Universitätsklinikum Hamburg-Eppendorf

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

Choice-Induced Biases in Perceptual Decision-Making

Dissertation
zur Erlangung des Doktorgrades PhD
an der Medizinischen Fakultät der Universität Hamburg

vorgelegt von
Anne Eugenia Urai
aus Utrecht

Hamburg, 2018

Angenommen von der Medizinischen Fakultät der Universität Hamburg am 16.01.2018

Veröffentlicht mit Genehmigung der Medizinischen Fakultät der Universität Hamburg.

Prüfungsausschuss, der Vorsitzende — *Prof. Dr. Tobias Donner*

Prüfungsausschuss, zweite Gutachterin — *Prof. Dr. Laura Busse*

Prüfungsausschuss, dritter Gutachter — *Prof. Dr. Christoph Mulert*

"Exactly!" said Deep Thought. "So once you do know what the question actually is, you'll know what the answer means."

— *Douglas Adams*, The Hitchhikers' Guide to the Galaxy

Contents

1 | Introduction

Our world is full of ambiguous sensory input, from which we have to extract relevant information to guide our behaviour. Imagine, for example, that you are walking through a foggy forest at dusk and catch a glimpse of a feline behind the trees. Not sure if the animal has already spotted you, you have to decide whether you have seen a large, ordinary house cat or a hungry tiger (Figure 1). The decision to approach it or to run as fast as you can is crucial for your survival, so you gather as much information as possible about the situation to arrive at the best decision. In this case, the relevant information is in the visual image as well as your knowledge about the world. If you are in your local park where you often see your neighbors' cat roam around, you have little reason to run. However, after reading a news article about a zoo break-out, you are much more likely to get away as fast as you can. This example shows that different types of information can *bias* decision processes, making our choices highly dependent on not just the current sensory input but also on our previous experiences.

Figure 1. When sensory information is uncertain, we use our prior knowledge to arrive at the most likely conclusion about the state of the outside world: in this case, is the animal a friendly cat or a dangerous tiger? © Craig Norton photography.

Appropriately responding to sensory evidence in the environment, and combining this information with our prior knowledge and expectations, is a crucial skill needed to survive and thrive in an uncertain world. How exactly do we use our previous experiences to guide decisions about incoming sensory information? And how do the networks of neurons in our brain work to efficiently transform sensory information and previous experiences into these decisions about the world around us?

Psychophysics and Signal Detection Theory

The question of how people and animals combine different types of information to guide their behaviour is central to psychology. Early psychological theory relied largely on introspection, and lacked methods to objectively quantify perceptual and decision processes (Read 2015). This changed in the 19th century with the birth of *psychophysics*, "an exact science of the relations between body and soul" (Fechner 1860), which set out to quantitatively relate properties of physical stimuli to subjective percepts and decisions.

Quantifying such perceptual decision processes requires rigorous experimental control over the input factors affecting them, combined with precise measurements of their outcome. With these inputs and outputs in hand, the psychophysicist can then formulate a model which specifies a law-like relationship describing the transformation happening within the observer. For example, the work by Weber was the first to quantify people's choice patterns as a function of the sensory input over many repeated trials (Weber 1846). This early work introduced the concept of the *psychometric function*, which relates repeated presentations of the same stimulus to the average response about the identity of this stimulus (Figure 2). This allows the experimenter to make inferences about underlying properties of the perceptual and decision-making machinery, such as the observer's sensitivity to the external stimulus, the shape of stimulus-response function, or the bias towards one over the other choice option.

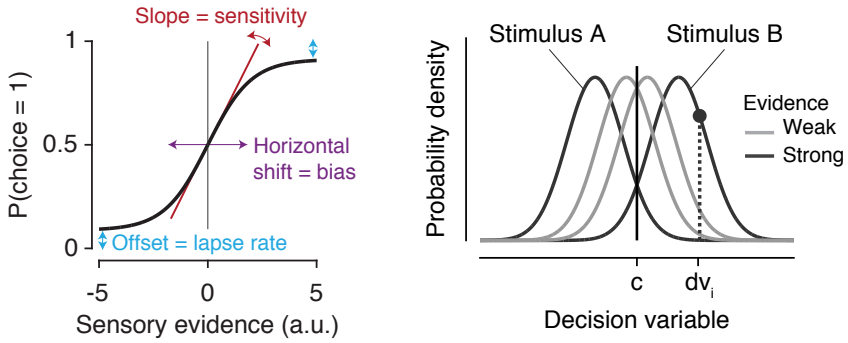

Figure 2. Left: The psychometric function quantifies separate aspects of choice behaviour. The slope of the function indicates the observer's perceptual sensitivity. The intercept indicates a horizontal shift of the psychometric function, reflecting a bias towards a specific choice independent of the sensory evidence. The vertical offsets from the two asymptotes indicate the fraction of stimulus-independent errors ('lapses'). Right: Signal detection theory. On each trial, the external stimulus is converted into an internal decision variable (DV). To make a binary choice, the observer compares the value of the DV to an internal choice criterion. Due to internal and external noise, a distribution of DVs over trials can be modeled as a Gaussian, the variance of which reflects the amount of noise present in the system. By increasing the available sensory evidence, DV distributions corresponding to the two choice options become more separable, resulting in better performance. By comparing the proportion of responses for the two external stimulus identities, SDT can deduce the observers' perceptual sensitivity (d') and her choice criterion (c). Reprinted from Urai et al. 2017.

Signal Detection Theory (SDT, Green and Swets 1966), was a major advance in formalizing the relationship between sensory input and choice behaviour. Originally developed in the field of radar detection, SDT specifies how to derive information about the presence of a signal which is embedded in noise. It provides a mathematical framework explaining several properties of the psychometric function, thereby separating

sensory information from the observer's decision rule. SDT is based on the idea that to make a single decision, the observer transforms the available sensory information into an internal *decision variable* (DV) which quantifies the decision-relevant evidence for one choice option over the other. For example, to answer the question 'did you see the dots on the screen move up- or downwards?', the observer takes the difference between sensory evidence for 'up' and 'down'. This DV is zero when there is no sensory information available to distinguish between the two stimulus classes, forcing the observer to guess. Importantly, the observer's internal DV is not a perfect representation of the true stimulus: it is corrupted by external and internal noise, which leads to an across-trial distribution of DVs for each external generative stimulus (Figure 2). When the two stimuli are easier to distinguish, for example by increasing the clarity of the image, these across-trial distributions become more separable. It is this separability that quantifies the observer's *perceptual sensitivity*, represented in SDT by the quantity d'.

Equipped with this decision variable, the observer then has to decide which of the two responses to give: has she seen upwards or downwards motion? Here, SDT postulates a *decision criterion*, which is the internal threshold to distinguish the two choice options. The placement of this criterion determines choice bias, or the tendency to choose one option over the other. The criterion can be determined by several factors: the relative cost or reward for one choice over another (the importance of this balance is apparent in medical decision-making, where misses often carry a much higher cost than false alarms) or the prior probability of one over the other option occurring. Importantly, SDT first formulated this explicit distinction between perception and the decision rule, by independently quantifying observers' sensitivity and criterion.

Signal Detection Theory offers several major benefits when analyzing data in perceptual choice experiments. First, it can estimate perceptual sensitivity in a bias-free way. Second, it specifies what the optimal criterion should be given a certain experimental context: for example, if an observer gets rewarded more for correctly responding 'up' then 'down', she should move her criterion such that she answers 'up' when she is not sure about the identity of the stimulus (Macmillan and Creelman 2004). Third, the psychometric function (Figure 2) can be interpreted as the cumulative distribution function of the underlying decision variable when multiple levels of external stimuli are presented. Assuming that Gaussian noise corrupts the decision process, the slope of the *probit* psychometric function is the inverse of the variance of the noise distribution (Klein 2001, see also **chapter 2.5**). The intercept of the psychometric function is related to the criterion placement, quantifying at which level of input the observer response to the two choice options with equal frequency. By fitting the psychometric function to data (Wichmann and Hill 2001b), the experimenter can thus make various deductions about the perceptual and decision processes taking place (Read 2015), making it one of the most used measurement models in experimental psychology. Lastly, within this simple framework another crucial variable in decision-making can be defined: the confidence in each binary choice. Such a model of confidence, and its application to the study of choice behaviour, is presented in **chapter 2.1**.

Sequential sampling models

One important aspect absent from SDT is the passage of time. Observers do not just make choices, but take a certain amount of time to report their judgments. This response time (RT) carries a wealth of information that can be used to disentangle the components of the decision-making process. The idea that reaction times can be used to infer the properties of mental processes dates back to the work of Donders (1868). He first reported that impulses take a measurable amount of time to travel through nerve tissue, and proposed that the time between stimulus and response can be used to make inferences about cognitive processes. For example, he measured how long it took observers to respond to the flash of a white light in cases when they knew a white light would be shown, or when they could see either a red or a white light and had to respond with one of both hands accordingly. The resulting difference of 150 milliseconds then reflected the time needed for the nervous system to arrive at a perceptual decision about the identity of the color (Donders 1868). This 'mental chronometry' was later formalized in a class of *sequential sampling models*, which represent the accumulation of evidence as a function of time, resulting in both a decision and its associated RT on each trial (Link and Heath 1975; Luce 1986).

The Drift Diffusion Model (DDM, Ratcliff 1978) is in many ways the standard sequential sampling model, in the sense that it is simple enough to be mathematically well characterized while at the same time able to capture a large variety of behavioural data (Bogacz et al. 2006). The DDM represents optimal behaviour in two-alternative forced choice (2AFC) decision tasks (Laming 1968), and in certain parameter regimes, a range of more complex sequential sampling models can be reduced to the DDM (Bogacz et al. 2006). Importantly, the DDM can be fit to psychophysical data (Ratcliff and Childers 2015), allowing for the recovery of generative parameters. This is not the case for all sequential sampling models, where more complete or biologically realistic models generally lose in parameter identifiability (Bogacz et al. 2006).

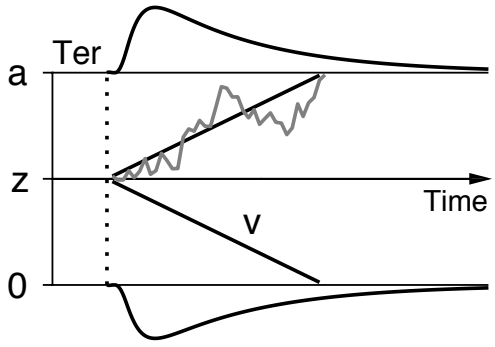

Figure 3. Schematic of the Drift Diffusion Model. In one example trial (grey), noisy sensory evidence is integrated over time into a decision variable, until reaching one of two choice bounds. Each reaction time is the sum of this decision process and a constant non-decision time Ter, resulting in reaction time distributions for each choice.

The DDM conceptualizes the process of evidence integration for 2AFC decisions as follows (Figure 3). Sensory evidence, corrupted by Gaussian noise, is integrated over time into a DV which represents the

relative evidence for one over the other choice. The average speed of accumulation across trials is quantified by the drift rate v. The evidence accumulation process terminates when the integrated DV reaches one of two bounds, corresponding to the two choice options. Reaction times are then the sum of the duration of this accumulation process and an additional non-decision time Ter that captures the initial encoding of sensory input as well as the time needed to execute the motor response. For unbiased decisions, the starting point of accumulation z is halfway between the two choice bounds 0 and a, and the drift rate is symmetric between the two choice outcomes. **Chapter 2.3** discusses the implementation of different types of bias in the DDM, and their fit to psychophysical data.

The idea of integrating evidence over time into a better decision is as powerful as generally applicable (Shadlen and Shohamy 2016), as underscored by its historical background. The sequential probability ratio test (SPRT), a formal definition of the optimal stopping rule when sequentially gathering information, was developed in the 1940's in the US (Wald 1947), and mainly employed for quality control in war-time logistics (Wallis 1980). Alan Turing independently developed the same statistical framework for dealing with the large amounts of data generated in the efforts to break the Enigma code (Good 1979). Specifically, the code-breakers accumulated information over different intercepted messages to infer the most likely settings of the Enigma rotor wheels, which changed every day (Gold and Shadlen 2002). These same sequential sampling models later formed a powerful foundation for the study of decision-making in the brain, allowing neural activity to be linked to specific algorithmic predictions of evidence accumulation models.

Neural basis of perceptual choice

To study how evidence is integrated in the brain and transformed into a choice, neuroscientists needs access to the decision-relevant information being accumulated. Many studies investigating decision-making in primates use visual tasks, of which the neural basis is relatively well understood. In influential early work on perceptual decision-making, monkeys performed a random dot motion task where they integrated the global motion direction of a cloud of moving dots, and reported the perceived motion direction by a saccade. Researchers then established that neural responses in visual area MT, known to be involved in motion processing, covary with the monkey's choice (Newsome et al. 1989; Britten et al. 1996). These observations and further computational modeling were used to deduce that the monkey used a small number of MT neurons to form his decision about the motion stimulus, and pointed to shared variability between neurons as the major source of noise determining the monkey's perceptual sensitivity (Shadlen et al. 1996).

When monkeys make a saccade to report their choice about the perceived motion direction of the cloud of dots, neurons in the lateral intraparietal area (LIP) were found to reflect integrated evidence towards one or the other choice (Roitman and Shadlen 2002). Specifically, neurons whose receptive field matched the location the monkey was preparing to saccade to showed a ramp-like increase in activity (but see Latimer et al. 2015), that was stronger for trials where the monkey responded faster. The firing rate in LIP neurons reached a typical fixed level at response execution, analogous to the pre-set decision bound in sequential

sampling models (Roitman and Shadlen 2002). This showed that neurons in motor-related areas reflected the integrated decision-variable posited by sequential sampling models (Gold and Shadlen 2007; Shadlen and Kiani 2013; Kira et al. 2015). Note that this does not, however, imply that LIP neurons are causally involved in the integration of decision-relevant evidence (Katz et al. 2016).

At the larger scale of neural mass signals, which can be recorded non-invasively in humans, similar electrophysiological read-outs of decision formation can be measured. Specifically, the lateralization of *beta-band* (12-36 Hz) oscillations over motor cortex tracks the gradual build-up of action preparation (Donner et al. 2009), in line with the idea that decision variables can be continually read out using motor signatures. A central-parietal positive event-related potential also tracks the temporal evolution of an ongoing decision variable, reaching a stereotypical threshold level at response execution (O'Connell et al. 2012; Twomey et al. 2015). Interestingly, this component predicts the timing but not the identity of the upcoming choice, suggesting it may also be a read-out of the observer's confidence (Urai and Pfeffer 2014). While motor-related build-up signatures are thus dependent on a known choice-response mapping (Bennur and Gold 2011; Filimon et al. 2013; de Lafuente et al. 2015; Twomey et al. 2016), they offer the unique advantage of allowing the neuroscientist to track not only when, but most importantly which decision is being prepared.

