REPRESENTATIONS DURING TASK ACTIVATIONS IN POSTERIOR PARIETAL AND MOTOR REGIONS RELATED TO BEHAVIORALLY QUAN-TIFIED SEQUENTIAL BIAS, SUGGESTING THAT TRANSIENT ACTIVATION OF CHOICE REPRESENTATIONS CONTAIN NEURAL CORRELATES OF SEQUENTIAL BIAS.

4.1 INTRODUCTION

Perceptual choices are influenced by a range of idiosyncratic preferences and heuristics, giving rise to biased behavior (Gardner 2019; Gold and Stocker

AUTHOR CONTRIBUTIONS: Conceptualization, T.H.D. and J.D.H.; Methodology, B.C.T., B.C., dG.J.W. and W.N.; Software, B.C.T., dG.J.W. and W.N.; Formal analysis, B.C.T.; Investigation; B.C.T.; Writing, B.C.T.; Visualization, B.C.T.; Supervision, T.H.D., J.D.H.; Funding acquisition, T.H.D.

2017; Rahnev and Denison 2018). One such bias, called Sequential bias, is the individual tendency to systematically repeat or alternate previous choices. Sequential biases were reported across several perceptual paradigms in humans (Fernberger 1920; Fischer and Whitney 2014; Liberman et al. 2014), rodents (Busse et al. 2011; Hermoso-Mendizabal et al. 2020; Odoemene et al. 2018), and non-human primates (Gold et al. 2008; Tsunada et al. 2019). Recent studies have started to characterize sequential bias at behavioral (Abrahamyan et al. 2016; Fritsche et al. 2017; Fründ et al. 2014; Urai et al. 2017), and computational (Bonaiuto et al. 2016; Marcos et al. 2013; Meyniel et al. 2016; Urai et al. 2019; Yu and Cohen 2009) levels. Sequential biases are strongly influenced by decision confidence (Braun et al. 2018; Desender et al. 2018; Drugowitsch et al. 2019; Lak et al. 2020; Urai et al. 2017) and phasic arousal (de Gee et al. 2017, 2019; Urai et al. 2017), are adaptable to stimulus statistics (Abrahamyan et al. 2016; Braun et al. 2018), and arise due to previous perceptual choices, but not the motor responses used to report them (Akaishi et al. 2014; Braun et al. 2018).

Sequential bias could influence the processing of decision-relevant evidence in a range of brain regions, from early sensory areas to higher-level decision-making areas to motor execution areas. In early visual cortex, the representation of orientation in the current trial was found to be selectively biased towards the previous choices (St John-Saaltink et al. 2016; but see Lueckmann et al. 2018). Posterior parietal cortex in mice contained sequential bias signals (Marcos et al. 2013; Scott et al. 2017), and optogenetically inactivating this region reduced the bias in behavior suggesting a causal role (Akrami et al. 2018; Hwang et al. 2017). Furthermore, activation in frontal eve-field was associated with updated choice estimate (Akaishi et al. 2014), tracking sequential bias, while neuronal activity in sensorimotor cortex in 12-30 Hz (beta band power), quantified using magnetoencephalography, predicted motor response alternation across trials (Pape and Siegel 2016). Together, these findings suggest that sequential bias signals in the brain are not localized to any region, but are found across multiple levels of the decision-making network. Yet, each of these studies focused on individual cortical regions, and a comprehensive characterization of sequential bias in choice representations across the entire cortex is missing.

In this study, we addressed this question using a combination of fMRI and whole-brain multivoxel pattern analysis methods. Human subjects performed a two alternative random dot motion direction discrimination task where stimuli across subsequent trials were independent, repeated, or alternated more often. We identified distinct networks of cortical regions containing transient choice representations at a single trial level, and sustained choice representations in the inter-trial intervals. We then mapped behaviorally quantified sequential biases onto the neural activity patterns in these regions, isolating behaviorally relevant effects. We found that choice representations in the current trial, but not in the inter-trial interval predict sequential bias across subjects suggesting that transient activation of choice information is closely linked to behavioral biases.

4.2 RESULTS

4.2.1 Behavior

Subjects performed a fixed duration version of the random dot motion direction discrimination task (Figure 1A) under two difficulty levels (see Methods). Subjects were faster to respond (mean difference in RT between high coherence and low coherence = 0.046 s, p < 0.001) and had greater accuracy (mean accuracy in high coherence = 85%, p < 0.001; in low coherence = 68%, p < 0.001; difference between conditions, p < 0.001) in high coherence condition compared to low coherence condition. Choice-response mapping was fixed during individual runs in a session, but counterbalanced across runs. Unknown to the subjects, stimulus transition probability (i.e., the change in stimulus direction from one trial to the next) was manipulated in separate runs resulting in three distinct task conditions: Random, Repeating, and Alternating (see Section 4.4.2).

4.2.2 Sequential biases are idiosyncratic

Previous studies identified that sequential biases vary substantially across individuals (Abrahamyan et al. 2016; Braun et al. 2018; Fründ et al. 2014; Urai et al. 2019). These idiosyncratic biases can be quantified as the shift in



Figure 1: A. Left, Behavioral task: Subjects judged the net direction (up vs down) of a random dot motion stimulus with varying motion strength and direction (see Section 4.4.2). Different events were indicated by different colors of the fixation cross during the course of a trial, which started with a baseline period (red fixation cross), followed by the onset of decision-relevant evidence (green fixation cross, coherence levels titrated to 65% or 80% accuracy of individual subjects). Subjects were prompted to provide a response when the fixation turned red (max. response time = 1.25 s), followed by the inter-trial interval (blue fixation cross). Random dot motion was present throughout the duration of a trial, with 0% coherence at all times except the evidence window. Right, pre-defined transition probabilities between motion direction in successive trials resulted in three different conditions: Random (transition probability = 0.5), Repeating (transition probability = 0.2), Alternating (transition probability = 0.8). **B.** Psychometric functions conditioned on previous choice illustrate sequential biases in Random condition as a horizontal shift between the two plots, for three example subjects showing a bias to repeat previous choice (left panel), switch from previous choice (right panel), and no history bias (central panel). **C.** Choice weights for previous trial in all three conditions (n = 26). Example subjects in B were highlighted using different markers in the Random condition.See also main text, and Section 4.4.

subjects' psychometric functions (relating signed motion energy to choice probability) fit separately by conditioning on the previous choice in the Random condition (where the stimulus directions in successive trials were drawn independently). Figure 1B shows three example subjects with a tendency to repeat their previous choice (left panel), alternate from their previous choice (right panel) from their previous choice, or show no bias (central panel). A recent study proposed a more comprehensive approach to quantifying sequential biases by explicitly modelling the contribution of past choices as an additive bias to the conventional fit of the psychometric function (Fründ et al. 2014). This approach was used to quantify the influence of previous choice on the current trial, separately for the three conditions (Figure 1C; see Section 4.4.5). Indeed, subjects exhibited idiosyncratic sequential biases in this task: mean previous choice weight \pm sem in the Random condition = -0.2122 \pm 0.033, and the previous choice weights across the group were not significantly different from zero, p = 0.22.

A recent study using session-wise manipulations in stimulus transition probability such as the one described here found that subjects adjusted the sign and strength of their sequential biases to match the statistics of different environments. Such adjustments were not present in our data at the group level. Specifically, there was no difference in previous choice weights between Repeating and Random conditions, mean = 0.0923, p = 0.70; and between Alternating and Random conditions, mean = -0.0038, p = 0.98.

4.2.3 Neural correlates of single trial choices

Cortical regions carrying single trial choice information in the Random condition were identified using multi-voxel pattern classification analysis using a searchlight procedure (de Gee et al. 2017; Haynes and Rees 2006; Hebart et al. 2012, 2016; Kriegeskorte et al. 2006). This analysis was restricted to data from the Random condition since autocorrelations in successive stimulus directions in Repeating and Alternating conditions make it difficult to identify regions encoding only choice information. Searchlight decoding was done separately for the two choice-response mapping runs, to identify cortical regions sensitive to both perceptual choice and motor response information. The resulting maps were averaged (see Section 4.4.6). This approach revealed three cortical regions encoding choice at a single trial level (Figure 2A): right posterior parietal cortex (r-PPC), primary motor cortex (M1), and dorsal Premotor cortex (PMd) (see *Choice-specific cortical regions of interest* in Section 4.4.6 for information about the anatomical labels underlying these ROIs).

To test whether the regions identified above also carry choice information in other conditions, we trained a classifier on single trial choice data from the Random condition and used it to predict the corresponding choices in the Repeating and Alternating conditions (see Section 4.4.6). All three regions



[△]Choice/Response mapping preserved

Figure 2: **A**. Searchlight decoding of single trial choice in Random condition. Maps show the regions with significant choice decoding (p < 0.05) after cluster correction. Decoding was done separately for the two choice-response mappings (see Section 4.4.6), and the resulting maps of precision scores (sensitive to both perceptual choice and motor response) were averaged. **B**. Generalization of choice decodability for regions identified in **A** in Repeating and Alternating conditions. Classifier trained on Random condition was used to predict choices in the other two conditions. **C**. Generalization of choice representations in **A** to previous trials in Random condition. Classifier trained on the current trial in Random condition was used to predict previous choices up to 5 lags. *, p < 0.05; **, p < 0.01; ***, p < 0.001. Dashed horizontal line, chance level performance of the classifier. Significance levels were obtained by comparing precision scores at the group level (n = 26) against 0.5 using permutation tests (see Section 4.4.7). See also Figure S1.

show single trial choice decodability in the Repeating condition (r-PPC, mean precision scores = 0.51, p = 0.036; M1, mean precision scores = 0.56, p < 0.001; PMd, mean precision scores = 0.55, p < 0.001), while M1 and PMd show single trial choice decodability in the Alternating condition (M1, mean precision scores = 0.57, p < 0.001; PMd, mean precision scores = 0.56, p < 0.001), establishing the generality of the choice representation in these regions across all experimental conditions (Figure 2B).

To test whether the choice representations in the above ROIs also generalize to the past trials, a classifier trained on current choice in the Random condition was used to predict previous choices up to 5 trials in the past (see Section 4.4.6). All three regions show above-chance decodability for up to 2 previous trials (3 previous trials for M1; Figure 2C), suggesting that the choice representations in these regions are relatively stable over multiple trials. This generalization effect was not due to choice correlations across trials, as shown by the lack of correlations between precision scores for predicting previous choice and choice repetition probability across subjects (Figure S1) in all ROIs.

4.2.4 Neural correlates of sustained activity during choice repetitions

Sustained BOLD (Blood Oxygenation Level Dependent) activity during choice repetitions, including the inter-trial intervals between the trials, was modelled as a choice 'streak' (see Section 4.4.6; Figure S2). Two additional nuisance regressors modelling trial-based activation and motor response were used so that the streak regressor is not influenced by any trial-related BOLD responses. Univariate analysis of the streak regressor showed deactivations in fronto-parietal regions during the streaks (Figure 3A), a network of regions that have been proposed to be involved in perceptual decision-making, and cognitive control (Cole and Schneider 2007; Kable and Glimcher 2009; Keuken et al. 2014).

Multivoxel pattern classification (decoding) was used to test whether the ROIs carrying choice information at a single trial level (Figure 2A) also carry sustained choice signals across a streak of choice repetitions. The decoding analysis was performed by concatenating data from all runs, irrespective of the conditions (see see Section 4.4.6). The potential confounds in the single trial choice decoding that arose due to autocorrelations in acrosstrial stimulus directions in Repeating and Alternating conditions does not apply here since individual streaks (separated by choice alternations) are independent of each other (see Figure S2). Choices could be decoded from sustained patterns of activity during choice streaks from M1, and PMd regions (Figure 3B), suggesting that these carry sustained choice information across multiple trials, including the inter-trial intervals. The ROI-based decoding analysis was supplemented by a searchlight analysis to identify additional fine-grained cortical regions carrying sustained choice information in the patterns of activity during choice streaks. This searchlight analysis revealed additional ROIs (Figure 3C): left Cuneus, Supplemental Motor Area (SMA),



Figure 3: **A.** Univariate analysis of streak regressor. **B.** Decoding of streak regressor in the ROIs identified in Figure 2A. *, p < 0.05; **, p < 0.01; ***, p < 0.001. Dashed horizontal line, chance level performance of the classifier. **C.** Searchlight decoding of choices from activity during streaks. Maps in **B** and **C** show the regions with significant regression coefficients and precision scores respectively (p < 0.05) after cluster correction. Significance levels were obtained by comparing precision scores at the group level (n = 26) against 0.5, and regression coefficients against 0 using permutation tests (see Section 4.4.7). Decoding of the streak regressor was done separately for the two choice-response mappings (see Section 4.4.6), and the resulting precision scores were averaged. Yellow outlines in **C** indicate regions with significant single trial choice decoding. See also Figure S2.