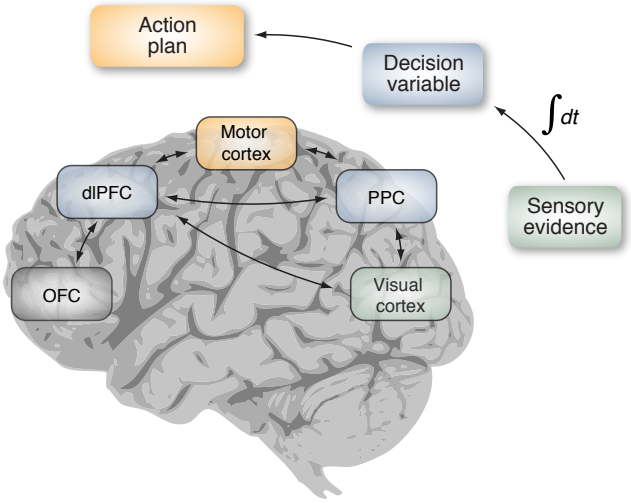

Figure 4. The neural basis of evidence accumulation in the human brain. Sensory evidence is sent from visual cortex to fronto-parietal decision-making circuits, which are tightly coupled to motor regions during decision formation. Black arrows denote information flow.

One key tool in the study of decision-making is the quantification of *choice probability*, the degree to which a neuron's activity covaries with the monkey's final choice over repeated presentations of the same stimulus. Significant choice probabilities have been reported for neurons throughout the visual cortex (Logothetis and Schall 1989; Britten et al. 1996; Nienborg and Cumming 2006). How does such correlation between a sensory neuron's firing rate and the animal's final choice arise? Stochastic fluctuations in firing rate, when shared between sensory neurons, can be integrated into the decision variable and thereby drive the animal's choice. In addition, feedback connections from decision circuits may affect the activity of

sensory neurons, causing them to be more correlated with the higher-level integrated decision variable.

To distinguish between these options, Nienborg and Cumming (2009) compared *psychophysical kernels*, which quantify at what point in time sensory evidence most affects an observer's final decision, with the timecourse of neural choice correlations. The researchers found that while choice probabilities increased over the course of a trial, sensory information presented later in the trial had a smaller effect on the monkey's final choice (Nienborg and Cumming 2009). These findings indicated that choice probabilities were at least partly caused by feedback from higher areas, rather than by sensory neurons causally affecting the decision. Later work using a hierarchical network model of spiking neurons corroborated that feedback connections from decision to sensory circuits could indeed explain the opposite timecourses of choice probability and psychophysical kernels (Wimmer et al. 2015; Urai and Murphy 2016). These findings are in line with the idea that commitment to a choice pushes the system in an attractor-like stable state, where new input has little effect on its state.

Intrinsic variability in brain and behaviour

The formal models of psychophysics and sequential sampling, combined with neurophysiological measures of sensory evidence and its integration, offer an exquisite tool to quantitatively study decision-making. They allow for investigating not just its constancy, but also its variability: humans and animals make different decisions each time they are presented with exactly the same information. A major goal of the neuroscience of decision-making is to tease apart the origins and consequences of this variation. Some of the randomness in behaviour may arise from irreducible unpredictability in neural functioning (Glimcher 2005). However, a substantial portion of behavioural variability is not noise, but rather a complex deterministic dependence on previous experiences, preferences and memories (Gold and Stocker 2017). What is the nature of this "central process which seems relatively independent of afferent stimuli" (Hebb 1949)? In other words, how exactly are our previous experiences, goals and preferences reflected in brain activity, and how do they affect choices?

A large body of work in psychology and neuroscience has investigated how decision-making is affected by factors beyond the immediate sensory input, ranging from attention (Gottlieb 2012), expectations (Summerfield and de Lange 2014) and reward (Hunt and Hayden 2017) to the trade-off between response speed and accuracy (Bogacz et al. 2010). By experimentally varying and tightly controlling such external factors, a picture has emerged of decision-making as a dynamic and flexible process, that can be adapted to the task context and observer's goals (Gold and Stocker 2017). However, those factors that we can experimentally manipulate tell us only half the story. To fully understand the rich interplay of causes that influence decision-making, we need a better grasp on *intrinsic* behavioural variability arising from sources we cannot control, but that we may be able to measure.

One important source of intrinsic variability is the history of *categorical choices*, which, rather than merely reflecting the endpoint of the decision-making process, bias ongoing brain activity and subsequent choice

patterns in precise ways. Even though experimenters randomize the order of presented stimuli, observers are consistently biased by their previous decisions. These intrinsic choice biases are ubiquitous in psychophysical datasets, and may account for some of the previously unexplained variability in neural activity and choice behaviour.

Serial dependencies in perceptual choice

Dependencies between subsequent perceptual decisions have been reported for almost a century (Fernberger 1920; Rabbitt 1968; Cross 1973), and have recently received renewed interest in neuroscience and psychology. Also called *sequential dependencies*, *decision inertia* or *choice hysteresis*, these *serial choice biases* have been observed in mice (Busse et al. 2011; Odoemene et al. 2017), rats (Licata et al. 2017), monkeys (Gold et al. 2008) and humans (Akaishi et al. 2014; Fründ et al. 2014; Kiyonaga et al. 2017). While the correlation between subsequent choices decreases with training (Gold et al. 2008), it remains present even in highly trained psychophysical observers who know that stimulus sequences are not auto-correlated (Fernberger 1920; Fründ et al. 2014).

Previous choices and stimuli have most effect on decisions when sensory evidence is unreliable (Fründ et al. 2014; Abrahamyan et al. 2016; Urai et al. 2017). In randomized laboratory experiments, the ideal strategy is to make decisions entirely based on decision-relevant information presented in the current trial: serial dependencies then result in suboptimal performance. Indeed, serial choice biases impair perceptual thresholds (Abrahamyan et al. 2016), although these effects may be small (Fründ et al. 2014). Previous choices also interact with external feedback about choice outcomes (Busse et al. 2011; Akaishi et al. 2014; Abrahamyan et al. 2016), and may relate to high-level cognitive variables and individual differences (Abrahamyan et al. 2016).

Why do sequential dependencies in choice behaviour arise? One intriguing suggestion is that observers learn that the world is generally stable over time, and bring this expectation with them to randomized laboratory experiments (Yu and Cohen 2008). They may learn the local transition probabilities between items in a sequence (Meyniel et al. 2016), and combine past observations using a process of leaky integration to guide their upcoming decisions (Sugrue et al. 2004; Glaze et al. 2015; Kim et al. 2017). This would lead their choices to be biased by streaks of repetitions and alternations that appear in randomly generated sequence of stimuli.

Serial dependencies could reflect processing in a range of brain regions, with candidates in early sensory, higher-level decision-making or motor-related brain areas. In a recent paper, St. John-Saaltink et al. (2016) used fMRI to show that in a 2AFC orientation discrimination task, previous choices bias the representation of orientation in early visual cortex. Task-irrelevant, salient stimuli that are presented between trials, which mask stimulus aftereffects in visual cortex, do not reduce serial dependencies (Akaishi et al. 2014). Together, these findings suggest that the effect of previous choices on activity in primary visual cortex may be due to feedback from downstream decision-making brain regions, with neural signatures that are perhaps similar

to those of top-down attention. Indeed, model-based fMRI shows that a network of frontal and parietal regions reflects trial-by-trial updates of expected choices, which track serial dependencies (Akaishi et al. 2014). At the motor end, the 'beta rebound', an increase in activity between 12-30 Hz that is observed after a hand movement and higher over contra- vs. ipsilateral motor cortex, carries over to the next trial and biases observers to alternate their choices (de Lange et al. 2013; Pape and Siegel 2016). These results suggest that serial dependencies may not be localized to one particular location in the brain, but might be present in various parts of the decision-making network.

One important potential contributor to biases in brain state and behaviour may be the release of modulatory neurotransmitters. These *neuromodulators* (such as noradrenaline, dopamine and acetylcholine) are phasically released during perceptual decision-making (Aston-Jones and Cohen 2005; de Gee et al. 2014), are diffusely released throughout the cortex, and can reflect remarkably specific cognitive and behavioural processes (Murphy et al. 2014; Kloosterman et al. 2015; Murphy et al. 2016a; de Gee et al. 2017). Decision-related phasic neuromodulation may thus be an interesting candidate marker for choice-induced biases in ongoing perceptual decision-making.

Outline of this thesis

A picture thus starts to emerge of serial choice patterns as a ubiquitous and likely fundamental phenomenon in decision-making. However, several open questions remain. First, how flexible are serial choice patterns? Can they be adapted to environments where stimulus sequences are structured? And do choices also bias perceptual inference within, rather than across trials? Second, at what stage of the decision-making process do serial choice biases arise? Do choice biases mainly arise from previous decisions, or their associated motor response? And do decisions affect later response preparation, or the perceptual inference process itself? At the neural level, do sensory, decision-making or motor signals explain serial choice biases? Third, how are serial choice patterns modulated by physiological and computational variables, such as arousal and confidence? This thesis investigates these questions from several angles, using psychophysics, computational modeling, pupillometry and magnetoencephalography.

Chapter 2.1 describes a statistical model of decision uncertainty, and shows data suggesting that both reaction times and pupil responses, a read-out of neuromodulatory brain state, represent such an uncertainty signal. This chapter then shows that these measures of decision uncertainty modulate serial choice biases, decreasing the probability of repeating previous choices.

In **chapter 2.2**, we show that serial choice biases can be adapted to environments where subsequent stimuli tend to repeat or alternative, in a way that again depends on decision uncertainty. Additionally, this chapter describes an experiment where choices and motor responses were dissociated, showing that choices but not motor responses drive serial choice patterns.

Moving from signal detection theory to sequential sampling models, **chapter 2.3** uses drift diffusion

modeling to tease apart biases in starting point and in the rate of evidence accumulation, finding that individual serial choice biases are mainly reflected in history-dependent bias in the rate of evidence integration.

Chapter 2.4 investigates the neural basis of serial choice patterns, using electrophysiological signatures of sensory evidence encoding and motor preparation to link individual serial biases to stages of the decision-making process. We show that effects of previous choices can be found in human electrophysiological signals both at the visual and motor end of the decision-making process.

Lastly, in **chapter 2.5** we investigate how binary choices halfway through a trial change how observers integrate subsequently presented evidence into a continuous estimation judgement. We show that observers are less sensitive to incoming sensory evidence that is inconsistent with a previous binary choice.

2.1 | Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias

Urai AE, Braun A, Donner TH. (2017) *Nature Communications*, 8:14637

Abstract

While judging their sensory environments, decision-makers seem to use the uncertainty about their choices to guide adjustments of their subsequent behaviour. One possible source of these behavioural adjustments is arousal: Decision uncertainty might drive the brain's arousal systems, which control global brain state and might thereby shape subsequent decision-making. Here, we measure pupil diameter, a proxy for central arousal state, in human observers performing a perceptual choice task of varying difficulty. Pupil dilation, after choice but before external feedback, reflects three hallmark signatures of decision uncertainty derived from a computational model. This increase in pupil-linked arousal boosts observers' tendency to alternate their choice on the subsequent trial. We conclude that decision uncertainty drives rapid changes in pupil-linked arousal state, which shape the serial correlation structure of ongoing choice behaviour.

 doi: 10.1038/ncomms14637. AUTHOR CONTRIBUTIONS: Conceptualization, A.E.U. and T.H.D.; Investigation, A.E.U.; Formal Analysis, A.E.U. and A.B.; Software, data curation and visualization, A.E.U.; Writing, A.E.U. and T.H.D.; Supervision, T.H.D.

Introduction

In perceptual and sensory-motor tasks, humans and animals behave as if they make use of decision uncertainty – the probability that a choice is correct, given the sensory evidence (Kepecs et al. 2008; Ma and Jazayeri 2014; Pouget et al. 2016). Theoretical accounts postulate that decision uncertainty should shape subsequent decision processing and, thereby, subsequent choice behaviour (Kepecs and Mainen 2012; Meyniel et al. 2015; Pouget et al. 2016). But how decision uncertainty is transformed into subsequent behavioural adjustments has, so far, remained elusive.

One prominent idea is that the brain broadcasts uncertainty signals across brain-wide neural circuits via low-level arousal systems (Dayan et al. 2000; Yu 2005; Meyniel et al. 2015). Arousal systems might be driven by uncertainty (Aston-Jones and Cohen 2005; Yu 2005; Dayan and Yu 2006; Nassar et al. 2012; Meyniel et al. 2015; de Berker et al. 2016), and they profoundly shape the global state of the brain through the action of modulatory neurotransmitters (Harris and Thiele 2011; Lee and Dan 2012; McGinley et al. 2015a). Uncertainty-dependent changes in global brain state, in turn, might translate into adjustments of choice behaviour. The goal of our study was to investigate whether arousal (1) reflects decision uncertainty in a perceptual choice task; and (2) predicts changes in subsequent choice behaviour.

Changes in central arousal state (as assessed by various measures of cortical dynamics) are tightly coupled to fluctuations in pupil diameter under constant luminance (Eldar et al. 2013; Reimer et al. 2014; McGinley et al. 2015a,b; Vinck et al. 2015). We here built on this connection and monitored pupil diameter as a proxy for central arousal state. We used a model based on statistical decision theory, illustrated in Figure 1, in which decision uncertainty is defined as the probability a choice is correct, given the available evidence (Pouget et al. 2016; Sanders et al. 2016). This operationalization of decision uncertainty obviates the need for subjective confidence reports (Kepecs and Mainen 2012), bridging to the insight from animal physiology that neurons in a number of brain regions encode decision uncertainty, as defined in Figure 1 (Kepecs et al. 2008; Komura et al. 2013; Lak et al. 2014; Teichert et al. 2014).

The model assumes that observers base their judgment of each stimulus on a noisy decision variable, sampled from a distribution that depends on the identity and strength of the stimulus (Figure 1a). Two-alternative forced choice tasks entail comparing this decision variable with a decision bound. When the decision variable happens to fall on the wrong side of the bound, errors occur. This happens more often for weaker stimuli, because the distributions corresponding to the two possible stimuli show higher overlap (Figure 1b). A monotonic function of the distance between the decision variable and the bound is a metric of decision confidence; uncertainty is its complement (Kepecs et al. 2008; Hebart et al. 2016; Sanders et al. 2016) (Figure 1a and Methods).