Orbito-Frontal Cortex (OFC), and right Intraparietal Sulcus (r-IPS) (see *Choice-specific cortical regions of interest* in Section 4.4.6 for information about the anatomical labels underlying these ROIs).

4.2.5 Neural correlates of choice predict sequential bias

Having identified the neural correlates of choice information at a single trial level (Figure 2), and during the inter-trial interval (Figure 3), the next question is if any of these choice-specific signals relates to subjects' sequential bias, quantified as the weight exerted by previous choices on the current choice (Figure 1C). To address this, precision scores from decoding choice at a single trial level in the Random condition, separately for current and previous trials, were correlated to the previous choice weights in Random condition across subjects (see Figure 4A for an example scatter plot showing correlation for r-PPC) in ROIs identified above (Figure 2A, Figure 3C). Correlating previous



Figure 4: **A.** Example scatter plot showing the between-subjects correlation between previous choice weights and precision scores for previous choice decoding, in r-PPC. **B, C.** Correlation coefficients between previous choice weights, and precision scores for decoding previous choice, current choice, and choice streaks in ROIs from single trial decoding (**B**) and singe streak decoding (**C**). *, p < 0.05; **, p < 0.01; ***, p < 0.001; p-values indicate the statistical significance of Spearman's correlations. See also Figure S3.

choice weights with precision scores from decoding previous choice (light colored bars in Figure 4B, Figure 4C) showed significant correlations across subjects in r-PPC (Spearman's rho = 0.52, p = 0.006), M1 (Spearman's rho = 0.60, p = 0.001), PMd (Spearman's rho = 0.40, p = 0.041), and OFC (Spearman's rho = 0.41, p = 0.038); and with precision scores from decoding current choice (dark colored bars in Figure 4B, Figure 4C) showed significant correlations across subjects in r-PPC (Spearman's rho = 0.51, p = 0.008), M1 (Spearman's rho = 0.42, p = 0.033), and r-IPS (Spearman's rho = 0.71, p < 0.001). These correlations were absent for previous stimulus weights (Figure S3). To predict whether the sustained choice representations between choice repetitions also predicted the strength of sequential biases, the precision scores from decoding choice streaks across all conditions were correlated to the mean previous choice weights across all conditions. Interestingly, sustained choice representations were not predictive of sequential bias in the ROIs considered (striped bars in Figure 4B, Figure 4C). Thus, trial engagement activates the neural correlates of sequential bias, but not sustained choice representations in the ROIs considered above.

4.2.6 Neural correlates of switch vs stay behavior

A dominant idea in foraging and hierarchical reasoning tasks is that the anterior cingulate cortex (ACC) underlies switch behavior in sequential foraging decisions (Hayden et al. 2011; Kolling et al. 2012), sequence learning



Figure 5: **A**. Univariate analysis contrasting switch and stay regressors. Maps show the regions with significant activation during choice alternations compared to choice repetitions (p < 0.05) after cluster correction. **B**. Maps showing the conjunction of different brain regions identified in this study. See also Figure S4.

(Findling et al. 2019), covert rule switches (Sarafyazd and Jazayeri 2019), behavioral shifts to exploitation (Quilodran et al. 2008), and volatility tracking (Rushworth and Behrens 2008). In our task, given the idiosyncratic sequential biases in participants' behavior, it is possible they relied on a network including ACC to track recent choice history and break a streak of consecutive choices. To test whether this was the case, we modelled all the trials where the choice switched or repeated (stayed) from the previous trial using separate regressors (see Figure S4) and contrasted these two regressors (see Section 4.4.6), with the hypothesis that the contrast would be 0 across different brain regions. Surprisingly, this contrast was positive in regions in intraparietal, primary motor, pre-motor, left-anterior insular, and anterior cingulate regions (Figure 5A) suggesting that these regions were more active during choice switches compared to choice repetitions.

Our complementary analyses identified a diverse set of cortical regions for different behavioral variables of interest, summarized in Figure 5B. Interestingly, pre-motor and primary motor regions, responsible for motor response preparation and execution respectively, seem to contain neural correlates for several of these variables, probably due to their downstream nature in the perceptual decision-making cortical hierarchy, making them candidate regions to contain choice-sensitive neural representations beyond simple motor responses.

4.3 DISCUSSION

We used fMRI to characterize choice representations during single trial activity, and during ongoing activity in the inter-trial intervals across a streak of choice repetitions. We found that sequential bias correlated with the decodability of previous and current choices in posterior parietal and motor regions, but not that of sustained representations between choice repetitions. Our findings suggest that sequential bias behavior is influenced by choice-selective signals during active trial engagement, but not during the pre-trial interval.

Earlier studies identified neural correlates of sequential bias in diverse brain regions in humans. Sensory representations in early visual areas during the current trial were biased towards previous percepts (van Bergen and Jehee 2019; St John-Saaltink et al. 2016), possibly reflecting decision-related feedback signals about previous choices from higher level cortical regions (Nienborg et al. 2012; Wimmer et al. 2015), to selectively direct attention towards previously chosen alternatives in the incoming information (Fritsche and de Lange 2019; Talluri et al. 2018). During the pre-stimulus interval, beta band activity (in 12 Hz – 30 Hz frequency) over sensorimotor cortex predicted motor response alternations (Pape and Siegel 2016), while activation in frontal eye-field reflected updated choice-estimate (Akaishi et al. 2014). Here, we extended this line of research by identifying the neural correlates of sequential bias in activity patterns across the entire cortex, for both task-evoked activity and ongoing activity in the inter-trial intervals. Using behavioral modelling, we rigorously quantified sequential bias in behavior, and related these behavioral effects to the neural correlates of choice signals to separate behaviorally relevant and irrelevant neural representations.

We identified brain regions carrying choice representations at a single-trial level in posterior parietal, and (pre-) motor regions (Figure 2A, Figure 2B). Previous studies of perceptual decision-making in humans identified robust choice-selective signals in an action-independent format in intraparietal, and prefrontal regions by decoupling perceptual choices from motor responses on a trial-by-trial basis (Hebart et al. 2012, 2016; Heekeren et al. 2006). The location of choice representations suggests that in our task, where the choice-

response mapping was not decoupled on a trial-by-trial basis, choice-related activity could be reliably mapped onto motor preparation signals. Indeed, a large body of work in non-human primates employed fixed choice-response mapping and found choice-predictive activity in a variety of regions involved in selection and preparation of motor responses (Christopoulos et al. 2018; Gold and Shadlen 2007). Recent studies employing a similar fixed choiceresponse mapping in humans also found choice-selective signals in anterior intraparietal and motor cortical regions (Donner et al. 2009; de Gee et al. 2017; Wilming et al. 2020), suggesting that the locus and format of choicerelated activity depends on the behavioral context of the task employed. The choice representations in the regions identified above generalized to previous choices, up to 2 trials in the past (Figure 2C), and choice correlations across successive trials could not explain these generalization effects.

Recent behavioral studies found that sequential bias is primarily driven by previous choices and not motor responses (Akaishi et al. 2014; Braun et al. 2018), by decoupling choices and the motor commands used to report these choices on a trial-by-trial basis. A recent study trained rats in olfactory and auditory tasks in randomly interleaved trials, and found that the dependence of choices on past trials transferred across modalities where rats tended to repeat the same action across trials, irrespective of changes in sensory modalities in successive trials (Lak et al. 2020). Under such conditions, choices are treated as a problem of movement selection and sequential bias could indeed modulate neural signals underlying the motor responses used to indicate behavioral choices. We found neural activity predictive of previous choices also in regions primarily responsible for preparing and executing motor responses (Figure 5B), in line with such a prediction.

We found sustained choice signals in Cuneus, intraparietal, orbitofrontal, and motor regions (Figure 3B, Figure 3C) during the choice streaks, demonstrating a persistence of choice information in neural activity several seconds after a decision has been reached. Persistence of choice information could arise from slow decay of choice-related activity during the ITI in attractor networks underlying decision-making, producing choice bias in subsequent trials (Bonaiuto et al. 2016). However, the sustained choice representations do not seem to be relevant for bias behavior observed in our task (Figure 4B, Figure 4C). The parietal and motor regions showing significant single trial choice decoding could be responsible for translating memory of previous choices into behavioral bias, explaining the lack of behavioral relevance of the sustained choice information. Indeed, there was little overlap between regions showing choice decoding at a single trial level and those showing decoding across streaks, suggesting that the 'memory' of the choice across repetitions resides in a different network of regions (Figure 3C). Future investigations should look into the interaction between these two distinct networks of regions through connectivity measures in various temporal windows of interest.

We used GLMs in our study to identify task-relevant BOLD activation by convolving various events of interest during a session with a canonical HRF, which is the standard approach in such studies (Lindquist 2008). However, BOLD hemodynamic responses were found to vary across individuals, and across brain regions in a single individual (Aguirre et al. 1998; Handwerker et al. 2004). Thus, a potential confound in our study is mismodelling of BOLD responses by the canonical HRF. The slow event related design of our study (with an ITI of 14 s), and our inclusion of all the relevant nuisance regressors ensure that the effects observed were not due to 'carry-over' of activation in one trial to the next, or to the ITI. A potential direction of future investigation is to model subject specific, and region-specific BOLD hemodynamic responses using deconvolution approaches (Donner et al. 2008), to replace the canonical HRF, and test the robustness of the results obtained.

Subjects in our study showed idiosyncratic choice-history biases in the Random condition (Figure 1B, Figure 1C), in-line with earlier studies (Abrahamyan et al. 2016; Braun et al. 2018; Urai et al. 2019). Across the group, the idiosyncratic biases did not adapt to the statistics in the Repeating and Alternating conditions (Figure 1C). Such an absence of adaptability stands at odds with the findings of two recent studies where subjects were shown to adapt to changes in trial order statistics (Abrahamyan et al. 2016; Braun et al. 2018). This suggests that at least one of the contextual factors that differentiate our study from those above could drive adaptability: feedback after every trial (Abrahamyan et al. 2016, or prior knowledge about the presence of different environments, combined with a session-wise manipulation (as opposed to run-wise manipulation in our study) of the transition probabilities (Braun et al. 2018). That being said, the characterization of the adaptability of sequential biases to different environments at behavioral and neural level in individual subjects is a topic of future investigation with this dataset.

While sequential biases are detrimental in conventional laboratory tasks of decision-making, they could indeed be beneficial under different contexts where regularities in spatial and temporal statistics inform subjects' choices (Glaze et al. 2015, 2018; Gold and Stocker 2017; Murphy et al. 2020; Yu and Cohen 2009). To fully exploit such regularities, the brain needs to track the state of the world over shorter and longer timescales. Indeed, studies have shown that distinct cortical regions accumulate information over multiple timescales (Murray et al. 2014), enabling temporally extended information processing (Honey et al. 2012). Thus, identifying the neural correlates of this sub-optimal bias could provide insights into the general principles underlying decision making.

4.4 METHODS

4.4.1 Subjects

33 healthy human volunteers (age (mean \pm standard deviation) = 26.2 \pm 3.6 years, no. of males = 8) participated in the study. All subjects had normal or corrected-to-normal vision, and gave written informed consent. The study was approved by the ethics committee of the Department of Psychology, at the Humboldt University in Berlin (reference number 2014-17). Subjects received monetary payment for participation in the study, plus a finishing bonus for completing all sessions. Subjects were scanned in two batches (20 subjects in the first batch, and 13 subjects in the second batch). 7 subjects were excluded from further analyses due to excessive head movements during the fMRI session (head movement exclusion criteria: relative translation motion in an axis > twice the voxel size along the axis).

4.4.2 Task and Procedure

Main experiment

Subjects performed an up vs down random dot motion direction discrimination task (Figure 1A, left panel). Each trial consisted of four consecutive intervals, each indicated by a different color of the fixation cross: (i) a baseline period of 2 s (red fixation cross), (ii) the coherent motion of the random dot stimulus for a fixed duration of 750 ms (green fixation cross), (iii) response period until the subject reported her response, up to a maximum of 1.25 s (red fixation cross), and (iv) the inter-trial interval (ITI; blue fixation cross). The ITI was fixed at 14 s during the fMRI sessions, and at 4 s during the training sessions. The random dot motion stimulus was present on the display at all times with 0% motion coherence, except during the 750 ms period of coherent motion.

Each subject participated in four sessions on separate days: one behavioral training session (about 2 hr), and three fMRI sessions (about 2.5 hrs each). During the training session, subjects were first trained on the task for 40 trials using high coherence stimuli. This was followed by 500 trials with auditory feedback after every trial, split into 10 blocks using stimuli with different levels of coherence under method of constant stimuli (MOCS) to determine individual psychometric thresholds. The training session ended with 3 blocks of 40 trials each without feedback, with stimulus coherence levels titrated to performance at 65% and 80% accuracy from MOCS.