This model predicts three signatures of decision uncertainty (Kepecs et al. 2008; Sanders et al. 2016): (1) uncertainty decreases with evidence strength for correct choices (blue line in Figure 1c) but, counter-intuitively, increases with evidence strength for incorrect choices (red line in Figure 1c); (2) uncertainty predicts a monotonic decrease in choice accuracy from 100 to 50% (Figure 1d); (3) higher uncertainty predicts

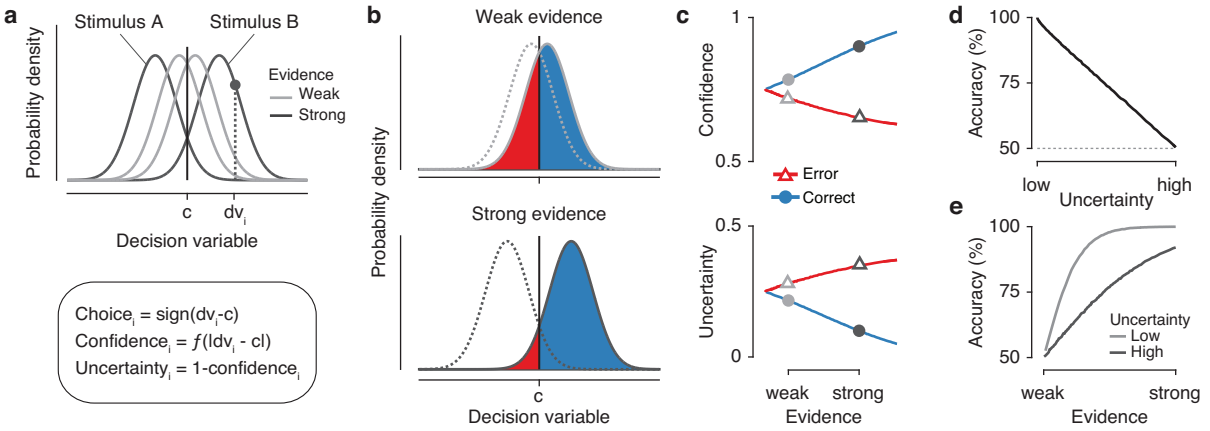

Figure 1. Operationalizing decision uncertainty. (**a**) Computations underlying choice and decision uncertainty. Due to noise, repeated presentations of a generative stimulus produce a normal distribution of internal responses centered at the mean of this generative stimulus. The internal response on each trial dv_i is a sample drawn from this distribution. It is compared to a decision bound or criterion c, to compute the binary choice as well as a graded measure of decision confidence (or its complement: uncertainty). (**b**) For two example levels of evidence strength, the average confidence is indicated by the shaded regions, separately for correct (blue) and error (red) trials. (**c**) Confidence (top) and uncertainty (bottom) as a function of evidence strength (100 bins), separately for correct and error trials. The two levels of evidence indicated by symbols (circles, triangles) correspond to the two example levels of evidence strength in panels **a**, **b**. (**d**) Accuracy as a function of decision uncertainty (100 bins). (**e**) Accuracy as a function of evidence strength (100 bins), separately for trials with high and low decision uncertainty (median split). For details, see Methods and (Kepecs et al. 2008; Sanders et al. 2016).

lower choice accuracy, even for the same evidence strength (Figure 1e). The opposite, monotonic scaling of uncertainty with evidence strength for correct and error trials (Figure 1c) also emerges from a variety of dynamic decision-making models, including race models (Kepecs et al. 2008), Bayesian attractor models (Bitzer et al. 2015), and biophysically detailed circuit models of cortical dynamics (Insabato et al. 2010; Wei and Stocker 2015).

We systematically manipulated the strength of sensory evidence and tested whether pupil responses exhibited the three signatures derived above. We then quantified the predictive effects of pupil-linked arousal on subsequent behaviour in terms of the key elements of the perceptual decision process: response time (RT), perceptual sensitivity, lapse rate, and choice bias. Choice bias was decomposed into an overall bias for one choice, and a serial bias dependent on the history of previous choices or stimuli. We found a predictive effect of pupil-linked arousal responses on serial choice bias.

Results

Pupil responses reflect decision uncertainty

27 human observers performed a 2-interval forced choice visual motion coherence discrimination task (Figure 2a and Methods). We applied motion energy filtering (Adelson and Bergen 1985) to the stochastic random dot motion stimuli, yielding a more fine-grained estimate of the decision-relevant sensory evidence contained in the stochastic stimuli than the nominal level of motion coherence (Figure 2b,c and Methods).

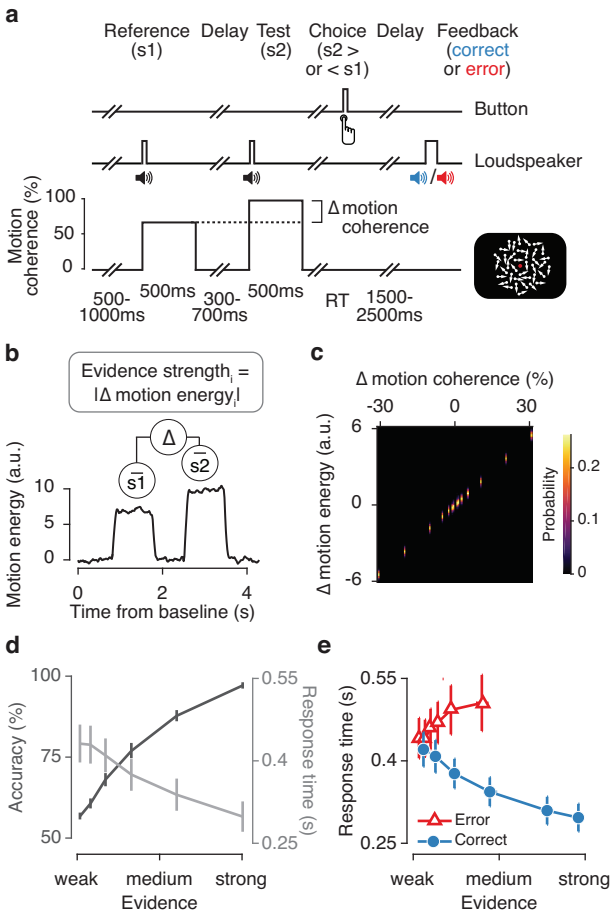

Figure 2. Perceptual choice task and behaviour. (**a**) Behavioural task. Dynamic random dot patterns were displayed throughout each trial. In two successive intervals (onset cued by beeps), the dots moved in one of the four diagonal directions (fixed per observer): A first 'reference' interval with always 70% motion coherence, and a second 'test' interval with varying levels of motion coherence, larger or smaller than the reference. Observers reported whether the test stimulus contained stronger or weaker motion than the reference by pressing one of two buttons. They received auditory feedback after a variable delay. (**b**) Quantifying evidence strength. Each random dot stimulus was convolved with a set of spatio-temporal filters (Adelson and Bergen 1985) to obtain a time course of motion energy. The difference between mean motion energy during test and reference intervals was used as a measure of single-trial measure evidence strength. (**c**) Probability distribution of evidence strength as a function of difference in nominal motion coherence. (**d**) Accuracy and median RT as a function of evidence strength (6 bins). (**e**) Median RT as function of evidence strength (6 bins), split by correct and error trials. (N=27, group mean ± s.e.m.)

The absolute value of this sensory evidence served as a single-trial measure of evidence strength (Figure 2b). As expected, stronger evidence yielded higher choice accuracy and faster responses (Figure 2d and Figure S2a).

In line with previous work (Sanders et al. 2016), RT exhibited all three signatures of decision uncertainty derived in Figure 1 above (Figure 2e and Figure S1b,c). This was true despite the interrogation protocol (Bogacz et al. 2006), in which the test stimulus had a fixed duration, its offset prompted the choice, and observers were instructed to maximize accuracy without speed pressure (response deadline was 3 seconds after test offset). Specifically, RT decreased with evidence strength on correct trials but increased with evidence strength on errors (Figure 2e). Further, RT predicted accuracy over a wide range, but not below

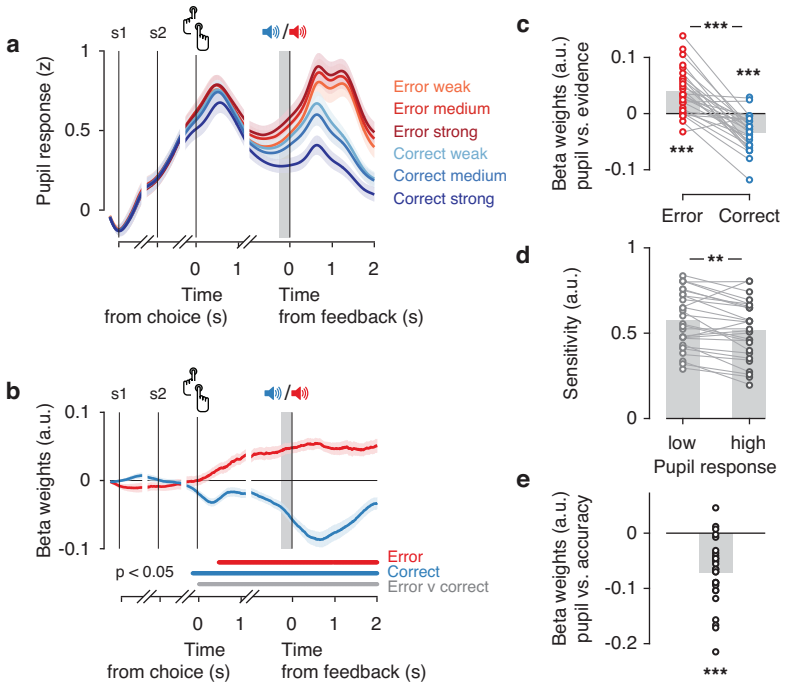

Figure 3. Pupil dilation after choice and before feedback reflects decision uncertainty. (**a**) Time course of pupil responses throughout the trial. Time courses were baseline-corrected and split by correct and error as well as three bins of evidence strength. Mean pupil dilation in the 250 ms before feedback (grey box) was used as a single-trial measure of pupil response. (**b**) Time course of uncertainty scaling in the pupil, computed as sample-by-sample regression of baseline-corrected pupil dilation onto evidence strength. Lower bars indicate $p < 0.05$ from a cluster-corrected permutation test, of the difference between each time course and zero, and between the two time courses. (**c**) Regression weights for the linear relationship between evidence strength and pupil responses. (**d**) Individual perceptual sensitivity, separately for lowest and highest pupil tertiles. (**e**) Individual logistic regression weights, using pupil responses to predict single-trial choice correctness. In **b**-**e** z-scored, log-transformed RTs were removed from the pupil signal via linear regression. *** $p < 0.001$, ** $p < 0.01$, permutation test. (N=27, group mean ± s.e.m.)

50% (Figure S1b), indicating that RT reflected decision uncertainty rather than error detection (Kepecs et al. 2008). We next assessed whether decision uncertainty also affected pupil-linked arousal.

The pupil dilated during decision formation, peaking just after the choice (button press) as observed in previous work (de Gee et al. 2014), and then dilated again after feedback (Figure 3a). Between these two peaks, dilation amplitudes diverged between different conditions, as predicted by decision uncertainty (compare with Figure 1c): Pupil responses were smallest after correct decisions based on strong evidence, they were overall larger after errors than correct choices, and largest after errors made on trials with strong evidence (Figure 3a).

To quantify the temporal evolution of uncertainty scaling in the pupil, we regressed baseline-corrected pupil time courses against each trial's evidence strength, separately for correct and error trials. From choice onwards, pupil dilation scaled positively with evidence strength on error trials, and negatively on correct trials (Figure 3b,c and Figure S3a). In other words, the scaling of the pupil response with evidence strength diagnostic of decision uncertainty emerged in the interval between choice and feedback. Consequently, this uncertainty scaling was not a response to the external information about choice correctness provided by the

external feedback, but rather reflected internal decision-related computations as described in Figure 1. For simplicity, we refer to the single-trial pupil dilation averaged across the 250 ms interval before feedback as 'pupil response' in the following.

The pupil response also exhibited the other two signatures of decision uncertainty predicted by the model in Figure 1. Larger pupil responses were accompanied by an overall lower choice accuracy (Figure 3e and Figure S3c), and psychophysical sensitivity was lower on trials with a larger pupil response (Figure 3d and Figure S3b). Specifically, the pupil response did not predict choice accuracy below 50%, suggesting that it did not signal the detection of errors (Figure S3c).

The scaling of the pupil response with decision uncertainty was not inherited from the analogous scaling of RT, but was also present after first removing (via linear regression) the trial-to-trial variations accounted for by RT (Figure S3d-f). Indeed, trial-to-trial correlations between pupil responses and RTs were generally small (Pearson correlation, average r: 0.087 range: -0.042 to 0.302, for log-transformed RT). For all subsequent analyses reported in this paper, we removed RT-fluctuations from the trial-to-trial fluctuations of single-trial pupil responses via linear regression (see Methods).

Pupil-linked arousal alters subsequent choice behaviour

We proceeded to test whether uncertainty-related pupil responses predicted changes in subsequent choice behaviour. It has been proposed that arousal signals control various aspects of learning and decision-making (Dayan et al. 2000; Aston-Jones and Cohen 2005; Yu 2005; Dayan and Yu 2006; Meyniel et al. 2015). In the context of our task, the choice parameters of interest were perceptual sensitivity (measured as the slope of the psychometric function, Figure S4a), lapse rate (measured as the vertical distance of the asymptotes of the psychometric function from 0 or 1, Figure S4a), bias (measured as the horizontal shift of the psychometric function, Figure S4a), and RT. For RT, we focussed on increases after error trials, an effect referred to as post-error slowing (Dutilh et al. 2012), which was found to be modulated by pupil-linked arousal in a speeded RT task (Murphy et al. 2016b). Choice bias was further decomposed into two parameters: overall bias (i.e., a general tendency towards one choice option, averaged across the entire experiment, Figure S4b) and serial bias (i.e., a local, choice history-dependent tendency towards one option that becomes evident when conditioning the psychometric function on the preceding choice, Figure S4c) (Fernberger 1920; Yu and Cohen 2008; Fründ et al. 2014). Because in our task (as common in laboratory choice tasks), the sensory evidence was independent across trials, any serial bias was maladaptive, reducing observers' performance below the optimum they could achieve given their perceptual sensitivity.