Each fMRI session started with a short training run of 80 trials, during which the coherence levels were varied using a QUEST staircase procedure (Watson and Pelli 1983). The starting points for the staircase procedure were set at the coherence levels determined during the training session. This training run served two purposes- to familiarise subjects with the task, and to titrate the coherence levels at 65% and 80% accuracy levels for presentation during the main experiment. The main experiment consisted of 6 runs of the direction discrimination task per session, with 40 trials per run. Choice-response mappings were counter-balanced across runs in each session i.e., subjects had to respond with their left (right) hand to indicate up (down) choice in three runs, and right (left) hand to indicate up (down) choice in the other three. The order of these runs was randomised in each session, and subjects received instructions on the display about the mapping at the beginning of the run. Subjects received cumulative feedback of their performance after each run.

Unknown to the subjects, we manipulated the transition probabilities between the two stimulus categories (up/down) over trials. Transition probability is defined as the probability of a change in the stimulus direction from the current trial to the next trial, regardless of coherence (Figure 1A, right panel). This resulted in three conditions per session: (i) Random (transition probability = 0.5), motion direction in each trial was chosen randomly independent of the direction in the preceding trial; (ii) Repeating (transition probability = 0.2), motion direction in the current trial switched from that of the previous trial with a 20% chance; and (iii) Alternating (transition probability = 0.8), motion direction in the current trial switched from that of the previous trial with a 80% chance. The three conditions, and choice-response mappings were counterbalanced across runs in each session such that subjects performed two runs- with one choice-response mapping each per condition.

In addition to the main experimental runs, each fMRI session consisted of two localizer runs- (i) motor response localizer run, and (ii) coherent motion direction localizer run. All fMRI runs started and ended with a baseline period of 14 s during which only the blue fixation cross was displayed.

Motor response localizer

The motor response localizer run had 30 trials, and in each trial, subjects were required to make a delayed motor response to a visual cue. The color of the fixation cross instructed subjects when to execute the corresponding motor response. Each trial started with a visual cue, which is either the word *LEFT* or *RIGHT*, presented above the central fixation cross for 300 ms. At the offset of the cue, the fixation cross turned red for 2 s during which subjects had to withhold their motor response. After the delay period, the fixation cross turned green following which subjects immediately made the corresponding motor response (left-hand response when the displayed word was *LEFT* and right-hand response when the displayed word was *RIGHT*). The fixation stayed green until the onset of the next trial after an inter-trial interval that varied between 2 s – 5 s.

Coherent motion direction localizer

The coherent motion direction localizer run had a controlled block design comprising 16 blocks of random dot motion stimuli. The blocks alternated between stimuli with 100% motion coherence and 0% motion coherence. The direction of coherent dots was sampled from up and down directions, counterbalanced across blocks, resulting in 4 blocks of upward coherent motion, 4 blocks of downward coherent motion, and 8 blocks of random motion. Subjects had to fixate on the central fixation cross at all times during the run. To motivate them to fixate, subjects performed an independent detection task- to make a motor response with their right hand whenever the fixation changed color briefly from blue to red (for 250 ms) at random points of time in the run.

4.4.3 Stimuli

Random dot kinematograms were generated using Psychtoolbox-3 (Kleiner et al. 2007), and presented on a 52.5 cm (37.5 cm in the training session) wide display with a resolution of 1024 x 768 pixels at the rate of 60 Hz. The display is placed at a distance of 142.5 cm (60 cm in the training session) from the subject. A fixation cross with each arm 5 pixels wide was presented at the center of the display. The random dot kinematograms consisted of a field of dots with a density of 6 dots/degree², presented within a circular aperture. The circular aperture was 5° in diameter with its center 3.5° below the center of the fixation cross. Each dot was 0.06° in diameter, had 100% contrast from the black background, and moved at 11.5°/sec. All dots were divided into 'signal dots' and 'noise dots', whose proportions defined the motion coherence levels. Signal dots were randomly selected on each frame, and moved in up or down directions. Signal dots that left the aperture wrapped around and reappeared on the other side. Signal dots had a limited 'lifetime' and were re-plotted in a random location after being on the display for ten consecutive frames. Noise dots were assigned a random location within the aperture on each frame, resulting in 'random position' noise with a 'different' rule (Scase et al. 1996). Three independent motion sequences were interleaved (Roitman and Shadlen 2002), making the effective speed of signal dots in the display 3.8°/sec.

4.4.4 Magnetic resonance imaging data acquisition

Each experimental session consisted of several MRI scans divided into the following groups: Echo Planar Imaging (EPI) scans during the main experimental task and localizer runs, a structural T1 scan using MPRAGE sequence for anatomical co-registration and cortical surface reconstruction, a high resolution T2-weighted structural scan to facilitate co-registration of EPI images and the high-resolution structural T1 scan, and field maps to correct for inhomogeneities in magnetic field.

Data were acquired with a 3 T Siemens TrioTim MRI scanner (with 3 T Siemens Prisma Fit in the second batch of subjects) equipped with a 12channel head coil. We measured blood oxygenation level-dependent (BOLD) changes in MRI signal intensity using an echo-planar pulse sequence with the following parameters: repetition time, TR = 2 s; echo time, TE = 30 ms; flip angle, 82°; in-plane resolution, 2.5 mm x 2.5 mm; slice thickness, 3 mm with a 25% distance factor; 31 slices (30 slices in the second batch of subjects), covering the whole cortex. Anatomical T1-weighted and T2-weighted scans were acquired with a voxel size 1 mm x 1 mm x 1 mm. Field maps were acquired with the following parameters: voxel size, 2.5 mm x 2.5 mm x 3 mm; repetition time, TR = 1400 ms (400 ms in the second batch of subjects); echo times, TE1 = 5.19ms, TE2 = 7.65 ms; phase encoding direction = *j*-.

4.4.5 Data Analysis: Behavior

Motion energy filtering

Even though we only used two coherence levels of random dot motion stimuli during the experiment, the stochastic nature of the stimulus lead to fluctuations in the sensory evidence within and across trials from the specified nominal coherence levels. To quantify choice-behavior as a function of the actual, rather than the nominal, single-trial evidence, we used motion energy filtering to estimate those fluctuations (Adelson and Bergen 1985; Urai et al. 2017), using the procedure described in (Urai and Wimmer 2016). This resulted in 45 discrete samples of motion energy per trial (750 ms long stimulus presented at a refresh rate of 60 Hz). The first 13 samples of the motion energy filter output corresponded to the 'rise time' of the filter (Kiani et al. 2008), and do not accurately reflect the actual motion energy values. Hence, the motion energy values of the remaining 32 samples were averaged to obtain a single motion energy value per trial.

Modelling sequential bias

Subjects' choice-behavior was modelled using logistic regression. Up and down categories in choices and stimuli were dummy coded as +1 and -1 respectively. The probability of making an up choice ($r_t=1$) on trial t with a given signed stimulus intensity $s_t=1$ (signed motion energy values, positive for up and negative for down stimulus categories) is described as:

$$P(r_t|s_t) = \gamma + (1 - \delta - \lambda)g(\delta + \alpha s_t)$$
(4.1)

where γ and λ were lapse rates for choices $r_t=1$ and $r_t=-1$ respectively, $g(x)=\frac{1}{1+e^{-x}}$ was the logistic function, δ was the bias term describing the overall bias for one choice, and α quantifies the perceptual sensitivity.

The model described in Equation 4.1 assumes that the choice on trial t is independent of trial history. To incorporate choice-history bias, we modified the bias term δ by adding a linear combination of different components of trial history (Braun et al. 2018; Fründ et al. 2014; Urai et al. 2017). The modified logistic regression model is described as follows:

$$P(\mathbf{r}_{t}|\mathbf{s}_{t}, \boldsymbol{h}_{t}) = \gamma + (1 - \delta - \lambda)g(\boldsymbol{\delta}\boldsymbol{h}_{t} + \alpha s_{t})$$
(4.2)

$$\delta h_{t} = \delta' + \sum_{k} w_{correct,k} h_{correct,k,t} + w_{error,k} h_{error,k,t}$$
 (4.3)

$$w_{choice,k} = w_{correct,k} + w_{error,k}$$

$$w_{stimulus,k} = w_{correct,k} - w_{error,k}$$
(4.4)

where δ' is the overall bias towards one choice, k is the condition (i.e., Random, Repeating, Alternating) $h_{correct,k,t}$ is the regressor for previous correct choice in condition k and takes the value +1 if the previous choice in condition k was correct and 0 otherwise, $h_{error,k,t}$ is the regressor for previous incorrect choice and takes the value +1 if the previous choice in condition k was incorrect and 0 otherwise, $w_{correct,k}$ and $w_{error,k}$ are the weights for $h_{correct,k,t}$ and $h_{error,k,t}$ respectively, $w_{choice,k}$ and $w_{stimulus,k}$ are the weights for previous choice and previous stimulus respectively, in condition k. The regressors $h_{correct,k,t}$ and $h_{error,k,t}$ are orthogonal to each other. Positive values of w indicate a bias to repeat the corresponding category, and negative values of w indicate a bias to switch from the corresponding category. The parameters of the regression model were fit by maximizing the log-likelihood L= $L=\sum_t \log P(r_t=1|\tilde{s_t},h_t)$ using an expectation maximization algorithm (Braun et al. 2018; Fründ et al. 2014).

4.4.6 Data Analysis: MRI

Preprocessing and analysis of the MRI data was done using custom-made software written in Python. A number of analysis steps relied on FSL (Jenkinson et al. 2012; Smith 2004), FreeSurfer (Dale 1999; Fischl et al. 1999), NiPype (Gorgolewski et al. 2011), Scikit-learn (Pedregosa et al. 2011), and Nilearn (Abraham et al. 2014).

Preprocessing

The T1-weighted anatomical scans for each subject were automatically segmented and inflated for visualization using FreeSurfer (Dale 1999; Fischl et al. 1999). The EPI scans in every run were then subjected to the following steps in sequence: (i) reorientation to match the approximate orientation of the standard MNI152 template image, (ii) brain extraction to remove non-brain tissue using BET tool in FSL, (iii) correcting for magnetic field inhomogeneities using B0 field maps and FUGUE (FMRIB's Utility for Geometrically Unwarping EPIs), (iv) image realignment to correct for small head movements (Jenkinson et al. 2002) using the EPI volume in the middle of the session as template, (v) linear detrending to correct for linear drifts and (vi) conversion to units of modulation (percent signal change) around the mean fMRI series. The preprocessed EPI volumes were then concatenated across runs to obtain a single BOLD time series per voxel per session.

Transformation matrices from EPI volumes to MNI space were computed in three steps: (i) the template EPI volume in every session (the volume in the middle of the session) was transformed to the T1-weighted space of the session using FLIRT (FMRIB's Linear Image Registration Tool) with 12 degrees of freedom and sinc interpolation (Jenkinson and Smith 2001), (ii) the T1-weighted image of the session was transformed to MNI space using FLIRT (affine transformation with 12 degrees of freedom and sinc interpolation), and (iii) the transformation matrices from (i) and (ii) were concatenated to obtain EPI volume to MNI space transformation. Finally, the inverse transforms from the above steps were computed to facilitate inverse transformation.

Univariate analyses

General approach. A generalized linear model (GLM) with the appropriate regressors of interest was used to estimate BOLD response amplitude of each voxel to various events during a session. Before estimating the regression coefficients, the BOLD timeseries and the individual regressors were z-scored, and the individual regressors were normalized to have a unit vector norm. In addition to the regressors of interest, one nuisance regressor per run was added to the GLM, consisting of a boxcar function whose onset and duration were the onset of the run and the duration of the run respectively.

Quantification of evoked BOLD responses during the trial. A GLM with two regressors per trial was used to estimate the trial-specific BOLD response amplitude of each voxel: the first regressor consisted of a stick function at the stimulus onset of the current trial, convolved with a canonical double-gamma hemodynamic response function (HRF), and the second nuisance regressor consisted of a stick function at stimulus onset of the remaining trials, convolved with the canonical HRF. This approach of fitting one GLM per trial was shown to produce trial-by-trial estimates that are more representative of the true activation magnitudes (Mumford et al. 2012).

Quantification of choice-specific activity during inter-trial intervals of choice streaks.The sustained component of the BOLD response amplitude during

choice repetition 'streaks' (see Figure S2) was modelled using a GLM. A streak consisted of two or more consecutive trials across which a choice was repeated. Each streak was modelled by one regressor, which was a boxcar function with onset at the stimulus onset of the first trial of the streak and offset at the stimulus onset of the last trial of the streak, convolved with the canonical HRF. In addition to the streak regressors, two nuisance regressors were used to model BOLD response amplitudes to trial-specific activation, and motor response: the former regressor consisted of a stick function at stimulus onset of all trials in a session, and the latter regressor consisted of a stick function at stimulus onset of all trials in a session and taking the values +1 and -1 for trials with Left- and Right-hand responses respectively, both convolved with a canonical HRF. To compute the whole brain map for group level analyses, the regression coefficients were first averaged across the streaks in a session. The session-wise maps thus obtained were transformed into MNI space, spatially smoothed (FWHM: 6 mm), and averaged across sessions to obtain a single map per subject.

Modelling Switch vs Stay choice behavior. Switch Vs Stay choice behavior was modelled using a GLM with two regressors: the first regressor modelled the switch trials in a session, and consisted of a stick function at stimulus onset of all trials where choice switched from the previous trial, convolved with the canonical HRF; the second regressor modelled the stay trials in a session, and is similar to the switch regressor, except that the stick function was at stimulus onset of all trials where choice was same as that of the previous trial (see Figure S4). In addition to the switch and stay regressors, one nuisance regressor were used to model BOLD response amplitudes motor response (see section above for a description of this regressor). Switch Vs Stay contrast for each session was computed by subtracting the corresponding regression coefficients of the Switch and Stay regressors. The contrast map was transformed into MNI space, spatially smoothed (FWHM: 6 mm), and averaged across sessions to obtain a single map per subject for group-level analyses.

Multivoxel pattern classification analyses

General approach. Multivoxel pattern classification was used to identify cortical regions encoding choice information. Binary classification was done using support vector machines (SVMs) with the following parameters: kernel='linear', regularization parameter (C)=1, tolerance=10⁻⁴, max. iterations=10⁶, class weight='balanced'. Before fitting the classifier, training data was balanced across both classes by randomly under-sampling the class with higher number of samples. Standardization was performed on each feature (i.e., voxel) of the training data using scaler transformation, which was then applied to the corresponding feature in the test data. The classifier was fit on the training data, and the fitted classifier was used to predict the classes of the test data. Classifier performance was quantified as the average of precision scores computed for the two classes of the test data separately. This fitting procedure was repeated 50 times to account for the randomness in under-sampling, and the reported metrics are the mean of the 50 iterations.

Classification was performed either using a searchlight-based approach or a region of interest (ROI)-based approach.

Searchlight-based approach. Searchlight-based multivoxel pattern classification (Kriegeskorte et al. 2006) was performed to identify fine-grained patterns encoding choice information across all the cortical regions (Hebart et al. 2012, 2016). A sphere of voxels (i.e., searchlight) was selected around a given voxel with a radius of 15 mm (\approx 603 voxels). Single trial regression coefficients obtained from the univariate analyses were extracted from these voxels, and these formed the pattern vectors for the classification of choices. Each pattern vector was assigned a label corresponding to the choice of the subject on that trial ('up' vs 'down'). The pattern vectors were iteratively split into training and test datasets using a 5-fold cross-validation scheme, by preserving the percentage of samples for each class in each fold. The mean cross-validated precision score across all folds for this searchlight was assigned to the center voxel of the searchlight sphere. This procedure was repeated iteratively for all voxels in the brain, yielding a map of precision scores per subject per session. This single trial choice decoding analysis was performed separately on the two runs of Random condition, to obtain one precision score map for each choice-response mapping run and the resulting maps were averaged. This combined individual map was, transformed into MNI space, spatially smoothed (FWHM: 8 mm) and averaged across sessions to obtain a single map per subject for group-level analyses.

Searchlight-based decoding of single streak choice-information was done using a similar approach as above with the following differences: a smaller searchlight with a radius of 10 mm (\approx 179 voxels) was used; the decoding analysis was performed by concatenating data across all runs (irrespective of the condition) in the session separately for the two choice-response mappings; the cross-validation scheme used was leave-one-run-out i.e., the pattern vectors were iteratively split such that the vectors from all runs but one constituted the training set, and those from the left-out run constituted the test set; and the spatial smoothing of individual maps was done with a FWHM=6 mm.

ROI-based approach. Cortical regions of interest were identified from the searchlight decoding procedure described above (see below). Three different decoding analyses were performed on these ROIs: (i) decoding current trial choice from current trial regression coefficients in Random condition, (ii) decoding previous trial choice from current trial regression coefficients in Random condition, and (iii) decoding current streak choice from current streak regression coefficients in all conditions. The decoding procedure for single trial (streak) choice is similar to that described in Multivoxel pattern classification analysis I with the only difference being that pattern vectors for the classification constituted all voxels from the ROI.

Generalization of single trial choice-decoding signals. Two different generalization analyses were performed on the cortical regions identified from searchlight decoding of single trial choice: (i) generalizing current trial choice decoding signals from Random condition to Repeating and Alternating conditions, and (ii) generalizing choice decoding signals from current trial to previous 5 trials in the Random condition.

In the former, the classifier was trained on current choice from current trial regression coefficients in the Random condition, and used to predict current choice from current trial regression coefficients in Repeating and Alternating conditions. This approach is similar to leave-one-run-out cross-validation, where the training data are from Random condition and test data are from either Repeating or Alternating condition.

To test the generalizability of current trial choice decoding signals to previous trials, the procedure was similar to that described above in the *ROI-based approach*, with the only difference being that the test data was used to predict the choices from each of the previous 5 trials instead of current choice.

Choice-specific cortical regions of interest

Choice-specific ROIs were defined as the regions showing significant abovechance decoding of single trial and single streak choice from the searchlightbased decoding (see *Searchlight-based approach* above) after performing group-level cluster correction (see Section 4.4.7). The ROIs, identified in the MNI standard space were transformed to session space, and all the ROIbased analyses were performed in session space. The anatomical labels for these ROIs were derived from a recent multimodal parcellation of cortical regions (Glasser et al. 2016.

Single trial choice decoding yielded the following ROIs: Posterior parietal cortex (r-PPC) comprising area 7PC in the right hemisphere; primary motor region (M1) comprising areas 1, 3b in both hemispheres and area 5L in the right hemisphere; dorsal premotor region (PMd) comprising areas 4, 6d in both hemispheres and area 6mp in the left hemisphere.

Single streak choice decoding yielded the following ROIs: cuneus (left-Cuneus) comprising areas DVT, V6 in the left hemisphere; intraparietal sulcus (r-IPS) comprising intraparietal areas LIPv, LIPd, IP1, MIP, AIP in the right hemisphere; supplementary motor area (SMA) comprising area 4 in both hemispheres and areas 5m, 24dd in the right hemisphere; orbitofrontal cortex (OFC) comprising areas OFC, pOFC, 13l in both hemispheres, areas 11l, 47s in the right hemisphere and areas 10pp, 10v in the left hemisphere.

4.4.7 *Statistical tests*

All the analyses reported were done separately in each session for all subjects. The individual metrics per subject for group level analyses were obtained by averaging corresponding metrics across all sessions. Nonparametric permutation tests were used to test for group-level significance of behavioral measures (Figure 1C), and classification precision scores in the ROI-based analyses (Figure 2B, Figure 2C, Figure 3B). This was done by randomly switching the labels of individual observations against zero (for behavioral measures), or 0.5 (for precision scores). This procedure was repeated 10,000 times and computing the difference between the two group means in each iteration, and a p value was obtained as the fraction of permutations that exceeded the observed difference between the means. p-values testing the significance of behavioral measures were computed using two-sided tests, and that of precision scores were computed using one-sided tests (observed precision score > 0.5).

Nonparametric permutation tests within the FSL Randomise implementation was used to test cluster-corrected single streak regression coefficients against 0 (Figure 3A), switch vs stay contrast against 0 (Figure 5A), and searchlight classification precision scores against 0.5 (Figure 2A, Figure 3C). FSL's randomize implemented 10,000 randomly generated permutation tests of the data to perform a Monte-Carlo style permutation test. This procedure was robust w.r.t inflated false-positive rates (Eklund et al. 2016). A cluster correction threshold of p < 0.05 was used in all cases.

Correlations between behavioral measures and classification precision scores (Figure 4, Figure S1, Figure S3) were quantified using Spearman's correlation.

4.5 ACKNOWLEDGMENTS

We thank Shuyan Liu for help with data collection, and Stefan Hetzer and Christian Labadie for their help in setting up the scanning protocols. This research was supported by grants from the German Research Foundation (DO 1240/2-1, DO 1402/2-2, DO 1240/3-1 and DO 1240/4-1, to THD). BCT's visit to Berlin for data collection was supported by a travel grant from the Boehringer-Ingelheim Fonds. We acknowledge computing resources provided by NWO Physical Sciences.

4.6 SUPPLEMENTARY FIGURES



Choice repetition probability vs generalisation precision scores (lag 1)

Choice repetition probability in Random condition

Figure S1: Scatter plots showing the lack of correlation between the precision scores for predicting previous choice using a classifier trained on the current choice, and choice repetition probability in Random condition in all three ROIs identified from single trial choice decoding. Choice repetition probability was obtained by computing the proportion of trials where the choice repeated from previous trial. Lack of correlations suggest that generalization of choice representation to previous trials was not an artefact of choice repetitions across trials but due to slow fluctuations of choice signals. Correlations, Spearman's correlation coefficients. Related to Figure 2.



Figure S2: Individual choice streaks consisted of two or more consecutive trials across which a choice was repeated. Choice streaks were modelled by one regressor: a boxcar function with onset at the stimulus onset of the first trial of the streak and offset at the stimulus onset of the last trial of the streak, convolved with the canonical HRF. Illustration shows a single run, where the choice-response mapping was fixed. First and second rows shows streak regressors for left and right responses respectively; individual regressors were concatenated for illustration purposes. The GLM also consisted of two nuisance regressors to model trial-specific activation (third row), and motor response (bottom row). Vertical lines, stimulus onset; red lines, trials with a left-hand response; blue lines, trials with a right-hand response. Related to Figure 3.



Figure S3: **A.** Example scatter plot showing the between-subjects correlation between previous stimulus weights and precision scores for previous choice decoding, in r-PPC. **B, C.**. Correlation coefficients between previous stimulus weights, and precision scores for decoding previous choice and current choice, in ROIs from single trial decoding (**B**) and singe streak decoding (**C**). *, p < 0.05; **, p < 0.01; ***, p < 0.001; p-values indicate the statistical significance of Spearman's correlations. Related to Figure 4.



Figure S4: Stay (top row), and switch (second row) regressors model all the trials where the current choice repeated and alternated from the previous choice respectively, and contrasting them quantified Switch vs Stay choice behavior. The GLM consisted of a nuisance regressor (bottom row) to model motor-response. Vertical lines, stimulus onset; red lines, trials with a left-hand response; blue lines, trials with a right-hand response. Related to Figure 5.

DISCUSSION

Human decision-making is susceptible to various biases, some which arise from previous decisions themselves. These choice-induced biases are pervasive, and are observed in a variety of decision-making paradigms ranging from simple sensory decisions to more complex reasoning. Although such biases were described centuries ago, our understanding of the underlying mechanisms remains incomplete. In this thesis, I investigated the mechanistic and neural basis of choice-induced biases in decision-making.

Chapter 2 identifies the mechanistic basis of confirmation bias, a wellknown bias of human decision-making. We used a novel psychophysics task protocol that required subjects to make a continuous estimation judgment about the average motion direction across two intervals of random dot motion stimulus. The direction of dots in the two intervals were sampled independently. In addition, subjects had to report a binary categorization judgment about the direction of dots after the first interval. We found a selective enhancement of sensitivity to evidence in the second interval that is consistent with the binary choice, akin to selective attention. Our results suggest that past choices selectively direct attention towards the chosen features in subsequent evidence, producing confirmation bias in behavior. This study thus provides a new framework to investigate the neural and mechanistic basis of confirmation bias.

Chapter 3 investigates the effect of choice commitment on evidence preceding and succeeding the choice. Using the same task protocol as above, we instructed subjects to refrain from reporting the binary judgment by making a choice-neutral button press in a subset of trials. Critically, this instruction came after subjects viewed the first interval. Subjects showed reduced overall sensitivity to subsequent evidence after a choice-commitment (reporting the binary judgment), compared to the choice-neutral button press. This overall reduction in sensitivity reliably predicted the extent of choice-induced confirmation bias in a given individual, suggesting that these two effects could be mediated by a common mechanism. We observed that the relative weighting of evidence in the first vs second interval in the estimation judgment changed depending on whether the intermittent response was a choice-commitment or a choice-neutral button press: reporting an intermittent choice after the first interval boosted sensitivity to that evidence, but at the cost of reducing sensitivity to subsequent evidence. Thus, reporting a binary choice midway through the trial could effectively flip the temporal weighting of decisionrelevant evidence from recency to primacy. Similar results were also observed for the numerical judgments, generalizing the effects of choice-commitment on evidence integration.