The pupil response predicted a reduction of serial choice bias (Figure 4a and Figure S5). When a choice was followed by a small pupil response, observers tended to repeat this choice on the next trial; when the previous pupil response was large, this serial bias was abolished (Figure 4a). This predictive effect was similar for correct and error trials (Figure S6a). An analogous pattern of predictive effects was observed when binning by previous trial RT: Fast, but not slow, RTs were followed by a tendency to repeat the previous choice

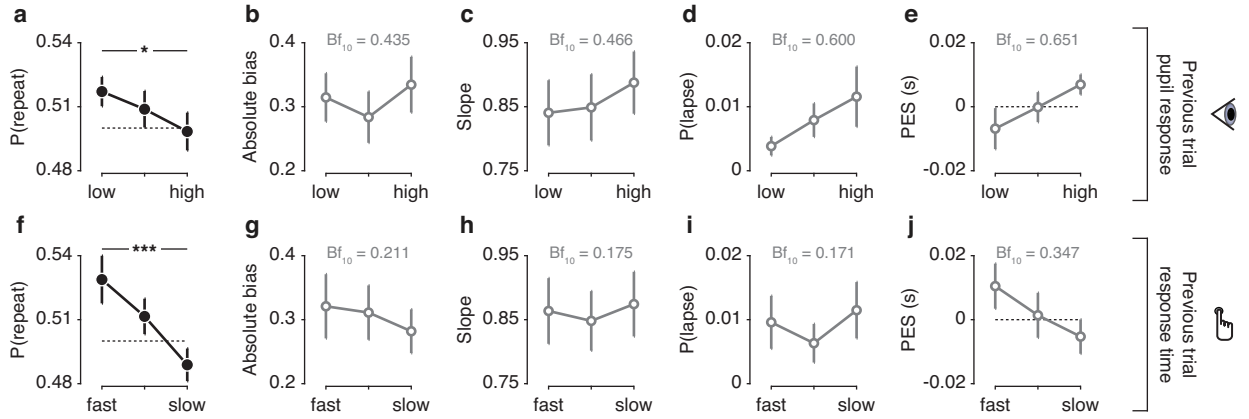

Figure 4. Pupil responses and RT predict reductions in serial choice bias. (**a**) Serial choice bias, quantified as the history-dependent shift of the psychometric function, for tertiles of previous trial pupil responses. (**b**) Absolute choice bias, measured as the intercept of a logistic psychometric function, for tertiles of previous trial pupil responses. (**c**) Perceptual sensitivity, measured as the slope of a logistic psychometric function, for tertiles of previous trial pupil responses. (**d**) Lapse rate, measured as the probability of stimulus-independent guesses, for tertiles of previous trial pupil responses. (**e**) Post-error slowing, for tertiles of previous trial pupil responses. (**f**-**j**) as in **a**-**e**, but for tertiles of previous trial RT. *** $p < 0.001$, * $p < 0.05$, main effect of pupil/RT bin on repetition probability computed from a one-way repeated measures ANOVA. Unfilled markers indicate $p > 0.05$, with Bf_{10} from a Bayesian repeated measures ANOVA written in panel. (N=27, group mean ± s.e.m.)

(Figure 4f and Figure S6b).

The pupil response predicted none of the other choice parameters on the next trial (assessed by one-way repeated measures ANOVA), neither overall choice bias (signed overall bias: $F_{(2,52)} = 0.939$, $p = 0.398$, $Bf_{10} = 0.221$; absolute value of overall bias: $F_{(2,52)} = 1.817$, $p = 0.173$, Figure 4b), nor perceptual sensitivity ($F_{(2,52)} = 1.936$, $p = 0.155$, Figure 4c), nor lapse rate ($F_{(2,52)} = 2.213$, $p = 0.120$, Figure 4d), nor RT (overall RT: $F_{(2,52)} = 3.232$, $p = 0.048$, $Bf_{10} = 1.207$; post-error slowing: $F_{(2,52)} = 2.056$, $p = 0.138$, Figure 4e). Variations in RT, likewise, did not predict a change in any of the other parameters of the decision process (Figure 4g-j, all $p > 0.05$). The overall pattern of results implies that observers did not simply act more randomly after large pupil responses or RT. Random button presses would have reduced sensitivity, in other words, decreased the slope of the psychometric function, contrary to our observations (Figure 4c,h). Rather, the pattern of results implies that, after large pupil responses or RT, observers' tendency towards one or the other choice became less history-dependent.

In sum, large pupil responses and slow RTs were neither followed by improved processing of sensory evidence (a common effect of attention, Ress et al. 2000), nor a change in overall response bias. Large pupil responses and slow RTs were followed by only minor (and statistically not significant) changes in stimulus-independent lapses as well as small adjustments in speed accuracy trade-off, as observed after response conflict, errors, or large pupil responses in speeded RT tasks (Botvinick et al. 2001; Gao et al. 2009; Murphy et al. 2016b). The weak effect on post-error slowing might be due to the use of an interrogation protocol in our study, which did not require observers to optimize their speed-accuracy trade-off (Bogacz et al. 2006). However, both RT and pupil-linked arousal had a robust effect on serial choice bias, reducing an overall repetition bias that predominated across the group of observers. This effect of both uncertainty-related

measures on the serial correlation structure of choice behaviour has so far been unknown. We therefore proceeded to model and comprehensively quantify this effect at the level of individual observers.

Pupil-linked arousal predicts choice alternation

To this end, we extended a previously established regression model of serial choice biases (Fründ et al. 2014) with pupil- and RT-dependent modulatory effects. The basic model (i.e., without modulatory terms) quantified the impact of the previous seven choices and stimuli on the current choice bias in terms of linear combination weights (Figure 5a, see Methods and (Fründ et al. 2014). We added to this model multiplicative interaction terms, that quantified how much the effect of previous stimuli and choices was modulated by either pupil response or RT on those same trials (Figure 5a). Simultaneously modeling the effects of both pupil responses and RT enabled us to estimate their independent impact on serial choice bias; we found the same results when fitting a separate regression model for each modulatory variable (Figure S7).

The model fits revealed robust, and idiosyncratic, patterns of serial choice biases in most observers (Figure 5c,d; see Figure S2b,c for individual sessions). As expected, the contribution of past stimuli and choices to current behaviour was strongest when sensory evidence was weak and decayed strongly with evidence strength (Figure 5b). The weight of the immediately preceding choice was generally stronger than the weight of the previous stimulus (Figure 5d). The effect of previous choices lasted up to seven trials in the past (corresponding to about 60 s, Figure 5c), but had the largest absolute magnitude on the preceding trial (Figure 5c, grey dashed line). There was large inter-individual variability in choice weights (Figure 5c,d). While the majority of observers systematically repeated their choices (purple symbols; 12 significant at $p < 0.05$), a good fraction tended to alternate their choices (orange symbols; 7 significant at $p < 0.05$).

Observers' serial choices biases were unrelated to the (small) serial correlations between stimuli. The transition probabilities between stimulus categories (i.e. s2<s1 or s2>s1) were close to 0.5 (range across observers: 0.475 to 0.508), and did not correlate with individual choice weights (Pearson correlation $r = 0.010$, $p = 0.960$, $Bf_{10} = 0.149$) or stimulus weights (Pearson correlation $r = -0.176$, $p = 0.381$, $Bf_{10} = 0.217$). Likewise, the auto-correlation of absolute motion coherence differences (i.e., absolute levels of evidence strength) was close to 0 (range across observers: -0.061 to 0.028) and did not correlate with individual choice weights (Pearson correlation $r = 0.123$, $p = 0.541$, $Bf_{10} = 0.179$) or stimulus weights (Pearson correlation $r = -0.142$, $p = 0.480$, $Bf_{10} = 0.190$).

Critically, pupil responses and RT both negatively interacted with the effect of previous choices (Figure 5e), in line with the observation that large pupil responses or long RTs were followed by less choice repetition (Figure 4a,f). By contrast, neither pupil responses nor RT interacted with the effect of the previous stimulus (Figure 5e). Pupil responses beyond one trial in the past, as well as baseline pupil diameter on the current trial, did not predict a modulation of serial biases (Figure S8). Moreover, these results were not accounted for by trial-to-trial variations in trial timing or the passage of time between trials (Figure S9).

The pupil response after feedback did not contain information predictive of serial choice bias, beyond the

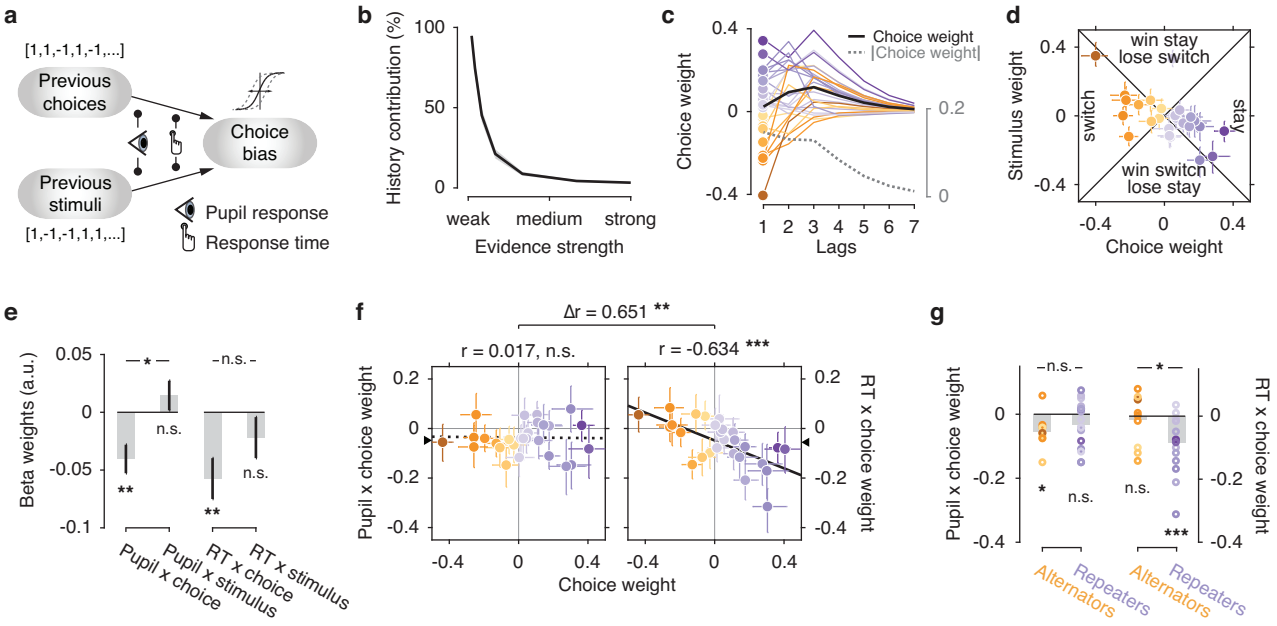

Figure 5. Modeling the modulation of serial choice bias. (**a**) Schematic of the regression model with modulatory terms. (**b**) The contribution of history terms (past choices and stimuli) as a fraction of the total variance in the decision variable (Fründ et al. 2014), decreased with stronger sensory evidence. (**c**) Choice weights for the previous 7 trials, obtained from the history model without modulatory terms. Each line corresponds to one observer. Purple, 'repeaters' with positive choice weight for lag 1. Orange, 'alternators' with negative choice weights for lag 1. Black line, group mean. Grey dashed line, group mean of absolute choice weight. (**d**) Choice weights at lag 1 plotted against the corresponding stimulus weights. Colored dots and error bars indicate individual observers ± 68% confidence intervals obtained from a bootstrap. See Methods for an interpretation of this graph in terms of behavioural strategy. (**e**) Regression weights for the interaction between previous pupil response or RT and previous choices or stimuli. N=27, group mean ± s.e.m. (**f**) Correlation between choice weights and their modulation by pupil dilation or RT. Colors indicate the choice weight as derived from the basic model in **d**. Error bars are 68% confidence intervals obtained from a bootstrap. The intercept of the least-squares regression line, corresponding to the mean beta weight across the group, is indicated with a triangle on the y-axis. (**g**) Beta weights for interaction between previous pupil response or RT and previous choices. Group split based on the sign of individual choice weights. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, n.s. $p > 0.05$, Pearson's correlation coefficient or permutation test.

information already present during the pre-feedback interval. The post-feedback pupil responses similarly predicted modulation of serial choice biases, but no longer did so when removing (via linear regression) variance explained by pre-feedback pupil responses from the post-feedback pupil signal (Figure S10).

While the modulatory effects associated with pupil responses and RT were both negative on average, such an overall reduction of the group-level repetition bias (Figure 4a,f) might be due to two alternative scenarios at the level of individual observers: either a reduction of each observer's intrinsic serial choice bias for repetition or alternation (referred to as 'bias reduction' hereafter); or, alternatively, a general boost of choice alternation, regardless of the observer's intrinsic serial bias (referred to as 'alternation boost'). We quantified intrinsic serial bias as each observer's choice weight (i.e., the main effect of the previous on the current choice estimated by our model). The bias reduction scenario predicts a negative correlation between choice weights and modulation weights across observers. The alternation boost scenario predicts negative individual modulation weights for all observers, independently of their corresponding choice weights (i.e., no correlation).

The analysis of these individual behavioural patterns revealed dissociable effects of pupil-linked arousal and RT (Figure 5f,g). Modulation weights for the pupil were negative for most observers, irrespective of their individual choice weight. When splitting all 27 observers into 'alternators' and 'repeaters' based on the sign of their intrinsic bias (i.e., choice weight), we found no correlation between individual modulation and choice weights (Figure 5f, Pearson correlation $r = -0.017$, $p = 0.935$, $Bf_{10} = 0.149$). Further, the modulation weights were negative for both subgroups, and not significantly different between them (Figure 5g). These observations are consistent with the idea that pupil-linked arousal generally boosted observers' tendency to alternate their choice on the next trial.

The alternation boost scenario for pupil responses was further supported by a striking contrast to RT-linked modulations, which were in line with the bias reduction scenario. The RT-linked modulation weights exhibited a strong negative correlation with individual choice weights (Figure 5f, Pearson correlation $r = -0.634$, $p < 0.001$, $Bf_{10} = 76.359$), were negative only for the group of repeaters, and differed significantly between alternators and repeaters (Figure 5g). Correspondingly, the correlations with individual choice weights were significantly different for pupil- and RT-modulation weights (Figure 5f). Moreover, RT-dependent bias reduction was most pronounced after incorrect choices, whereas the pupil-dependent alternation boost was most pronounced after correct choices (Figure S11).

In sum, the modulatory effects associated with post-decision pupil-linked arousal and RT both shaped the serial correlation structure of choices, but in distinct ways: pupil-linked arousal generally promoted choice alternation, regardless of the observer's intrinsic bias, whereas RT-linked processes generally reduced observers' intrinsic bias.