Chapter 4 investigates the neural correlates of sequential bias in perceptual decisions. Subjects did a up vs down random dot motion direction discrimination task, while we measured their brain activity using fMRI. Unknown to the subjects, we manipulated the across-trial stimulus transition probabilities resulting in three conditions: a repeating condition where stimulus directions repeated more often across successive trials, an alternating condition where stimulus directions alternated between the two alternatives more often across successive trials, and a random condition where stimulus directions were drawn independently in each trial. We quantified behavioral metrics of sequential bias using computational models. We used multivoxel pattern analysis approaches to identify choice-selective signals across the whole cortex. Most of the analysis of neural signals in this chapter was limited to the random condition, to identify choice signals in the brain unaffected by the across-trial dependencies in the stimulus. We first identified brain regions containing choice-selective signals at a single-trial level in posterior parietal, and motor cortical regions using multivoxel pattern analysis methods across the whole cortex. The choice-selective signals thus identified in these regions generalized up to 2 choices in the past, suggesting that neural correlates of current choice also contain correlates of previous choices. We then identified sustained choice-selective signals during the inter-trial intervals between choice repetitions in a distinct set of parietal, motor and frontal brain regions. We found that behavioral metrics of sequential bias correlated with the strength of transient signals selective to current and previous choices in posterior parietal and motor regions across subjects, but not with the

sustained choice-selective signals. Our findings imply that neural correlates of sequential bias in humans can be identified in transient activation of choice signals during trial engagement in the above regions.

5.1 CHOICE-INDUCED BIASES IN DECISION-MAKING

The physical structure of the natural environment contains spatiotemporal dependencies that can be exploited by the organism to efficiently navigate their world. Indeed, several theories postulate that brains exploit this structure to constrain the computations performed in real time (Barlow 1961) by incorporating these dependencies as priors (Clark 2013; Fiser et al. 2010; Heeger 2017). Reliance on such dependencies in the form of priors may be beneficial in real-life situations where feedback about choices we make is not immediately available, or is unreliable (Varrier et al. 2019; Weilnhammer et al. 2018). However, most laboratory tasks treat trials as independent events where dependencies in choices can be detrimental to behavior (Gardner 2019).

This thesis investigated two choice-induced biases in decision-making: confirmation bias, and sequential bias. In the framework of sequential sampling models, previous choices could bias the processing of current decisionrelevant evidence either by shifting the offset (Mulder et al. 2012), or by altering the rate of the accumulation process (Ratcliff and McKoon 2008). Our analysis revealed that confirmation bias arises via a selective modulation of the gain of subsequent evidence, consistent with a multiplicative change in the rate of accumulation. Recent computational investigations showed that sequential bias arises due to a change in the rate of evidence accumulation due to previous choices (Urai et al. 2019). Neural basis of sequential bias investigated in this thesis also showed that sequential bias behavior of individual subjects relates to the transient activation of choice signals during trial engagement, but not to the sustained choice signals prior to stimulus onset. Together, these results suggest that similar computational mechanisms (i.e., a change in the rate of evidence accumulation) could underlie sequential bias and confirmation bias. Future research should verify this prediction by quantifying both biases within individual subjects using comprehensive

computational investigations.

A selective modulation of the gain of sensory evidence is commonly observed when subjects selectively attend to certain features of incoming evidence (Maunsell and Treue 2006). Thus, selective attention could be a plausible mechanism by which choices bias subsequent information (Talluri et al. 2018; Urai et al. 2019). In the brain, this effect can be realised either through a feedback of decision-related signals from higher level brain areas to low-level sensory regions (Wimmer et al. 2015), or as a selective read-out of sensory representations (Gilbert and Li 2013) corresponding to previous choices, while the incoming sensory evidence is encoded in an unbiased fashion. These two possibilities should be disentangled by future investigations by directly comparing the behavioral and neural signatures of these choice-induced biases to selective attention.

5.2 CONFIRMATION BIAS IN PERCEPTUAL DECISION-MAKING

Investigating confirmation bias in perceptual decisions, where sensory evidence underlying the decision-making process was tightly controlled, allowed us to precisely identify the effects of binary choices on the integration of subsequent decision-relevant evidence. We reported similar behavioral signatures of confirmation bias across two different decision-making paradigms: low-level perceptual decisions, and higher-level numerical integration, suggesting that perceptual decisions and high-level judgments might share similar mechanisms underlying the confirmation bias (Summerfield and Tsetsos 2012). Thus, future investigations of confirmation bias could use insights from decades of research into perceptual decisions to understand its neural basis.

Identifying the neural basis of confirmation bias also has implications in psychiatry. Confirmation bias is shown to affect the learning rates for reward prediction errors (Palminteri et al. 2017), which are signaled by the dopaminergic system in the brain (Schultz et al. 1997). Indeed, genetic polymorphisms in the dopamine system predicted individual differences in susceptibility to
confirmation bias in healthy subjects (Doll et al. 2011). Furthermore, patients with schizophrenia and psychosis were less prone to rely on prior information, making them less susceptible to confirmation bias than healthy controls (Doll et al. 2014; Jardri et al. 2017; Weilnhammer et al. 2020). Thus, identifying the neural and mechanistic basis of confirmation bias could help further our understanding of these psychiatric disorders. An interesting future direction is to use novel task protocols such as the one used in this thesis to identify behavioral markers in people susceptible to schizophrenia and psychosis.

5.3 TEMPORAL WEIGHTING OF PERCEPTUAL EVIDENCE

Previous empirical studies reported myriad differences in the nature of temporal weighting, ranging from primacy, to recency, to non-monotonic weights. While these studies differ in a number of ways, none of them convincingly explain which factors underlie different weighting profiles. Computational investigations using the LCA, for example, showed that the dynamic interplay between leak, and mutual inhibition in the decision-making networks produces empirically observed profiles: an inhibition-dominant state produces primacy, a leak dominant state produces recency, a perfect balance between leak and inhibition produces flat weighting, and a dynamic time-dependent change in the balance produces non-monotonic weights (Bronfman et al. 2016). Using a novel experimental protocol presented in this thesis, we were able to produce flexible temporal weighting behavior within each individual, by changing the nature of intermittent response, keeping all other variables like stimulus duration, and the type of stimulus unchanged in two different experiments. We observed a phasic increase in pupil diameter following choice-commitment in perceptual decisions. The increase in pupil size could reflect a transient gain modulation across the cortex through phasic activation of central arousal systems (Aston-Jones and Cohen 2005). Similar gain modulations at choice-commitment have been observed in neurophysiological recordings of choice-selective neurons in interrogation designs (Niyogi and Wong-Lin 2013; Roitman and Shadlen 2002). Thus, choice-commitment acts as an internal signal modulating the state of the decision-making networks in the brain.

5.4 CONCLUDING REMARKS

In this thesis, using novel psychophysics task protocols, I have shown that choices are not merely the end points of a deliberation, but act as contextual factors biasing subsequent decisions. Identifying choice-induced biases in low-level perceptual decisions suggests that these biases are integral to decision-making, and comprehensive investigations into the biases in perceptual decisions provides insights into the fundamental mechanisms of decision-making. The insights presented in this thesis advance the idea that choices act as cues to selectively redirect attention towards the chosen features in subsequent evidence. The precise link between selective attention and choice-induced biases in decision-making needs to be systematically explored by future investigations.

GENERAL SUMMARY

ENGLISH SUMMARY

Our decisions are influenced by biases, some of which arise from our past choices themselves. Real-world decisions are embedded in sequences of other decisions, making them even more susceptible to such choice-induced biases. These choice-induced biases manifest in many forms. In this thesis, I investigated the mechanistic basis of such biases by studying how intermittent decisions influence subsequent decision-relevant evidence, and how neural signals sensitive to previous choices affect current choices. Understanding the mechanistic and neural basis of these biases provides crucial insights into the fundamental processes underlying decision-making. These questions were addressed in perceptual decisions, a class of decisions that require making inferences using simple sensory information, like judging the average direction of a stream of dots displayed on a computer screen.

Confirmation bias is a well-known and pervasive bias of human reasoning, where they tend to ignore evidence that doesn't agree with their beliefs. We tested if humans show confirmation bias for simple sensory evidence, and thereby establish a lower-limit to the bias. Human subjects saw a stream of moving dots on a computer screen, and reported a category judgment whether on average the dots moved to the left, or right of a reference line on the screen. They then saw a second stream of moving dots, which might move in a slightly different direction. Finally, they estimated the average direction across the two streams as precisely as possible. We found that the direction of dots in the second stream consistent with the initially chosen category had a greater influence on the estimations, and directions inconsistent with the category had little influence. Our findings suggest that choices selectively direct attention towards the chosen alternatives in the upcoming evidence. These findings were generalized to judgments of symbolic numerical evidence, establishing the robustness of the results. Observing confirmation bias in simple sensory decisions that do not carry any meaning to the subjects suggests that this bias is fundamental to our decision-making behavior.

In the same task above, subjects were asked to refrain from reporting their category judgment after the first stream in a few trials. We then compared the two conditions within the same group of subjects: overt commitment vs noncommitment to the categorical judgment. We found that choice commitment was followed by an increase in pupil size, and a decrease in overall sensitivity to evidence in the second stream. The increase in pupil size reflects a transient release of neuromodulatory signals that possibly push the decisionmaking networks in the brain to a state less sensitive to new evidence. On the other hand, with-holding from choice commitment made subjects more sensitive to the second stream, but at the expense of reduced sensitivity to the first stream. Thus, commitment to a categorical proposition alters the way in which previous and subsequent decision-relevant evidence is processed. Furthermore, the choice-induced overall reduction in sensitivity in a given individual reliably predicts the extent of choice-induced confirmation bias in the same individual, suggesting that these two biases could be mediated by a common mechanism.

Sequential bias is another form of choice-induced bias commonly observed in laboratory tasks of decision-making where previous decisions tend to affect current decisions, even when subjects know that past decisions do not influence current sensory information. We measured the brain activity of human subjects making simple perceptual choices, and related the neural activity to their choice behavior. We identified two sets of brain regions with choice-specific signals: one set that contained transient choice-specific signals about the current and previous choices, the strength of which predicted the subjects' sequential bias; and another set that contained sustained choice-specific signals even in the absence of sensory information whenever subjects repeated their previous choices. Taken together, the results presented in this thesis suggest that rather than being the end-points of a deliberation, choices act as contextual factors biasing how we perceive the world at a fundamental level.

DEUTSCHE ZUSAMMENFASSUNG¹

Unsere Entscheidungen werden durch kognitive Verzerrungen (engl. "cognitive bias") beeinflusst, von denen einige aus unseren vergangenen Entscheidungen selbst entstehen. Im Alltag sind Entscheidungen in eine Reihe von anderen Entscheidungen eingebettet, was sie noch anfälliger für solche, durch Entscheidungen ausgelöste Verzerrungen, macht. Diese entscheidungsinduzierten Verzerrungen manifestieren sich in unterschiedlichen Formen.

In dieser Dissertation habe ich die mechanistische Grundlage solcher kognitiver Verzerrungen erforscht. Dazu habe ich untersucht, wie zwischenzeitliche Entscheidungen anschließende, entscheidungsrelevante Information beeinflussen und wie neuronale Signale, die sensibel für vorangegangene Entscheidungen sind, gegenwärtige Entscheidungen beeinflussen. Das Verständnis für die mechanistische und neuronale Grundlage kognitiver Verzerrungen gibt wesentliche Einblicke in die fundamentalen Prozesse, die der Entscheidungsfindung zugrundeliegen. Diesen Fragen wurde anhand perzeptueller Entscheidungen nachgegangen, einer Form von Entscheidungen, die Schlussfolgerungen aus einfachen sensorischen Informationen, wie zum Beispiel der Beurteilung der durchschnittlichen Richtung von flimmernden Punkten in einem visuellen Stimulus, erfordert.

Die Bestätigungstendenz (engl. "confirmation bias") ist eine bekannte und allgegenwärtige kognitive Verzerrung, wobei Information, die nicht mit den eigenen Ansichten übereinstimmt, ignoriert wird. Wir haben getestet, ob Menschen die Bestätigungstendenz bei einfacher sensorischer Information zeigen und eine Untergrenze für diese Verzerrung etabliert.

Den Probanden² wurden dabei visuelle Stimuli in Form von flimmernden Punkten auf einem Computerbildschirm gezeigt und sie sollten kategorisch beurteilen, ob sich die Punkte in Bezug auf eine Referenzlinie auf dem Bildschirm im Durchschnitt nach rechts oder nach links bewegten. Danach sahen sie einen zweiten Stimulus aus flimmernden Punkten, welche sich in eine geringfügig andere Richtung bewegen konnten. Schließlich beurteilten die Probanden die durchschnittliche Richtung der Punkte beider Stimuli so genau

¹ Thanks to Gina Monov and Thomas Pfeffer for help with translation.

² Es sind stets Personen männlichen und weiblichen Geschlechts gleichermaßen gemeint; aus Gründen der einfacheren Lesbarkeir wird im Folgenden nur die männliche Form verwendet

wie möglich. Wir konnten zeigen, dass die Richtung der Punkte im zweiten Stimulus einen größeren Einfluss auf die Beurteilung hatte, wenn sie mit der zuerst gewählten Kategorie konsistent war, während ihr Einfluss gering war, wenn sie mit der zuerst gewählten Kategorie inkonsistent war. Unsere Ergebnisse deuten darauf hin, dass Entscheidungen unsere Aufmerksamkeit bei aufkommender Information selektiv auf die zuvor gewählte Alternative lenken. Die Robustheit dieser Ergebnisse wurde durch die Verallgemeinerung auf die Beurteilung von symbolischer, numerischer Evidenz gezeigt. Die Beobachtung der Bestätigungstendenz in einfachen, sensorischen Entscheidungen, die keine Bedeutung für die Testpersonen haben, legt nahe, dass diese Form kognitiver Verzerrung ein grundlegender Bestandteil unseres Entscheidungsverhaltens ist.

In der gleichen Aufgabe wie oben sollten Probanden bei einigen Versuchen die Angabe eines kategorischen Urteils nach dem ersten Stimulus unterlassen. Anschließend haben wir die zwei Bedingungen innerhalb derselben Gruppe von Probanden verglichen: offenkundige Festlegung vs. keine Festlegung auf ein kategorisches Urteil. Wir haben herausgefunden, dass die Festlegung auf eine Entscheidung zu einer Zunahme der Pupillengröße und einer Abnahme der allgemeinen Empfindlichkeit gegenüber der Information im zweiten Stimulus führt. Die Pupillenerweiterung reflektiert eine vorübergehende Freisetzung neuromodulatorischer Signale im Gehirn, die möglicherweise Netzwerke der Entscheidungsfindung in einen Zustand versetzen, der weniger empfindlich gegenüber neuer Evidenz ist. Andererseits machte das Zurückhalten der Festlegung auf eine Entscheidung die Probanden sensibler gegenüber dem zweiten Stimulus, allerdings auf Kosten von reduzierter Sensibilität für den ersten Stimulus. Demnach verändert die Festlegung auf eine kategorische Aussage die Art und Weise, in der vorangegangene und nachfolgende entscheidungsrelevante Information verarbeitet wird. Außerdem kann die, durch die Entscheidung hervorgerufene, allgemeine Verringerung der Sensibilität eines Individuums das Ausmaß der entscheidungsinduzierten Bestätigungstendenz desselben Individuums zuverlässig vorhersagen, was vermuten lässt, dass die beiden kognitiven Verzerrungen durch einen gemeinsamen Mechanismus vermittelt sein können.

Die sequentielle kognitive Verzerrung (engl. "sequential bias") ist eine weitere Form entscheidungsinduzierter Verzerrung, die häufig unter experimentellen Bedingungen beobachtet wird, wobei Entscheidungen in einem vorangegangenen, unabhängigen Durchgang (engl. "trial") der Aufgabe dazu neigen, die Entscheidung im aktuellen Durchgang zu beeinflussen – auch wenn die Probanden wissen, dass ihre vorherige Entscheidung keine Auswirkung auf die aktuelle sensorische Information hat. Wir haben die Gehirnaktivität von Menschen während einfacher Wahrnehmungsentscheidungen gemessen und ihre neuronale Aktivität in Bezug zu ihrem Entscheidungsverhalten gesetzt. Wir haben zwei Gruppen von Hirnregionen mit entscheidungsspezifischer Aktivität identifiziert: die eine beinhaltet vorübergehende entscheidungsspezifische Signale über die aktuelle und vorangegangene Entscheidung, dessen Stärke die sequentiellen Verzerrung der Probanden vorhersagte; die andere beinhaltetet anhaltende entscheidungsspezifische Signale, auch in der Abwesenheit von sensorischer Information, immer wenn Probanden ihre vorherigen Entscheidungen wiederholten.

Zusammengefasst legen die in dieser Dissertation dargestellten Ergebnisse nahe, dass Entscheidungen weniger der Endpunkt einer Überlegung sind, sondern dass sie eher als kontextuelle Faktoren wirken, die verzerren, wie wir die Welt wahrnehmen.

BIBLIOGRAPHY

- Abraham A et al. (2014). Machine learning for neuroimaging with scikit-learn. *Front Neuroinform* 8:14.
- Abrahamyan A et al. (2016). Adaptable history biases in human perceptual decisions. *PNAS* 113:E3548–E3557.
- Adelson E and Bergen J (1985). Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A, JOSAA* 2:284–299.
- Aguirre G, Zarahn E, and D'Esposito M (1998). The Variability of Human, BOLD Hemodynamic Responses. *NeuroImage* 8:360–369.
- Akaishi R, Umeda K, Nagase A, and Sakai K (2014). Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia. *Neuron* 81:195–206.
- Akrami A, Kopec C, Diamond M, and Brody C (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* 554:368– 372.
- Aston-Jones G and Cohen J (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience* 28:403–450.
- Barber B (1961). Resistance by Scientists to Scientific Discovery. Science 134:596-602.
- Barlow H (1961). Possible principles underlying the transformations of sensory messages. *Sensory Communication*. Cambridge, MA: MIT Press, pp. 217–34.
- van Bergen R and Jehee J (2019). Probabilistic Representation in Human Visual Cortex Reflects Uncertainty in Serial Decisions. *J Neurosci* 39:8164–8176.
- Bertelson P (1965). Serial Choice Reaction-time as a Function of Response versus Signal-and-Response Repetition. *Nature* 206:217–218.
- Bogacz R and Cohen J (2004). Parameterization of connectionist models. *Behavior Research Methods, Instruments, & Computers* 36:732–741.
- Bogacz R et al. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* 113:700–765.
- Bonaiuto J, de Berker A, and Bestmann S (2016). Response repetition biases in human perceptual decisions are explained by activity decay in competitive attractor models. *eLife* 5:e20047.

- Botvinick M et al. (2001). Conflict monitoring and cognitive control. *Psychological Review* 108:624–652.
- Brascamp J et al. (2008). Multi-Timescale Perceptual History Resolves Visual Ambiguity. *PLoS ONE* 3:e1497.
- Braun A, Urai A, and Donner T (2018). Adaptive History Biases Result from Confidence-Weighted Accumulation of past Choices. *J Neurosci* 38:2418–2429.
- Brehm J (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology* 52:384–389.
- Breton-Provencher V and Sur M (2019). Active control of arousal by a locus coeruleus GABAergic circuit. *Nat Neurosci* 22:218–228.
- Britten K, Shadlen M, Newsome W, and Movshon J (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Britten K et al. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience* 13:87–100.
- Bronfman Z, Brezis N, and Usher M (2016). Non-monotonic Temporal-Weighting Indicates a Dynamically Modulated Evidence-Integration Mechanism. *PLOS Comput Biol* 12:e1004667.
- Bronfman Z et al. (2015). Decisions reduce sensitivity to subsequent information. *Proc R Soc B Biol Sci* 282:20150228.
- Bruce C and Goldberg M (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *Journal of Neurophysiology* 53:603–635.
- Brunton B, Botvinick M, and Brody C (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340:95–98.
- Busse L et al. (2011). The Detection of Visual Contrast in the Behaving Mouse. *J Neurosci* 31:11351–11361.
- Chambon V et al. (2019). Choosing and learning: outcome valence differentially affects learning from free versus forced choices. *bioRxiv*:637157.
- Cheadle S et al. (2014). Adaptive Gain Control during Human Perceptual Choice. *Neuron* 81:1429–1441.
- Chen M and Risen J (2010). How choice affects and reflects preferences: revisiting the free-choice paradigm. *Journal of Personality and Social Psychology* 99:573–594.
- Christopoulos V, Kagan I, and Andersen R (2018). Lateral intraparietal area (LIP) is largely effector-specific in free-choice decisions. *Scientific Reports* 8:1–13.
- Cisek P, Puskas G, and El-Murr S (2009). Decisions in changing conditions: the urgencygating model. *The Journal of Neuroscience* 29(37):11560–11571.
- Clark A (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36(3):181–204.

- Cole M and Schneider W (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *NeuroImage* 37:343–360.
- Cross D (1973). Sequential dependencies and regression in psychophysical judgments. *Perception & Psychophysics* 14:547–552.
- Dale A (1999). Optimal experimental design for event-related fMRI. *Human brain mapping* 8:109–114.
- Desender K, Boldt A, and Yeung N (2018). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychological Science* 29:761–778.
- Ding L and Gold J (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral Cortex* 22:1052–1067.
- Doll B, Hutchison K, and Frank M (2011). Dopaminergic Genes Predict Individual Differences in Susceptibility to Confirmation Bias. *J Neurosci* 31:6188–6198.
- Doll B et al. (2014). Reduced susceptibility to confirmation bias in schizophrenia. *Cogn Affect Behav Neurosci* 14:715–728.
- Donner T, Sagi D, Bonneh Y, and Heeger D (2008). Opposite neural signatures of motion-induced blindness in human dorsal and ventral visual cortex. *The Journal of Neuroscience* 28:10298–10310.
- Donner T, Siegel M, Fries P, and Engel A (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology* 19:1581– 1585.
- Drugowitsch J, Wyart V, Devauchelle AD, and Koechlin E (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron* 92:1–14.
- Drugowitsch J, Mendonça AG, Mainen ZF, and Pouget A (2019). Learning optimal decisions with confidence. *PNAS* 116(49):24872–24880.
- Eckhoff P, Wong-Lin K, and Holmes P (2009). Optimality and Robustness of a Biophysical Decision-Making Model under Norepinephrine Modulation. *J Neurosci* 29:4301–4311.
- Efron B and Tibshirani R (1986). Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. *Statistical Science* 1:54–75.
- Eisenstein E (1975). Aneural Systems and Neurobiology: A Point of View. *Aneural Organisms in Neurobiology*. Ed. by Eisenstein E and ed, pp. 1–3.
- Eklund A, Nichols T, and Knutsson H (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proc Natl Acad Sci USA* 113:7900–7905.
- Fernberger S (1920). Interdependence of judgments within the series for the method of constant stimuli. *Journal of Experimental Psychology* 3:126–150.

- Festinger L (1957). *A theory of cognitive dissonance*. Stanford: Stanford University Press.
- Findling C et al. (2019). Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nat Neurosci* 22:2066–2077.
- Fischer J and Whitney D (2014). Serial dependence in visual perception. *Nature Neuroscience* 17:738–743.
- Fischhoff B and Beyth-Marom R (1983). Hypothesis evaluation from a Bayesian perspective. *Psychological Review* 90:239–260.
- Fischl B, Sereno M, and Dale A (1999). Cortical Surface-Based Analysis: II: Inflation, Flattening, and a Surface-Based Coordinate System. *NeuroImage* 9:195–207.
- Fiser J, Berkes P, Orban G, and Lengyel M (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci* 14:119–130.
- Fleming S, Putten E, and Daw N (2018). Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience* 21:617–624.
- Fritsche M and de Lange F (2019). The role of feature-based attention in visual serial dependence. *Journal of Vision* 19:21–21.
- Fritsche M, Mostert P, and de Lange F (2017). *Opposite Effects of Recent History on Perception and Decision. Current Biology* 27:590–595.
- Fründ I, Wichmann F, and Macke J (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision* 14(9).
- Gardner J (2019). Optimality and heuristics in perceptual neuroscience. *Nat. Neurosci.* 22:514–523.
- de Gee J, Knapen T, and Donner T (2014). Decision-related pupil dilation reflects upcoming choice and individual bias. *PNAS* 111:E618–625.
- de Gee J et al. (2017). Dynamic modulation of decision biases by brainstem arousal systems. *eLife* 6:e23232.
- de Gee J et al. (2019). Phasic arousal suppresses biases in mice and humans across domains of decision-making. *bioRxiv*:447656.
- Gilbert C and Li W (2013). Top-down Influences on Visual Processing. *Nature Reviews Neuroscience* 14(5):350–363.
- Glasser M et al. (2016). A multi-modal parcellation of human cerebral cortex. *Nature* 536:171–178.
- Glaze C, Kable J, and Gold J (2015). Normative evidence accumulation in unpredictable environments. *eLife* 4:1–26.
- Glaze C et al. (2018). A bias-variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour* 2:213–224.