Discussion

Decisions about an observer's sensory environment do not only depend on the momentary sensory input, but also on the behavioural context (Gold and Shadlen 2007). One such contextual factor is the history of preceding choices and stimuli, which robustly biases even highly trained decision-makers (Fründ et al. 2014). Although such serial choice biases were first identified in psychophysical tasks about a century ago (Fernberger 1920), their determinants have remained poorly understood. Previous treatments of serial choice biases have conceptualized experimental history as sequences of binary external events (past stimulus identities, choices, or feedback) (Fründ et al. 2014; Abrahamyan et al. 2016). We here established that these serial biases were also modulated by the decision-maker's pupil-linked arousal state on the previous trial, which, in turn, reflected the uncertainty about the observer's choice.

Several important features of our approach allowed us to move beyond previous work linking human pupil dynamics to uncertainty and performance monitoring. First, different from most previous studies, we here unravelled the temporal evolution of uncertainty information in the pupil response, enabling inferences about not only the existence, but also the time course of this information (see O'Reilly 2013 for a similar approach). Second, the model-based definition of decision uncertainty we used helped dissociate decision

uncertainty from error detection, which has previously been linked to pupil dilation (Wessel et al. 2011). In a two-choice task, a signal encoding decision uncertainty should predict behavioural performance over a range from 100% to 50% correct (corresponding to 50% for the maximum uncertainty signal, or larger when encoding is imprecise). By contrast, an error detection signal should predict performance over the range 100% to 0% correct (Kepecs et al. 2008). Our measurements were more consistent with decision uncertainty than error detection (Figure S3c). Third, in our task, decision uncertainty critically depended on internal noise (the primary source of the variance in Figure 1a). By contrast, previous studies linking uncertainty to pupil dynamics (Preuschoff et al. 2011; Nassar et al. 2012; O'Reilly 2013; de Berker et al. 2016) have used tasks in which the primary source of uncertainty was in the observers' environment. Last, in contrast to most previous pupillometry studies (Preuschoff et al. 2011; de Gee et al. 2014; Lempert et al. 2015) we comprehensively quantified the predictive effects of pupil-linked arousal on the parameters of choice beyond the current trial, thereby complementing recent work on the effects of pupil-linked arousal on learning (Nassar et al. 2012; O'Reilly 2013). Taken together, our results critically advance the understanding of how internal decision uncertainty is encoded in pupil-linked arousal in humans, in a way that builds a direct bridge to single-unit recording studies of decision uncertainty in animals (Kepecs et al. 2008; Komura et al. 2013; Lak et al. 2014; Teichert et al. 2014).

The neural sources of task-evoked pupil responses at constant luminance are not yet fully identified (McDougal and Gamlin 2008), but mounting evidence points to the noradrenergic locus coeruleus (LC) (Murphy et al. 2014; Varazzani et al. 2015; Joshi et al. 2016) (a core component of the brain's arousal system (Aston-Jones and Cohen 2005)) as well as the superior and inferior colliculi (Wang and Munoz 2015). Microstimulation of all three structures triggers pupil dilation (Joshi et al. 2016). Among these structures, activity of the LC (spontaneous or evoked by electrical stimulation) is followed by pupil dilation at the shortest latency (Joshi et al. 2016). The LC also has widespread, modulatory projections to the cortex implicated in regulating central arousal (Aston-Jones and Cohen 2005). Dopaminergic and cholinergic systems, which are closely connected with the LC (Sara 2009), are likewise implicated in central arousal state (McGinley et al. 2015a) and may also contribute to task-evoked pupil responses.

We propose that decision-makers' uncertainty about their choices might shape serial choice biases by recruiting pupil-linked neuromodulatory systems. Frontal brain regions encoding decision uncertainty send descending projections to several of these systems (Aston-Jones and Cohen 2005; Sara 2009), which in turn project to large parts of the cortex, including networks of regions involved in perceptual inference and decision-making (Siegel et al. 2011). Neuromodulators like noradrenaline can profoundly alter the dynamics and topology of cortical networks (Marder 2012; Eldar et al. 2013; Polack et al. 2013; McGinley et al. 2015a). Thus, these brainstem arousal systems might be in an ideal position to transform variations in decision uncertainty into adjustments of choice behaviour (Yu 2005; Meyniel et al. 2015).

The behavioural effect of pupil-linked arousal might be explained by at least two (not mutually exclusive) scenarios. First, arousal responses might promote choice alternation at the level of response preparation, by altering the state of the motor system (de Lange et al. 2013). Second, the arousal response might modulate

the decision stage – specifically the dynamic updating of beliefs about the upcoming evidence, for example by shifting the criterion (assumed to be constant in signal detection theory, Figure 1) from one choice to the next. When this criterion is shifted in the direction opposite to the last choice, alternation ensues. In line with these ideas, changes in pupil-linked arousal state can indeed translate into specific behavioural effects (Eldar et al. 2013; de Gee et al. 2014), presumably by interacting with selective cortical circuitry (Donner and Nieuwenhuis 2013).

Our current observations are not easily reconciled with existing theoretical accounts of the impact of phasic arousal on decision-making. One account posits that threshold crossing of the decision variable triggers phasic noradrenaline release, facilitating the translation of the decision into a behavioural response (Aston-Jones and Cohen 2005). In contrast to our observations, this framework focuses on functional effects of phasic arousal within the same trial, rather than subsequent ones, and it predicts improvements in sensitivity and/or RT (Cavanagh et al. 2014), rather than changes in bias. Other accounts have proposed that phasic noradrenaline release facilitates a 'network reset' (Bouret and Sara 2005), enabling the transition of neural decision circuits to a new state (Dayan and Yu 2006). Our group-level finding that high pupil-linked arousal reduces serial biases might be interpreted as the discarding of post-decisional activity traces due to network reset (Karlsson et al. 2012; Tervo et al. 2014). However, our analysis of individual choice patterns revealed that pupil-linked arousal boosted alternation also in those observers who already exhibited a tendency to alternate their choices, which is not easily reconciled with the network reset idea.

Previous theories of arousal and neuromodulation have coarsely distinguished between two timescales of arousal fluctuations: Tonic fluctuations over the course of seconds to hours, and phasic responses on a sub-second timescale, time-locked to rapid cognitive acts (Aston-Jones and Cohen 2005; Yu 2005; Dayan and Yu 2006). Changes in tonic arousal occur spontaneously (Steriade 2000; McGinley et al. 2015a), and might also track changes in task utility or uncertainty (Aston-Jones and Cohen 2005; Yu 2005; Nassar et al. 2012; de Berker et al. 2016). Pupil-linked changes in tonic arousal strongly shape the operating mode of cortical circuits, including early sensory cortices, on slow timescales (McGinley et al. 2015a). Phasic pupil-linked arousal responses, on the other hand, predict behaviour related to the same transient cognitive processes that drive them (Einhäuser et al. 2008; Preuschoff et al. 2011; de Gee et al. 2014). The uncertainty-linked pupil responses we identified here built up slowly after choice and predicted choice behaviour several seconds later. Thus, our current results suggest that pupil-linked arousal systems are driven by, and interact with, cognitive processes also at intermediate timescales; faster than tonic arousal, but more sustained than task-evoked phasic responses.

The dissociation between pupil- and RT-linked modulatory effects (Figure 5f and Figure S11) on serial choice bias suggests that decision uncertainty signals were propagated along distinct central neural pathways, one linked to pupil responses and the other to RT, which then shaped serial choice biases in different ways. Even if the same uncertainty signals fed into these pathways, they might have become decoupled through independent internal noise. Specifically, it is tempting to speculate that the pupil-linked alternation boost reflected neuromodulator release from brainstem centers (such as noradrenaline from the LC, Tervo

et al. 2014), whereas RT-linked bias reduction was driven by frontal cortical areas involved in explicit performance monitoring and top-down control (such as anterior cingulate cortex, Botvinick et al. 2001; Yeung et al. 2004; Ebitz and Platt 2015). Top-down effects of prefrontal cortex on decision-making (Botvinick et al. 2001; Miller and Cohen 2001) are commonly associated with explicit strategic effects that are adaptive within the experimental task. Indeed, the RT-linked modulation of serial bias was adaptive, in that it generally reduced observers' intrinsic serial bias. By contrast, pupil-linked arousal modulated serial choice patterns in a way that was maladaptive for part of the observers (the alternators). This finding might be related to the observation that maladaptive serial choice biases remain prevalent even in highly trained observers who know the statistics of the task (Fernberger 1920; Fründ et al. 2014). Taken together, the dissociation between pupil- and RT-linked effects suggest that serial choice biases result from a complex interplay between low-level, pupil-linked arousal systems and higher-level systems for strategic control. Future studies should pinpoint the neural systems underlying these distinct effects, as well as their interactions (Tervo et al. 2014).

In conclusion, our study identified decision uncertainty as a high-level driver of phasic arousal, and it uncovered a role of this pupil-linked arousal response in shaping the dynamics of serial choice biases – a pervasive but often ignored characteristic of human decision-making. These insights shed new light on the link between decision uncertainty, pupil-linked arousal state, and serial dependencies in decision-making. They set the stage for further investigations into the neural bases of arousal-dependent modulations of serial choice behaviour.

Methods

Operationalizing decision uncertainty

In signal detection theory, a decision variable dv_i is drawn on each trial from a normal distribution $\mathcal{N}(\mu, \sigma)$ with μ corresponding to that trial's sensory evidence and σ reflecting the internal noise. In Figure 1, we used the range of single-trial motion energy values $[-6, 6]$ as our μ. We estimated σ from the data using a probit psychometric function fit on data combined across observers. The probit slope $\beta = 0.367$, where its inverse $\sigma = 2.723$ reflected the standard deviation of the dv distribution. The decision bound c was set to 0, reflecting an observer without overall choice bias. The two pairs of distributions in Figure 1 were generated using $\mu = -1$ and $\mu = 1$ for weak evidence, and $\mu = -4$ and $\mu = 4$ for strong evidence. To calculate the relationship between evidence strength and decision uncertainty (Figure 1c), we simulated a normal distribution of dv for each level of evidence strength, with $\mu = [0, 6]$ and $\sigma = 2.723$. Since these uncertainty computations are symmetrical with respect to choice identity, we visualized only the pattern corresponding to $\mu > 0$ (stimulus B in Figure 1a). All samples from such a distribution were split into correct and error parts based on their position with respect to the decision bound c. For each combination of evidence strength and choice, the average uncertainty level is w

$$uncertainty = 1 - \frac{1}{n}\sum_{i=1}^{n} f(|dv - c|) \quad (1)$$

where f is the cumulative distribution function of the normal distribution

$$f(x) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{x}{\sigma\sqrt{2}}\right)\right] \quad (2)$$

which transforms the distance between dv and c into the probability of a correct response (Lak et al. 2014).

We simulated ten million trials based on the range of evidence in the data, and for each we computed a binary choice, the corresponding level of decision uncertainty, and the accuracy of the choice. Figure 1c-e visualises the relationship between evidence strength, uncertainty and choice accuracy in these simulated data.

Participants and sample size

Twenty-seven healthy human observers (10 male, aged 23 $\pm$ 5.2 years) participated in the study. The ethics committee at the University of Amsterdam approved the study, and all observers gave their informed consent. We included all observers in each analyses presented in the paper. Each observer participated in one practice session and five main experimental sessions, each of approximately two hours and comprising 500 trials of the task. The number of observers was selected to allow for robust analyses of individual differences, as in previous pupillometry work from our lab (de Gee et al. 2014), and the total number of trials per observer was chosen to allow for robust psychometric function fits and detection of subtle changes in the fit parameters.

Task and procedure

Observers performed a 2-interval forced choice motion coherence discrimination task at constant luminance (Figure 2a). Observers judged the difference in motion coherence between two successively presented random dot kinematograms (RDKs): a constant reference stimulus (70% motion coherence) and a test stimulus (varying motion coherence levels specified below). The intervals before, in between, and after (until the inter-trial interval) these two task-relevant stimuli had variable duration (numbers in Figure 2a) and contained incoherent motion. A beep (50 ms, 440 Hz) indicated the onset of each (test and reference) stimulus. After offset of the test stimulus, observers had 3 seconds to report their judgment (button press with left or right index finger, counterbalanced across observers). After a variable interval (1.5-2.5 s), a feedback tone was played (150 ms, 880 Hz or 200 Hz, feedback-tone mapping counterbalanced across observers). Dot motion was stopped 2-2.5 s after feedback, with stationary dots indicating the inter-trial interval, during which observers were allowed to blink their eyes. Observers self-initiated the next trial by button press (range of median inter-trial intervals across observers: 0.68 to 2.05 s).

The difference between motion coherence of test and reference was taken from three sets: easy (2.5, 5, 10, 20, 30), medium (1.25, 2.5, 5, 10, 30) and hard (0.625, 1.25, 2.5, 5, 20). All observers started with the easy set. We switched to the medium set when their psychophysical thresholds (70% accuracy defined by a

cumulative Weibull fit) dropped below 15%, and to the hard set when thresholds dropped below 10%, in a given session.

Motion coherence differences were randomly shuffled within each block. We applied a counterbalancing scheme ensuring that within a block, each stimulus category (s2 > or < s1) was followed by itself or its opposite equally often (Brooks 2012). The algorithm generated sequences of 53 trials, of which the first 50 were used per block.

Random dot kinematograms

Stimuli were generated using PsychToolbox-3 (Kleiner et al. 2007) and presented on a 22" CRT monitor with a resolution of 1024 x 768 pixels and a refresh rate of 60 Hz. A red 'bulls-eye' fixation target (Thaler et al. 2013) of 1.5° diameter was present in the centre of the screen. Dynamic random noise was presented in a central annulus (outer radius 12°, inner radius 2°) around fixation. The annulus was defined by a field of dots with a density of 1.7 dots/degrees^2, resulting in 768 dots on the screen in any given frame. Dots were 0.2° in diameter and had 100% contrast from the black screen background. All dots were divided into 'signal dots' and 'noise dots', whose proportions defined the varying motion coherence levels. Signal dots were randomly selected on each frame, and moved with 11.5°/s in one of four diagonal directions (counterbalanced across observers). Signal dots that left the annulus 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 screen for 4 consecutive frames. Noise dots were assigned a random location within the annulus 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°/s.

Motion energy filtering

Due to the stochastic nature of the dynamic random dot kinematograms, the sensory evidence fluctuated within and across trials, around the nominal motion coherence level. To quantify behaviour and pupil responses 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).

Two spatial filters, resembling weighted sinusoids in opposite phase, were defined by

$$f_1(x,y) = \cos^4(a)\cos(4a)\exp(-\frac{y'^2}{2\sigma_g^2}) \tag{3}$$

$$f_2(x,y) = \cos^4(a)\sin(4a)\exp(-\frac{y'^2}{2\sigma_g^2}) \tag{4}$$

where $a = \tan^{-1}(x'/\sigma_c)$. The parameters $\sigma_g = 0.05$ and $\sigma_c = 0.35$ defined the carrier sinusoid and the Gaussian envelope, respectively, in line with the response properties of MT neurons (Kiani et al. 2008).

The coordinate system (x, y) was rotated to match the stimulus' target direction or its 180° opposite. Two temporal filters were defined by

$$g_1(t) = (kt)^{n_s} \exp(-kt) \left[\frac{1}{n_s!} - \frac{(kt)^2}{(n_s+2)!}\right] \quad (5)$$

$$g_2(t) = (kt)^{n_f} \exp(-kt) \left[\frac{1}{n_f!} - \frac{(kt)^2}{(n_f+2)!}\right] \quad (6)$$

Where $k = 60$ reflected the envelope of the temporal filters, and $n_s = 3$ and $n_f = 5$ controlled the width of the slow and fast filters, respectively (Kiani et al. 2008). A pair of spatio-temporal filters in quadrature pair was obtained by $f_1 g_1 + f_2 g_2$ and $f_2 g_1 - f_1 g_2$. We convolved each filter with the single-trial random dot movies. The resulting values were squared, and summed together across the two filters (Adelson and Bergen 1985).

This filtering procedure was performed for each observer's individual target direction as well as its 180° opposite. To avoid cardinal biases in motion perception, we used the four diagonals as target directions counterbalanced across observers. Outputs of the two filtering operations were subtracted to yield a direction-selective signal over time (Kiani et al. 2008). To obtain a single measure of sensory evidence per trial, we averaged over all timepoints within each stimulus interval, and took the difference between motion energy in the first and second interval as our measure of single-trial sensory evidence. Evidence strength was defined by taking the absolute value of this sensory evidence, collapsing over the two stimulus identities (Figure 2b).

Pupillometry

Observers sat in a dark room with their head in a chinrest at 50 cm from the screen. Horizontal and vertical gaze position, as well as the area of the pupil, were monitored in the left eye using an EyeLink 1000 desktop mount (SR Research, sampling rate: 1,000 Hz). The eye tracker was calibrated before each block of 50 trials.

Missing data and blinks, as detected by the EyeLink software, were padded by 150 ms and linearly interpolated. Additional blinks were found using peak detection on the velocity of the pupil signal and linearly interpolated. We estimated the effect of blinks and saccades on the pupil response through deconvolution, and removed these responses from the data using linear regression using a procedure detailed in ref (Knapen and de Gee 2016). The residual pupil time series were band-pass filtered using a 0.01 to 10 Hz second-order Butterworth filter, z-scored per run, and resampled to 100 Hz. We epoched trials, and baseline corrected each trial by subtracting the mean pupil diameter 500 ms before onset of the reference stimulus.

We included all trials from all five main sessions (i.e., excluding the practice session) in the analyses reported in this paper. The time series of consecutive trial-wise stimuli, choices, RTs and pupil responses was necessary for the regression model of serial bias modulation. Observers were well-practiced in the task structure after the practice session. As a consequence, they made few blinks during the trial intervals (on average across observers, only 7.7% of trials contained more than 50% interpolated samples). The

percentage of interpolated trials did not correlate with the estimated effect of pupil responses on serial choice bias (r = -0.268, p = 0.175, Bf_{10} = 0.369).

Quantifying pupil timecourses and single-trial responses

To characterize the time-course of uncertainty encoding in the pupil response, we regressed across-trial evidence strength onto each sample of the baseline-corrected pupil signal, separately for correct and error trials (Figure 3b). The design matrix for this regression also included an intercept and three nuisance covariates: (i) log-transformed RTs, (ii) sample-by-sample horizontal gaze coordinates and (iii) sample-by sample vertical gaze coordinates. We tested the significance of this regression timecourse using cluster-based permutation statistics (Blair and Karniski 1993).

We took the mean baseline-corrected pupil signal during 250 ms before feedback delivery as our single-trial measure of pupil response. Because of the temporal low-pass characteristics of the sluggish peripheral pupil apparatus (Hoeks and Ellenbroek 1993), trial-to-trial variations in RT can cause trial-to-trial in pupil responses, even in the absence of amplitude variations in the underlying neural responses. To specifically isolate trial-to-trial variations in the amplitude (not duration) of the underlying neural responses, we removed components explained by RT via linear regression

$$\mathbf{y}' = \mathbf{y} - (\mathbf{y}^{\mathrm{T}}\mathbf{r})\mathbf{r} \tag{7}$$

where $\mathbf{y}$ was the original vector of pupil responses, $\mathbf{r}$ was the vector of the corresponding single-trial RTs (log-transformed and normalized to a unit vector), and T denoted matrix transpose. The residual $\mathbf{y}'$ thus reflected pupil responses, after removing variance explained by trial-by-trial RTs. This residual pupil response was used for all analyses reported in the main text.

Quantifying post-error slowing

We quantified post-error slowing, for tertiles of previous trial pupil responses, as described by Dutilh et al. 2012. Briefly, we selected those error trials that were both preceded and followed by a correct trial, and subtracted the pre-error RT from the associated post-error RT. This procedure ensured that estimates of post-error slowing could not be driven by error-unrelated, intrinsic fluctuations in RT over the course of a session. Before this subtraction, we removed trial-by-trial evidence strength from RTs using linear regression, to account for the large variations in RT with stronger sensory evidence (Figure 2d).

Quantifying the psychometric function

We modeled the psychometric function (Figure S4a) as follows. The probability of a particular response $r_t = 1$ on trial t was described as:

$$P(r_t = 1|\tilde{s_t}) = \gamma + (1 - \gamma - \lambda)g(\delta + \alpha\tilde{s_t}) \tag{8}$$

where λ and γ were the probabilities of stimulus-independent errors ('lapses'), $\tilde{s_t}$ was the signed stimulus intensity (here: signed sensory evidence as in Figure 2b), $g(x) = 1/(1+e^{-x})$ was the logistic function, α was perceptual sensitivity, and δ was a bias term. The free parameters γ, λ, α and δ were estimated by minimizing the negative log-likelihood of the data (using Matlab's *fminsearchbnd*). We constrained γ and λ to be identical, so as to estimate a single, choice-independent lapse rate.

For the quantification of serial choice bias (Figure S5), we binned the data by previous choices and by previous pupil responses or RT. For each of those subsets of trials, we fit the psychometric function (equation 8) to the choices on the subsequent trials. The resulting bias term α was transformed from log-odds into probability by $p = e^{\alpha}/(1+e^{\alpha})$. This quantified $P(r_t = 1)$ for ambiguous evidence (i.e., strength of zero). Collapsing these values across the two choice options (shown separately in Figure S5) yielded the pooled measure of choice repetition probability in Figure 4a,f.

Quantifying perceptual sensitivity using cumulative Weibull function fits

In Figure 3d, S1c and S3b, we fit a cumulative Weibull function to accuracy as a function of evidence strength. The probability of a correct response $c_t = 1$ on trial t was defined as:

$$P(c_t = 1|s_t) = 1 - (0.5 - \lambda) f\left((\frac{s_t}{\theta})^{\beta}\right) \tag{9}$$

where s_t was the absolute evidence strength, $f(x) = (1 - e^{-x})$ was the cumulative Weibull function, λ was the lapse rate, θ was the threshold indicating at which level of evidence strength an accuracy of ~80% is achieved, and β was the slope of the cumulative Weibull function. The free parameters θ, β and λ were estimated by minimizing the negative log-likelihood of the data (using Matlab's *fminsearchbnd*). Perceptual sensitivity was then defined as $1/\theta$.

Modeling the modulation of serial choice bias by uncertainty-dependent variables

We modeled the pupil- and RT-linked modulation of serial choice bias by extending an established regression approach (Fründ et al. 2014). The basic regression model extended the psychometric function model (equation 8) by means of a history-dependent bias term $\delta_{hist}(\mathbf{h}_t)$, which was a linear combination of previous stimuli and choices

$$P(r_t = 1|\tilde{s_t}\mathbf{h}_t) = \gamma + (1 - \gamma - \lambda) g(\delta(\mathbf{h}_t) + \alpha\tilde{s_t}) \tag{10}$$

with

$$\delta(\mathbf{h}_t) = \delta' + \delta_{hist}(\mathbf{h}_t) = \delta' + \sum_{k=1}^{K} \omega_k h_{kt} \tag{11}$$

where the bias term $\delta(\mathbf{h}_t)$ was the sum of the overall bias δ' (see equation 8) and the history bias $\delta_{hist}(\mathbf{h}_t) = \sum_{k=1}^{K} \omega_k h_{kt}$, where ω_k were the weights assigned to each previous stimulus or choice. We here modeled

$$\mathbf{h}_t = (r_{t-1},\ r_{t-2},\ r_{t-3},\ r_{t-4},\ r_{t-5},\ r_{t-6},\ r_{t-7},\ z_{t-1},\ z_{t-2},\ z_{t-3},\ z_{t-4},\ z_{t-5},\ z_{t-6},\ z_{t-7}) \tag{12}$$

as a concatenation of the last seven responses and stimuli (Fründ et al. 2014). This procedure allowed us to quantify the effect of past trials on current choice processes (Figure 5c). We convolved every set of seven past trials with three exponentially decaying basis functions (Fründ et al. 2014). Positive history weights ω_k indicated a tendency to repeat the previous choice, or to make a choice that matched the previous stimulus. Negative weights described a tendency to alternate the corresponding history feature.

To model the effect of pupil-linked uncertainty on history biases, we extended this model by adding a multiplicative interaction term $\sum_{k=1}^{K} \omega'_k h_{kt} p_{kt}$, which described the interaction of pupil responses with the choice and stimulus identities at the last seven lags:

$$P(r_t = 1|\tilde{s_t}, \mathbf{h}_t, \mathbf{p}_t) = \gamma + (1 - \gamma - \lambda) g(\delta(\mathbf{h}_t, \mathbf{p}_t) + \alpha \tilde{s_t}) \tag{13}$$

with

$$\delta(\mathbf{h}_t, \mathbf{p}_t) = \delta' + \delta_{hist}(\mathbf{h}_t, \mathbf{p}_t) = \delta' + \sum_{k=1}^{K} \omega_k h_{kt} + \omega'_k h_{kt} p_{kt} + \omega''_k p_{kt} \tag{14}$$

where ω'_k were the history x pupil interaction weights, ω''_k were the pupil weights and

$$p_{kt} = (p_{t-1},\ p_{t-2},\ p_{t-3},\ p_{t-4},\ p_{t-5},\ p_{t-6},\ p_{t-7}\ )$$

was a concatenation of the last seven pupil responses. The term $\omega''_k p_{kt}$ acted as a nuisance covariate. To simultaneously model the effects of pupil responses and log-transformed RT, our model also included RT and history x RT terms, generated using the same procedure.

All parameters were fit using an expectation maximization algorithm. To assess whether individual observers were significantly influenced by their experimental history, we ran 1,000 iterations of permuting all trials, fitting the full model, and subsequently comparing the likelihood of the intact model to this null distribution (where permutation nullifies true history effects) (Fründ et al. 2014). Confidence intervals for individual regression weights were obtained from a bootstrapping procedure.

Serial bias and outcome-dependent choice strategies

The history weights for past stimuli and responses allowed us to characterize different decision strategies (Fründ et al. 2014) (Figure 5d). Positive weights associated with the previous choice, or the previous stimulus

category, indicate a tendency to repeat this previous choice, or to make a choice corresponding to the previous stimulus, respectively. Negative weights correspond to a tendency to alternate previous choice or stimulus. In the left and right triangle of this strategy space, the magnitude of the response weight is larger than the magnitude of the stimulus weight. Consequently, strategies are dominated by the previous choice and can be simply defined as choice alternation (left) or choice repetition (right).

In the upper and in the lower triangle, the magnitude of the stimulus weight is larger than the magnitude of the response weight, so strategies are dominated by the identity of the previous stimulus (which is only known to the observer as a function of their previous response and feedback). In the upper and lower triangle, strategies are thus defined by the sign of the stimulus weight. In the upper triangle stimulus weights are positive, indicating a tendency to repeat the previous stimulus. On a correct trial, previous choice and stimulus are equal and therefore, repeating the previous stimulus is equal to repeating the previous choice (a win-stay strategy). On errors, the previous choice is opposite to the previous stimulus and repeating the previous stimulus is equal to alternating the previous choice (lose-switch strategy). Conversely, in the lower triangle stimulus weights are negative, reflecting a tendency to alternate the previous stimulus. This implies a tendency to alternate the previous choice if the previous choice was correct (win-switch strategy) and a tendency to repeat the previous choice in case of a previous error (lose-stay strategy).

The weights for previous choices and stimuli can easily be combined to obtain weights reflecting a tendency to repeat previous correct or incorrect choices (Figure S6). Specifically, correct weights are defined by $choice + stimulus$, and error weights by $choice - stimulus$ (Fründ et al. 2014). The same holds for modulation weights. This transformation is identical to fitting a model with regressors for previous successes and failures (Busse et al. 2011; Abrahamyan et al. 2016).

Statistical tests

We used non-parametric permutation testing to test for the group-level significance of individually fitted parameter values (Figure 3 and Figure 5e,g). We randomly switched labels of individual observations either between two paired sets of values, between one set of values and zero, or between two unpaired groups. After repeating this procedure 10,000 times, and computing the difference between the two group means on each permutation, the p-value was the fraction of permutations that exceeded the observed difference between the means. All p-values reported were computed using two-sided tests.

In Figure 4, we split the data into tertiles of pupil response or RT, and computed next trial serial choice bias, signed choice bias, overall choice bias, perceptual sensitivity, lapse rate, RT and post-error slowing in each bin. We used a repeated-measures ANOVA to test for the main effect of bin on each dependent variable. We further used Bayes Factors (Bf), obtained from a Bayesian one-factor ANOVA (Rouder et al. 2012), to support conclusions about null effects observed. Bf_{10} quantifies the evidence in favour of the null or the alternative hypothesis, where $Bf_{10} < \frac{1}{3}$ or > 3 is taken to indicate substantial evidence for H_0 or H_1, respectively. $Bf_{10} = 1$ indicates inconclusive evidence. We similarly computed Bf_{10} for correlations, based

on the Pearson correlation coefficient (Wetzels and Wagenmakers 2012).