- Gold J, Law CT, Connolly P, and Bennur S (2008). The Relative Influences of Priors and Sensory Evidence on an Oculomotor Decision Variable During Perceptual Learning. *Journal of Neurophysiology* 100:2653–2668.
- Gold J and Shadlen M (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394.
- Gold J and Shadlen M (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16.
- Gold J and Shadlen M (2003). The Influence of Behavioral Context on the Representation of a Perceptual Decision in Developing Oculomotor Commands. *J Neurosci* 23:632–651.
- Gold J and Shadlen M (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience* 30:535–574.
- Gold J and Stocker A (2017). Visual Decision-Making in an Uncertain and Dynamic World. *Annual Review of Vision Science* 3:227–250.
- Gorgolewski K et al. (2011). Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python. *Frontiers in Neuroinformatics* 5:13.
- Goris R et al. (2017). Dissociation of Choice Formation and Choice-Correlated Activity in Macaque Visual Cortex. *The Journal of Neuroscience* 37:5195–5203.
- Green D and Swets J (1966). *Signal Detection Theory and Psychophysics*. John Wiley and Sons.
- Haefner R, Berkes P, and Fiser J (2016). Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron* 90:649–660.
- Handwerker D, Ollinger J, and D'Esposito M (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *NeuroImage* 21:1639–1651.
- Hanks T, Ditterich J, and Shadlen M (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience* 9:682–689.
- Hayden B, Pearson J, and Platt M (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939.
- Haynes JD and Rees G (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience* 7:523–534.
- Hebart M, Donner T, and Haynes JD (2012). Human visual and parietal cortex encode visual choices independent of motor plans. *NeuroImage* 63:1393–1403.
- Hebart M, Schriever Y, Donner T, and Haynes JD (2016). The Relationship between Perceptual Decision Variables and Confidence in the Human Brain. *Cerebral Cortex* 26:118–130.
- Heeger D (2017). Theory of cortical function. PNAS 114(8):1773-1782.

- Heekeren H et al. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *PNAS* 103:10023–10028.
- Hermoso-Mendizabal A et al. (2020). Response outcomes gate the impact of expectations on perceptual decisions. *Nature Communications* 11:1–13.
- Herrmann K, Heeger D, and Carrasco M (2012). Feature-based attention enhances performance by increasing response gain. *Vision Research* 74:10–20.
- Honey C et al. (2012). Slow Cortical Dynamics and the Accumulation of Information over Long Timescales. *Neuron* 76:423–434.
- Huettel S, Song A, and McCarthy G (2009). *Functional Magnetic Resonance Imaging*. 2nd ed. Massachusetts: Sinauer.
- Huk A and Shadlen M (2005). Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. *J Neurosci* 25:10420–10436.
- Hwang E, Dahlen J, Mukundan M, and Komiyama T (2017). History-based action selection bias in posterior parietal cortex. *Nature Communications* 8(1242).
- Izuma K and Murayama K (2013). Choice-induced preference change in the freechoice paradigm: a critical methodological review. *Front. Psychology* 4:41.
- Jardri R, Duverne S, Litvinova A, and Denève S (2017). Experimental evidence for circular inference in schizophrenia. *Nat Commun* 8:14218.
- Jazayeri M and Movshon J (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience* 9:690–696.
- Jazayeri M and Movshon J (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* 446:912–915.
- Jazayeri M and Movshon J (2010). Integration of sensory evidence in motion discrimination. *Journal of Vision* 7:1–13.
- Jenkinson M, Bannister P, Brady M, and Smith S (2002). Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage* 17:825–841.
- Jenkinson M and Smith S (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* 5:143–156.
- Jenkinson M et al. (2012). FSL. NeuroImage 62:782–790.
- Joshi S, Li Y, Kalwani R, and Gold J (2016). Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron* 89:1–14.
- Kable J and Glimcher P (2009). The Neurobiology of Decision: Consensus and Controversy. *Neuron* 63:733–745.
- Kahneman D (2011). Thinking, Fast and Slow. New York: Farrar, Straus and Giroux.

- Kahneman D and Tversky A (1973). On the psychology of prediction. *Psychological Review* 80:237–251.
- Kanai R and Verstraten F (2005). Perceptual manifestations of fast neural plasticity: Motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Research* 45:3109–3116.
- Kass R and Raftery A (1995). Bayes Factors. *Journal of the American Statistical Association* 90:773–795.
- Katz L, Yates J, Pillow J, and Huk A (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* 535:285–288.
- Kelly S and O'Connell R (2013). Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. *J Neurosci* 33:19434–19441.
- Keuken M et al. (2014). Brain networks of perceptual decision-making: an fMRI ALE meta-analysis. *Front Hum Neurosci* 8:445.
- Kiani R, Cueva C, Reppas J, and Newsome W (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology* 24:1542–1547.
- Kiani R, Hanks T, and Shadlen M (2008). Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment. *J Neurosci* 28:3017–3029.
- Kim J and Shadlen M (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Kim T, Kabir M, and Gold J (2017). Coupled Decision Processes Update and Maintain Saccadic Priors in a Dynamic Environment. *J Neurosci* 37:3632–3645.
- Kiyonaga A, Scimeca J, Bliss D, and Whitney D (2017). Serial Dependence across Perception, Attention, and Memory. Trends Cogn Sci 27:590–595.
- Kleiner M et al. (2007). What's new in Psychtoolbox-3. Perception 36:1-16.
- Kohn A (2007). Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology* 97:3155–3164.
- Kolling N, Behrens T, Mars R, and Rushworth M (2012). Neural Mechanisms of Foraging. *Science* 336:95–98.
- Krekelberg B, Boynton G, and van Wezel R (2006). Adaptation: from single cells to BOLD signals. *Trends in Neurosciences* 29:250–256.
- Kriegeskorte N, Goebel R, and Bandettini P (2006). Information-based functional brain mapping. *PNAS* 103:3863–3868.
- Kuhn D (1989). Children and adults as intuitive scientists. *Psychological Review* 96:674–689.
- Lak A et al. (2020). Reinforcement biases subsequent perceptual decisions when confidence is low: a widespread behavioral phenomenon. *eLife* 9:e49834.

- de Lange F, Rahnev D, Donner T, and Lau H (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *The Journal of Neuroscience* 33:1400–1410.
- Liberman A, Fischer J, and Whitney D (2014). Serial Dependence in the Perception of Faces. *Current Biology* 24:2569–2574.

Lindquist M (2008). The Statistical Analysis of fMRI Data. Statist Sci 23:439-464.

- Liu F and Wang XJ (2008). A Common Cortical Circuit Mechanism for Perceptual Categorical Discrimination and Veridical Judgment. *PLoS Computational Biology* 4:e1000253.
- Logothetis N, Pauls J, and Augath M (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- Lueckmann JM, Macke J, and Nienborg H (2018). Can Serial Dependencies in Choices and Neural Activity Explain Choice Probabilities? *J Neurosci* 38:3495–3506.
- Luo T and Maunsell J (2015). Neuronal Modulations in Visual Cortex Are Associated with Only One of Multiple Components of Attention. *Neuron* 86:1182–1188.
- Luo T and Maunsell J (2018). Attentional Changes in Either Criterion or Sensitivity Are Associated with Robust Modulations in Lateral Prefrontal Cortex. *Neuron* 97:1382– 1393.
- Luu L and Stocker A (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife* 7:e33334.
- Macke J and Nienborg H (2019). Choice (-history) correlations in sensory cortex: cause or consequence? *Current Opinion in Neurobiology* 58:148–154.
- Marcos E et al. (2013). Neural Variability in Premotor Cortex Is Modulated by Trial History and Predicts Behavioral Performance. *Neuron* 78:249–255.
- Maunsell J and van Essen V. D (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *Journal of Neurophysiology* 49:1127–1147.
- Maunsell J and Treue S (2006). Feature-based attention in visual cortex. *Trends in Neurosciences* 29:317–322.

Meyniel F, Maheu M, and Dehaene S (2016). Human Inferences about Sequences: A Minimal Transition Probability Model. *PLOS Computational Biology* 12:e1005260.

- Miller E and Cohen J (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience* 24:167–202.
- Mountcastle V, Talbot W, Sakata H, and Hyvärinen J (1969). Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *Journal of Neurophysiology* 32:452–484.
- Mulder M et al. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience* 32(7):2335–2343.

- Mumford J, Turner B, Ashby F, and Poldrack R (2012). Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *NeuroImage* 59:2636–2643.
- Murphy P et al. (2020). Normative Circuit Dynamics Across Human Cortex During Evidence Accumulation in Changing Environments. *bioRxiv*:924795.
- Murray J et al. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience* 17:1661–1663.
- Newsome W, Britten K, and Movshon J (1989). Neuronal correlates of a perceptual decision. *Nature* 341:52–54.
- Nickerson R (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2:175–220.
- Nienborg H and Cumming B (2010). Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Current Opinion in Neurobiology* 20:376–381.
- Nienborg H and Cumming B (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459:89–92.
- Nienborg H, R. Cohen M, and Cumming B (2012). Decision-Related Activity in Sensory Neurons: Correlations Among Neurons and with Behavior. *Annual Review of Neuroscience* 35:463–483.
- Nienborg H and Roelfsema P (2015). Belief states as a framework to explain extraretinal influences in visual cortex. *Current Opinion in Neurobiology* 32:45–52.
- Niyogi R and Wong-Lin K (2013). Dynamic Excitatory and Inhibitory Gain Modulation Can Produce Flexible, Robust and Optimal Decision-making. *PLOS Computational Biology* 9:e1003099.
- Norman K, Polyn S, Detre G, and Haxby J (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences* 10(9):424–430.
- Odoemene O, Pisupati S, Nguyen H, and Churchland A (2018). Visual Evidence Accumulation Guides Decision-Making in Unrestrained Mice. *J Neurosci* 38:10143– 10155.
- Oostenveld R, Fries P, Maris E, and Schoffelen JM (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational intelligence and neuroscience* 2011(156869).
- Ossmy O et al. (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Current Biology* 23:981–986.
- Oswald M and Grosjean S (2004). Confirmation bias. *Cognitive Illusions. A Handbook on Fallacies and Biases in Thinking*. Ed. by Pohl R. New York: Psychology Press.
- Palminteri S, Lefebvre G, Kilford E, and Blakemore SJ (2017). Confirmation bias in human reinforcement learning: Evidence from counterfactual feedback processing. *PLOS Computational Biology* 13:e1005684.

- Pape AA and Siegel M (2016). Motor cortex activity predicts response alternation during sensorimotor decisions. *Nature Communications* 7(13098).
- Pearson J and Brascamp J (2008). Sensory memory for ambiguous vision. Trends in Cogn Sci 12:334–341.
- Pedregosa F et al. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Philiastides M, Auksztulewicz R, Heekeren H, and Blankenburg F (2011). Causal Role of Dorsolateral Prefrontal Cortex in Human Perceptual Decision Making. *Current Biology* 21:980–983.
- Popper K (1959). *The Logic of Scientific Discovery*. Hutchinson.
- Preuschhof C et al. (2006). Neural Correlates of Vibrotactile Working Memory in the Human Brain. *J Neurosci* 26:13231–13239.
- Pyszczynski T and Greenberg J (1987). Toward an Integration of Cognitive and Motivational Perspectives on Social Inference: A Biased Hypothesis-Testing Model. *Advances in Experimental Social Psychology*. Ed. by Berkowitz L. Academic Press, pp. 297–340.
- Quilodran R, Rothé M, and Procyk E (2008). Behavioral Shifts and Action Valuation in the Anterior Cingulate Cortex. *Neuron* 57:314–325.
- Rahnev D and Denison R (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences* 223:1–66.
- Raposo D, Kaufman M, and Churchland A (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience* 17:1784–1792.
- Ratcliff R (1978). A theory of memory retrieval. Psychological Review 85:59–108.
- Ratcliff R and McKoon G (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation* 20:873–922.
- Raviv O, Ahissar M, and Loewenstein Y (2012). How Recent History Affects Perception: The Normative Approach and Its Heuristic Approximation. *PLOS Computational Biology* 8:e1002731.
- Reimer J et al. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature Communications* 7:13289.
- Reynolds J and Heeger D (2009). The Normalization Model of Attention. *Neuron* 61:168–185.
- Roitman J and Shadlen M (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475– 9489.
- Rolls E and Deco G (2010). The Noisy Brain. Oxford University Press.
- Rollwage M, Dolan R, and Fleming S (2018). Metacognitive Failure as a Feature of Those Holding Radical Beliefs. *Current Biology* 28:4014–4021.

- Romo R et al. (2000). Sensing without Touching: Psychophysical Performance Based on Cortical Microstimulation. *Neuron* 26:273–278.
- Romo R et al. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci* 5:1217–1225.
- Ross L and Anderson C (1982). Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments. *Judgment under Uncertainty*.
 Ed. by Kahneman D, Slovic P, and Tversky A. Cambridge University Press, pp. 129–152.

Rowan T (1990). Functional Stability Analysis Of Numerical Algorithms.

- Rushworth M and Behrens T (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience* 11:389–397.
- Salzman C, Murasugi C, Britten K, and Newsome W (1992). Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12:2331–2355.
- Sarafyazd M and Jazayeri M (2019). Hierarchical reasoning by neural circuits in the frontal cortex. *Science* 364:(6441), eaav8911.
- Scase M, Braddick O, and Raymond J (1996). What is Noise for the Motion System? *Vision Research* 36:2579–2586.
- Schultz W, Dayan P, and Montague P (1997). A Neural Substrate of Prediction and Reward. *Science* 275:1593–1599.
- Schwarz G (1978). Estimating the Dimension of a Model. *The Annals of Statistics* 6:461–464.
- Scott B et al. (2017). Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron* 31:5989–6000.
- Shadlen M and Kiani R (2013). Decision making as a window on cognition. *Neuron* 80:791–806.
- Shadlen M and Newsome W (1996). Motion perception: seeing and deciding. *PNAS* 93:628–633.
- Shadlen M and Newsome W (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Siegel M, Buschman T, and Miller E (2015). Cortical information flow during flexible sensorimotor decisions. *Science* 348:1352–1355.
- Siegel M et al. (2007). High-Frequency Activity in Human Visual Cortex Is Modulated by Visual Motion Strength. *Cereb Cortex* 17:732–741.
- Smith S (2004). Overview of fMRI analysis. The British Journal of Radiology 77:S167–S175.

St John-Saaltink E, Kok P, Lau H, and de Lange F (2016). Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex. *J Neurosci* 36:6186–6192.

- Stocker A and Simoncelli E (2007). A Bayesian Model of Conditioned Perception. *Advances in neural information processing systems 2007*, pp. 1409–1416.
- Summerfield C and Tsetsos K (2012). Building Bridges between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms. *Frontiers in Neuroscience* 6:70.
- Talluri B et al. (2018). Confirmation Bias through Selective Overweighting of Choice-Consistent Evidence. *Current Biology* 28:3128–3135.
- Tegenthoff M et al. (2005). Improvement of Tactile Discrimination Performance and Enlargement of Cortical Somatosensory Maps after 5 Hz rTMS. *PLOS Biology* 3:e362.
- Thura D, Beauregard-Racine J, Fradet CW, and Cisek P (2012). Decision making by urgency gating: theory and experimental support. *Journal of neurophysiology* 108(11):2912–2930.
- Tsetsos K, Chater N, and Usher M (2012). Salience driven value integration explains decision biases and preference reversal. *PNAS* 109:9659–9664.
- Tsetsos K et al. (2016). Economic irrationality is optimal during noisy decision making. *PNAS* 113:3102–3107.
- Tsunada J, Cohen Y, and Gold J (2019). Post-decision processing in primate prefrontal cortex influences subsequent choices on an auditory decision-making task. *eLife* 8:46770.
- Tuchman B (1984). *The March of Folly: From Troy to Vietnam*. Random House Publishing Group.
- Tversky A and Kahneman D (1974). Judgment under Uncertainty: Heuristics and Biases. *Science* 185:1124–1131.
- Urai A, Braun A, and Donner T (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nature Communications* 8(14637).
- Urai A, de Gee J, Tsetsos K, and Donner T (2019). Choice history biases subsequent evidence accumulation. *eLife* 8:46331.
- Urai A and Wimmer K (2016). Spatiotemporal motion energy filtering: a Matlab implementation. Zenodo.
- Usher M and McClelland J (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review* 108:550–592.
- Varrier R, Stuke H, Guggenmos M, and Sterzer P (2019). Sustained effects of corrupted feedback on perceptual inference. *Sci Rep* 9:5537.
- van Veen V, Krug M, Schooler J, and Carter C (2009). Neural activity predicts attitude change in cognitive dissonance. *Nat Neurosci* 12:1469–1474.
- Wang XJ (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron* 36:955–968.
- Wang XJ (2008). Decision Making in Recurrent Neuronal Circuits. Neuron 60:215-234.

- Wason P (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology* 12:129–140.
- Watson A and Pelli D (1983). Quest: A Bayesian adaptive psychometric method. *Perception & Psychophysics* 33:113–120.
- Weilnhammer V, Stuke H, Sterzer P, and Schmack K (2018). The Neural Correlates of Hierarchical Predictions for Perceptual Decisions. *Journal of Neuroscience* 38(21):5008–5021.
- Weilnhammer V et al. (2020). Psychotic Experiences in Schizophrenia and Sensitivity to Sensory Evidence. *Schizophrenia Bulletin* sbaa003.
- Wichmann F and Hill N (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics* 63:1293–1313.
- Wilming N, Murphy P, Meyniel F, and Donner T (2020). Disentangling Decision-related Feedforward and Feedback Signals in Human Cortex. *bioRxiv*:929893.
- Wimmer K et al. (2015). Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat Commun* 6:6177.
- Wong KF and Wang XJ (2006). A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *J Neurosci* 26:1314–1328.
- Wyart V, de Gardelle V, Scholl J, and Summerfield C (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* 76:847–858.
- Yu A and Cohen J (2009). Sequential effects: Superstition or rational behavior? *Advances in Neural Information Processing Systems 21*, pp. 1873–1880.
- Zamboni E, Ledgeway T, McGraw P, and Schluppeck D (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B: Biological Sciences* 283:20160263.
- Zuckerman M, Knee C, Hodgins H, and Miyake K (1995). Hypothesis confirmation: The joint effect of positive test strategy and acquiescence response set. *Journal of Personality and Social Psychology* 68:52–60.
- Zylberberg A, Barttfeld P, and Sigman M (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience* 6:79.

ACKNOWLEDGMENTS

This thesis is not just the result of *my* scientific work over the past few years, but also of all the support and encouragement I received from a wonderful group of people I've had the pleasure of being associated with.

First and foremost, I want to thank my mom and dad for providing me with a loving household, a good education, for exposing me to the richness and diversity of various cultures through our travels, for teaching me to appreciate the good in the world, for saying no at times, and for giving me a wonderful person as a sibling. All these (and much more) helped me figure out who I am, to appreciate what I have, and to be kind. Thank you for all the sacrifices you made, for the valuable criticism of my actions, and for going beyond your comfort zone to accept some of my life-choices (including my choice to do a PhD). I want to thank my sister, Sree Gayathri Talluri, for being my best friend, and for her unconditional love. Thank you for being there for me through thick and thin, and for the occassional advice when I felt lost. I want to thank the rest of my family, including Priya's, and the Challagali's for their trust and love.

Tobias, thank you for being my supervisor and my mentor. Every interaction I had with you, whether it was about an ongoing/future project, or an interesting paper, or about life in general, made me wiser, and (continue to) profoundly shaped my thinking. I am grateful for your encouragement and support during times of self-doubt, for your constructive criticism of my work, for your patience, and intellectual breadth. Your curiosity, critical thinking, and enthusiasm for science are a perennial source of inspiration to me. You transformed this electrical engineer into a Cognitive Neuroscientist, for which I am forever grateful.

To all my colleagues in Donner lab, thank you for being a part of my PhD journey. You made me feel very welcome to the Neuroscience community, and have been extremely helpful with your feedback to improve my sceintific skills, writing, and presentation. The last 4.5 years have been some of the best times of my life so far, in no small part due to all of you: Peter, for being my best-friend, and for being the constant companion in all my outings in Hamburg - life in Hamburg would not have been so interesting without you; Niklas, for sharing your outlook of life, your quantitative expertise, and for notifying me about the opening in Donner lab during your visit to Goettingen- without you, I would not have ended up in the lab; Peter, and Niklas again, for keeping our office time interesting with your conversations, and for your patient responses to my unprompted questions during work hours; Konstantinos, for your scientific and non-scientific advice, and your acute modelling insights; Jan Willem, for sharing your positive outlook of life, for being a good friend, and for the sail-boat experience on the Alster; Anne, for being the best co-author and colleague one can wish for; Anke, for showing me the nightlife in Hamburg- I look forward to doing the same in Berlin; Rudy, for showing how to be a kid at heart, and for the entertaining fantasies about maintaining a miniature safari in the lab; Tom, for introducing me to ping-pong, and together with Rudy, showing the joy of doing science with your best-friend; Kobe, for being a good ping-pong partner, and for the great company in Berlin at FENS Forum; Thomas, for introducing me to Martino's pizza, and for the thoughtful conversations about applying science in policy-making; Olympia, for the help with coding my first fMRI experiment, and for the words of wisdom ("no amount of hyper-parameter tuning can pull out a non-existing effect"); Genis, for the entertaining board game nights; Keno, for your contagious cheerfulness, and for making the ice-skating trip to Applebek happen; Gina and Natalia, for motivating me to be a little more organised through practice; Gina again, for helping me translate the summary of this thesis into German; Carolina, for your wisdom on project management; Alena, for the great cocktails and special treatment at Red Lounge, and for teaching me how to make a Mojito; Christopher, for the ping-pong challenges to the whole lab; Maryam, for taking care of our plants while I was away.

I want to thank John-Dylan Haynes for giving me the opportunity to collect data in Berlin, and for his contagious smile that put me at ease whenever I was worried if the data collection would go as planned. I want to thank members of the Haynes lab for making me a part of their group during my stay in Berlin: Carsten, for your help with the fMRI project in general, and for many enlightening discussions on a range of topics; Shuyan, for your assistance with the data collection, and for the interesting conversations that helped me get through the long scanning hours; Andrea and Riccardo, for the exhilarating experience of watching the Euro cup 2016 final at Tempelhof airfield; Riccardo and Felix, for showing me how to operate the eye-tracker in the scanner; Doris, for organising the happy-hour hangouts on Friday evenings.

I want to thank Andreas Engel, and the scientific and non-scientific staff at the department of Neurophysiology for their support during my PhD. I want to thank Hilke Petersen for her help with handling the paperwork, and Sina Trautman-Lengsfeld for organising events to help PhD students in the department. Also, I want to thank my thesis committee, Prof. Christian Buechel, Prof. Claus Hilgetag, and my external thesis examiner, Prof. Philipp Sterzer for their critical insights, and evaluation of my work.

Outside of the university, I want to thank my friends in Hamburg, and beyond. Especially Marie, Pepe, and Anton, for offering me a place to stay during my first two months in Hamburg- you made it so easy to adjust to the new city, and I can never forget your hospitality; Siva, Anuradha, and Sashin, for your company- it is always a pleasure to be able to hold conversations in my mother tongue, and for the south Indian food, whenever I got home sick; Srushti Mahamuni, for organising the get-togethers, and parties- I am super excited about the new adventure you and Jan Willem are about to undertake as new parents; and Frau Strauss, for her kindness, and her cheerful nature, and for all the stories about the history of Goettingen- it still feels surreal that I lived in the same room that once housed the first female PhD student in Germany, Dorothea von Rodde-Schloezer, in 1787(!). Outside of Hamburg, I want to thank my close friends who were always there for me whenever I needed them, and for motivating me to take one more step when I felt down: Harish Alagani, Rajesh Noothi, Sparsh Singhai, and Sangameswara Karumanchi.

I wouldn't have taken up science if not for the encouragement and support of some excellent teachers: Vani Kompalli, and Uma Maheswari, for the foundations in maths, and sciences; Gopinath Pillai at IIT Roorkee for supervising my Bachelor's project; Robert Fisher, James Bednar, and Peggy Series for your support during my Masters at the University of Edinburgh; Stefan Treue, and Suresh Krishna for starting me with psychophysics at the Deutsches Primatenzentrum in Goettingen: I feel incredibly fortunate to have trained, and learned from you.

Finally, I want to thank Priyanka, the love of my life, for everything she did for me- I am looking forward to growing old with you, and for everything the future has to offer us.

CURRICULUM VITAE

Lebenslauf wurde aus datenschutzrechtlichen Gründen entfernt

EIDESSTATTLICHE ERKLÄRUNG

Ich versichere ausdrücklich, dass ich die Arbeit selbständig und ohne fremde Hilfe verfasst, andere als die von mir angegebenen Quellen und Hilfsmittel nicht benutzt und die aus den benutzten Werken wörtlich oder inhaltlich entnommenen Stellen einzeln nach Ausgabe (Auflage und Jahr des Erscheinens), Band und Seite des benutzten Werkes kenntlich gemacht habe.

Ferner versichere ich, dass ich die Dissertation bisher nicht einem Fachvertreter an einer anderen Hochschule zur Überprüfung vorgelegt oder mich anderweitig um Zulassung zur Promotion beworben habe.

Ich erkläre mich einverstanden, dass meine Dissertation vom Dekanat der Medizinischen Fakultät mit einer gängigen Software zur Erkennung von Plagiaten überprüft werden kann.

Hamburg, May, 2020

Bharath Chandra Talluri