The p-value for the difference between the two correlation coefficients (choice weight by pupil modulation weight vs. choice weight by RT modulation weight), shown in Figure 5f, was obtained through permutation testing. To generate a null distribution of no difference, we randomly switched (or not, dependent on a simulated coin flip) each observer's RT and pupil modulation weights, after which we computed the between-subject correlation between choice weights and pupil modulation weights as well as between choice and RT modulation weights. Repeating this procedure 10,000 times generated a distribution of the difference in correlation, under the null hypothesis of no difference.

Data availability

All raw and processed data, as well as the code to reproduce all analyses and figures, are available at `http://dx.doi.org/10.6084/m9.figshare.4300043`.

Acknowledgements

We thank O'Jay Medina for assistance with data collection, all members of the Donnerlab for valuable discussions, and Konstantinos Tsetsos, Jan Willem de Gee, Niklas Wilming, Camile Correa, Florent Meyniel and Sander Nieuwenhuis for helpful comments on the manuscript. We acknowledge computing resources provided by NWO Physical Sciences. This research was supported by the German Academic Exchange Service (DAAD) and G.-A. Lienert Foundation (to A.E.U.) and the German Research Foundation (DFG), SFB 936/A7, SFB 936/Z1, DO 1240/2-1 and DO 1240/3-1, and European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project) (to T.H.D.).

Supplementary Figures

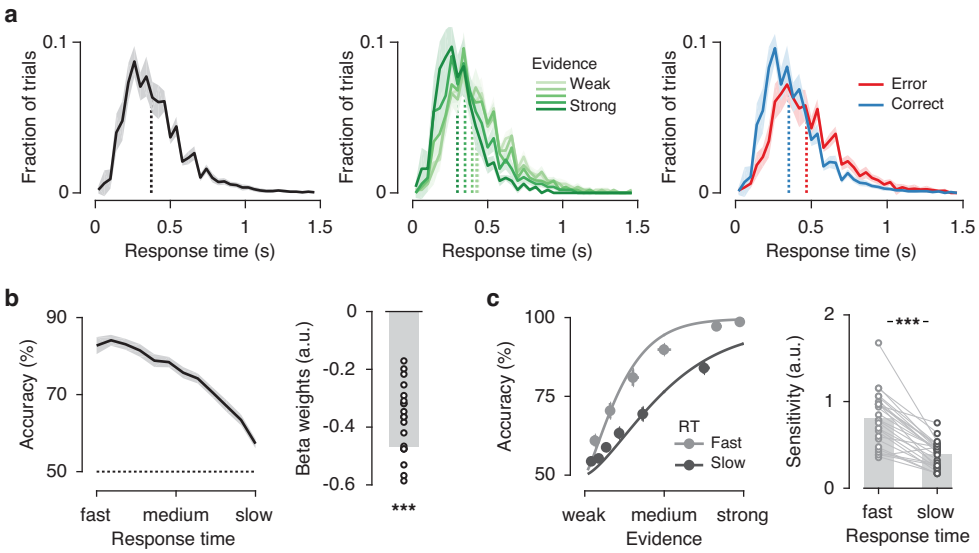

Figure S1. RTs scale with decision uncertainty. (**a**) RT distributions from stimulus offset, shown for all trials (left), split into five bins of evidence strength (middle), and separately for correct and error trials (right). For each observer, the number of trials was counted in each 40-ms wide bin from 0 to 1.5 seconds after stimulus offset, and normalized by the total number of trials. Shaded error bars indicate group median and inter-quartile range. Dotted line indicates group mean of individual RT medians. (**b**) RT predicted choice accuracy over a range from about 85% to about 60% correct, and not below chance level (50%). This relationship is consistent with decision uncertainty, but not error detection, which predicts accuracies of a range from 100% to 0% correct. Left: Accuracy for 12 bins of RT, shaded error bars indicate group mean ± s.e.m. Right: Individual logistic regression weights, using RT to predict single-trial accuracy. (**c**) Slow RTs reflected lower perceptual sensitivity. Left: Average cumulative Weibull psychometric function fits and data points (group mean ± s.e.m.), separately for the lowest and highest RT tertiles. Right: Individual perceptual sensitivity, separately for lowest and highest RT tertiles. In **b-c**, we z-scored and log-transformed RTs within each block and removed trial-to-trial variability shared with pupil responses via linear regression before computing statistics. *** $p < 0.001$, permutation test. (N=27, group mean ± s.e.m.)

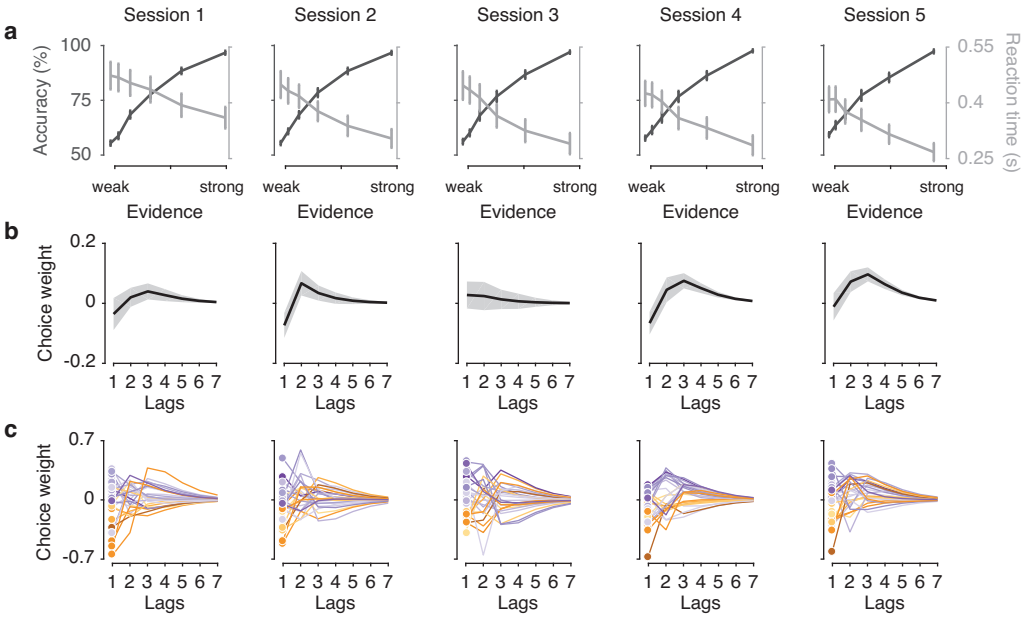

Figure S2. Behaviour over sessions. Data from each observer were collected over 5 main experimental sessions of 500 trials. Data from the practice session are not shown here. After discarding trials in which no response was recorded, each session contained an average of 498 trials (range 465-500). (**a**) Psychometric and chronometric functions, as in Figure 2d, separately for each session. (**b**) History kernels as in Figure 5c, separately for each session. (N=27, group mean ± s.e.m.) (**c**) Individual history kernels as in Figure 5c, separately for each session. Colors indicate the choice weight as derived from the model in Figure 5c-d, fit across all sessions combined. To complement these visual representations of behaviour over sessions, we computed repetition probability for three bins of pupil responses (Figure 4a), separately in each of the five sessions. Using a repeated measures ANOVA, we found no main effect of session ($F_{4,104}$ = 1.591, p = 0.182, Bf_{10} = 0.078) nor an interaction between session and pupil bin ($F_{8,208}$ = 1.333, p = 0.229, Bf_{10} = 0.023) on repetition probability. This analysis indicates that history biases do not detectably change over the course of learning, adding further evidence to the idea that serial choice biases are stable, individual traits.

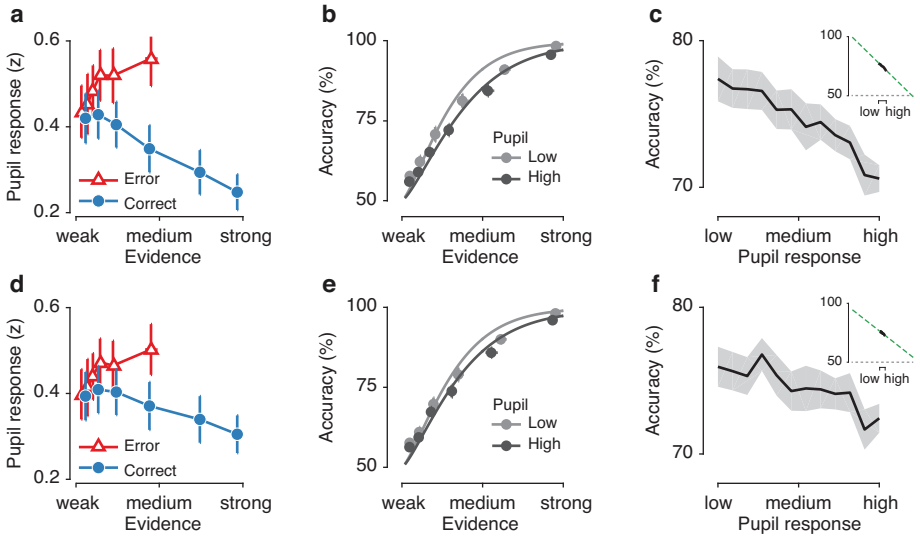

Figure S3. Pupil responses scale with decision uncertainty. Testing for all three signatures of decision uncertainty derived from the model in Figure 1. (**a**) Pupil responses scaled oppositely with evidence strength on correct and error trials. (**b**) High pupil responses reflected lower perceptual sensitivity. Average cumulative Weibull psychometric function fits (see Methods) and data points, separately for the lowest and highest tertiles of pupil responses. (**c**) Accuracy as a function of pupil responses (12 bins). Pupil responses predicted uncertainty over a range between 100% and 50% correct, but clearly not below 50% correct. This scaling is more consistent with decision uncertainty than with error awareness (which predicts accuracies down to 0%). Note that the analysis is limited by noise corrupting the single-trial pupil measurements. To address this issue, we fit a line to the data in **c**, extended its negative range to reach 100% accuracy, and then extended its positive range, with an equal distance. The result, shown in the inset, provided a rough estimate of the relationship expected, based on our result, if single-trial pupil-linked arousal could be measured without noise. Again, this analysis indicates that the scaling of pupil responses with accuracy is more consistent with decision uncertainty than with error awareness. (**d-e**) Same as **a-c**, after removing trial-by-trial fluctuations in log-transformed RT from the pupil signal using linear regression. The scaling of the pupil response with decision uncertainty was not inherited from the analogous scaling of RT. (N=27, group mean ± s.e.m.)

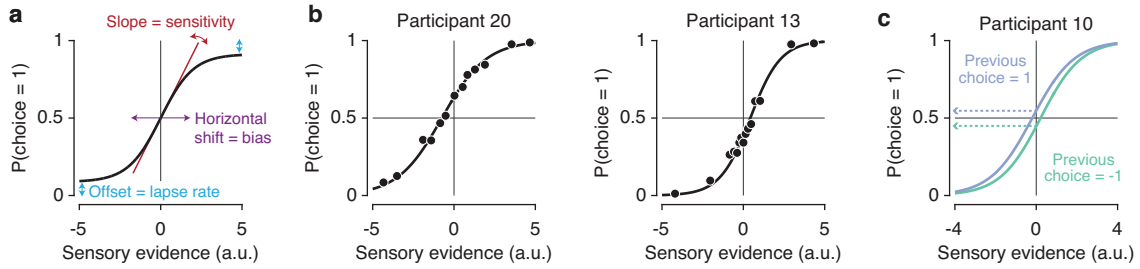

Figure S4. Quantifying choice bias using psychometric function fits. (**a**) A logistic psychometric function quantifies separate aspects of choice behaviour. The slope of the function indicates the observer's perceptual sensitivity. The intercept indicates a horizontal shift of the psychometric function, reflecting a bias towards a specific choice independent of the sensory evidence. The vertical offsets from the two asymptotes indicate the fraction of stimulus-independent errors ('lapses'). See also Methods. (**b**) Example psychometric functions with corresponding data points, for an example observer with a bias towards $\mathrm{choice}1$ (left) and an observer with a bias towards choice_{-1} (right). (**c**) History-dependent choice bias. Example observer with a tendency to repeat the previous choice.

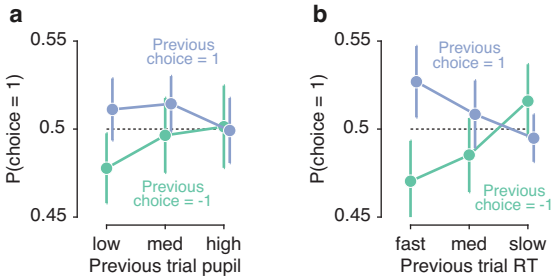

Figure S5. Pupil modulation of history-dependent choice bias. (**a**) Modulation of repetition probability by previous trial pupil response. $P(choice = 1)$ was computed from the intercept of the logistic function (see Methods), for tertiles of previous trial pupil responses. (**b**) as in a, but for tertiles of previous trial RT. The two choice identities were collapsed to obtain the measure of repetition probability in Figure 4a,f. (N=27, group mean ± s.e.m.)

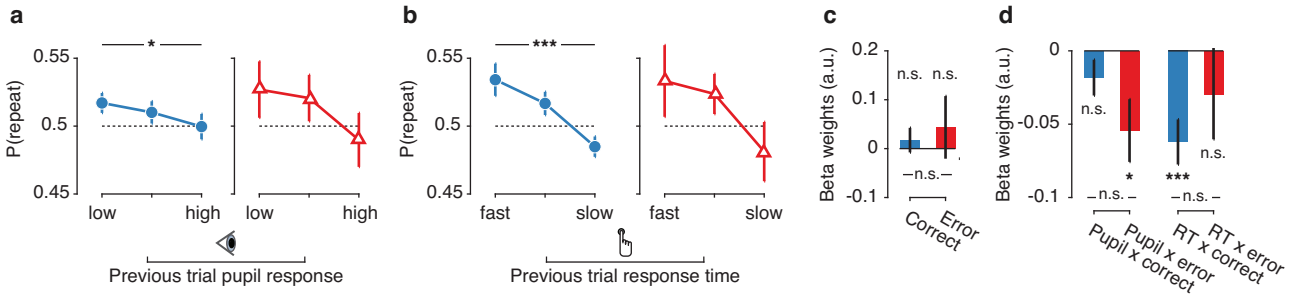

Figure S6. Pupil responses and RT modulate serial choice bias within categories of accuracy. (**a**) Repetition probability, for tertiles of previous trial pupil responses, separately for correct (left) and error (right) trials. (**b**) As in **a**, but for tertiles of previous trial RT. (**c**) Beta weights for repeating a previous correct vs. incorrect choice (see Methods). (**d**) Pupil- and RT modulation weights for repeating a previous correct vs. incorrect choice. Beta weights were obtained from the model show in Figure 5e-g, with pupil- and RT-linked modulatory terms included in the same regression model. Statistics indicate the main effect of a one-way ANOVA (**a, b**) or a permutation test (**c, d**). *** $p < 0.001$,* $p < 0.05$, n.s. $p > 0.05$. (N=27, group mean ± s.e.m.) These results indicate that the modulatory effect of pupil responses (and RT) on serial choice biases was not purely driven by higher pupil responses on error trials. Instead, serial choice bias was modulated by trial-to-trial fluctuations in pupil-linked arousal within categories of trial outcomes.

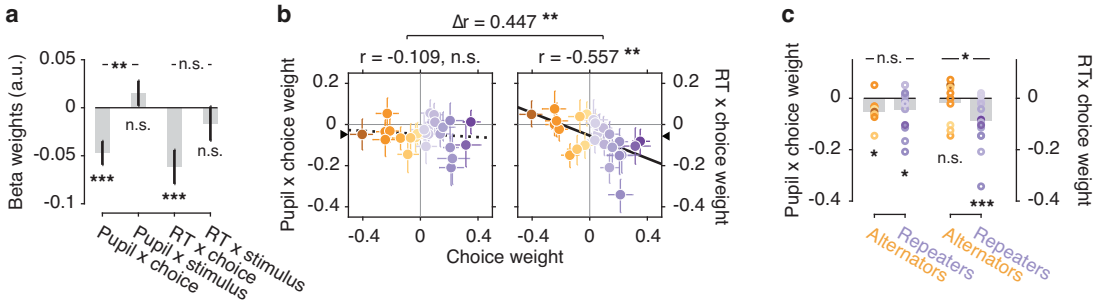

Figure S7. Modeling results do not depend on simultaneous fitting of both pupil and RT. Running two separate regression models, one including only pupil response and one only including RT as a modulatory variable, gives the same results as shown in Figure 5 (where the two were included in the regression model simultaneously). (**a-c**) as in Figure 5e-g, but with data obtained from two separate regression models.

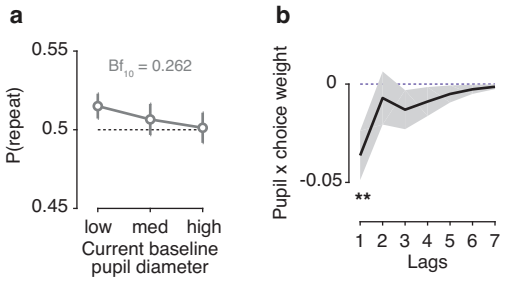

Figure S8. Predictive effect is specific for pupil response on the preceding trial. (**a**) Baseline pupil diameter on the current trial did not predict a modulation of serial choice bias. Repetition probability, for tertiles of current trial baseline pupil diameter (main effect of one-way repeated measured ANOVA, $F_{2,52}$ = 1.164, p = 0.320). (**b**) Pupil modulation of choice bias was only significant (** p < 0.01) across the group of observers at lag 1 (same data as Figure 5e), and did not reach significance beyond one trial in the past. This finding indicates that the modulation of choice biases by pupil responses was more short-lived than the overall serial choice biases shown in Figure 5c. (N=27, group mean ± s.e.m.)

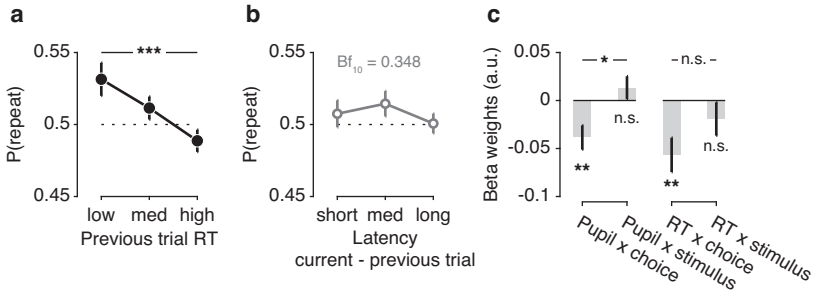

Figure S9. Serial choice biases are not explained by variations in interval timing. (**a**) To measure the passage of time between trials, we computed the latency between the onset of each test stimulus and the onset of the next trial?s test stimulus. These latencies correlated with RTs (mean Spearman's ± 0.296, range 0.086 to 0.726). Removing these trial-by-trial latencies from RT (using linear regression) did not abolish the effect of RTs on serial choice bias (main effect of RT bin, $F_{2,52}$ = 10.846, p < 0.001, Bf_{10}ha = 225.756). (**b**) Latencies did not predict a modulation of serial choice bias (main effect of latency bin, $F_{2,52}$ = 1.541, p = 0.224, Bf_{10} = 0.349). These results suggest that the uncertainty component of RTs, rather than the passage of time between trials, modulated serial choice bias. (**c**) We tested whether the modulation of serial bias by pupil response could be explained by trial-to-trial variations in the jittered interval between s1 and s2, or between button press and feedback delivery. When these random variations were long, they could cause larger pupil responses, irrespective of the amplitude of the underlying neural input, by driving the peripheral pupil apparatus for a longer duration. We removed these trial-to-trial interval durations from pupil responses using linear regression, and reran the analysis shown in Figure 5e. Although pupil responses were weakly correlated to the interval between s1 and s2 (mean Spearman's ρ -0.007, range -0.055 to 0.047, significant in 3 out of 27 observers) and the interval between button press and feedback (mean Spearman's ρ 0.056, range -0.025 to 0.290, significant in 13 out of 27 observers), removing this variance from trial-by-trial pupil responses did not change the predictive effect of pupil responses on serial choice bias. Statistics indicate the main effect of a one-way ANOVA (a, b) and permutation test (c). *** p < 0.001, ** p < 0.01, * p < 0.05, n.s. p < 0.05. (N=27, group mean ± s.e.m.)

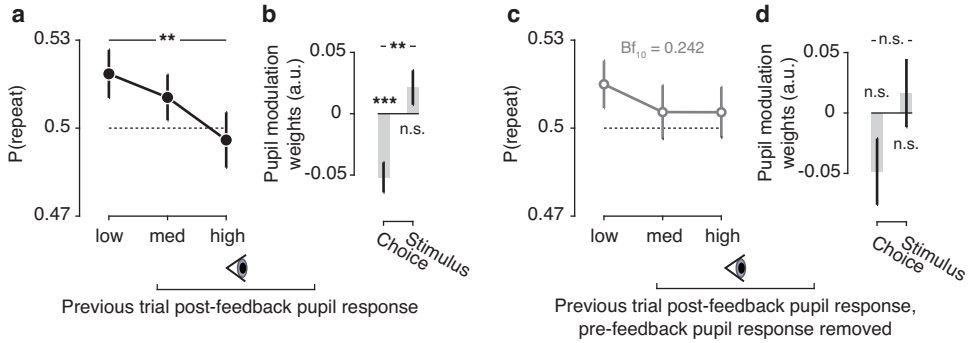

Figure S10. Serial choice biases are not explained by post-feedback pupil responses. To test whether serial choice biases were modulated by pupil responses to the feedback tone beyond pre-feedback uncertainty signaling, we computed post-feedback values as the mean pupil diameter 515-765 ms after feedback tone delivery. This window was defined as the peak of the grand average pupil response, and its length set equal to our pre-feedback window. (**a**) Serial choice bias for tertiles of previous trial post-feedback pupil responses (main effect of pupil bin, $F_{2,52}$ = 5.479, p = 0.007, Bf_{10} = 6.014). (**b**) Beta weights for the interaction between previous trial post-feedback pupil response and choice or stimulus, as in Figure 5e. (**c**) We removed the effect of single-trial pre-feedback from the post-feedback signal using linear regression. The residual reflected the effect of feedback on uncertainty scaling in the pupil, after taking into account the scaling already present before the feedback tone. Serial choice bias, for tertiles of residual pupil responses (main effect of pupil bin, $F_{2,52}$ = 1.063, p = 0.353) (d) Modulation weights for post-feedback pupil responses, with pre-feedback pupil responses added as a covariate in the same regression model. The information about serial biases was already contained in the pupil signal before feedback delivery. Statistics indicate the main effect of a one-way ANOVA (**a, c**) and permutation test (**b, d**) . *** p < 0.001, ** p < 0.01, n.s. p > 0.05. (N=27, group mean ± s.e.m.)

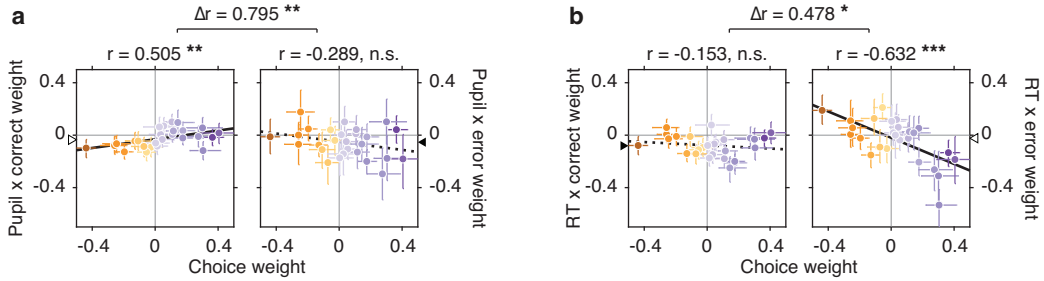

Figure S11. Differential gating of individual choice modulation by trial outcome. Pupil-and RT-linked modulations of serial choice bias were differentially gated by trial outcome. We computed correct and error modulation weights from choice and stimulus modulation weights (see Methods). (**a**) Correlation between choice weights and pupil modulation weights, separately for correct and incorrect choices. (**b**) Correlation between choice weights and RT modulation weights, separately for correct and incorrect choices. Colors indicate the choice weight as derived from the basic model in Figure 5c. Error bars indicate a 68% confidence interval obtained from a bootstrap. Triangles mark the intercept of a linear regression line; filled triangles indicate a group-level effect different from zero (as in Figure S6d). *** p < 0.001, ** p < 0.01, * p < 0.05, n.s. p > 0.05. Figure 5f shows that RT reduced observers' intrinsic serial biases while pupil responses generally promoted choice alternation. These results further dissociate these modulatory effects, in showing that they were 'gated' by trial outcome in distinct ways: Large pupil-linked arousal pushed observers to increase their intrinsic serial bias after correct trials, as indicated by the positive correlation in a. After error trials, on the other hand, a correlation of the opposite sign was observed indicating that across trial outcomes, these two effects nullified and lead to an overall boost in alternation. This stood in sharp contrast to the group-level effect of RT, which predicted a reduction in intrinsic serial bias across the group. This effect was strongly present after error trials (**b**), suggesting an adaptive control mechanism could be at work only after negative feedback is received. After correct trials, high RTs indicated a slight reduction in bias, but this negative correlation was not significant across the group.

2.2 | Confidence-dependent accumulation of past decision variables biases perceptual choice

Braun A, Urai AE, Donner TH. (2017) *bioRxiv*, 172049.

Abstract

Perceptual decision-making is biased by previous events, including the history of preceding choices: Observers tend to repeat (or alternate) their choices more often than expected by chance. Here, we find that such sequential choice biases result from the accumulation of action-independent, internal decision variables. Human observers performed different variants of a visual motion discrimination task around psychophysical threshold. First, we decoupled perceptual choices and motor responses on a trial-by-trial basis, to disentangle their relative contributions to sequential biases. The impact of previous choices outweighed the impact of previous motor responses and previous stimuli. Second, observers performed the task in both random and biased environments (containing systematic tendencies towards either repetition or alternation of motion directions) and in the absence of external feedback about the correctness of their choices. Observers adapted their sequential choice biases to the biased environmental statistics, in a way that indicated memory for previous choices and predicted their overall performance. We further found that this adaptation was enhanced by the confidence about the correctness of previous choices, consistent with the idea that decision variables driving choice encode both the categorical choice, as well as the graded confidence associated with it. Thus, our results are consistent with the idea that high-level decision variables, dissociable from both, sensory input and motor output, are accumulated across choices towards biases for upcoming choices. These insights constrain the candidate neural sources of sequential choice biases.

 doi: 10.1101/172049. AUTHOR CONTRIBUTIONS: A.B., A.E.U. and T.H.D. designed research; A.B. performed research; A.B. analyzed data; AEU re-analyzed the data with separate code; A.B., A.E.U. and T.H.D. wrote the paper; T.H.D. supervised research. At press time, the manuscript was under revision at *Journal of Neuroscience*.

Introduction

Organisms' judgments about their environment do not only depend on the current sensory input, but are systematically biased by the history of preceding choices and stimuli (Fründ et al. 2014). Several studies have found that observers repeat (or alternate) their perceptual choices more often than expected by chance (Fernberger 1920; Gold et al. 2008; de Lange et al. 2013; Akaishi et al. 2014; Fischer and Whitney 2014; Abrahamyan et al. 2016; Pape and Siegel 2016; St. John-Saaltink et al. 2016; Fritsche et al. 2017; Urai et al. 2017). Such sequential biases are ubiquitous across domains of decision-making (Leopold et al. 2002; Allefeld et al. 2013; Padoa-Schioppa 2013).

One idea holds that sequential choice biases result from accumulation of environmental states across decisions. Computational models postulate that sequential choice biases result from the accumulation of decision-related neural signals (Glaze et al. 2015; Bonaiuto et al. 2016). In that way, observers may update their prior beliefs about the next stimulus and can adapt behavior to their environment (Yu and Cohen 2008; Glaze et al. 2015). In standard perceptual choice tasks used in the laboratory, stimulus sequences are typically uncorrelated. In this context, sequential biases in observers' choices degrade their performance (Abrahamyan et al. 2016). However, more recent studies introduced perceptual choice tasks with correlated stimulus sequences (Goldfarb et al. 2012; Glaze et al. 2015; Abrahamyan et al. 2016; Kim et al. 2017). Here, sequential biases should improve performance.

Real-life decisions are often associated with uncertainty about their outcome, due to weak evidence and the absence of immediate feedback. Then, decision-makers can only accumulate their own inferences about the state of the environment. These inferences are encoded in the so-called 'decision variable', on which decision-makers base their categorical choices (Bogacz et al. 2006; Gold and Shadlen 2007). Neural correlates of the decision variable are distributed across many brain regions (Gold and Shadlen 2007; Siegel et al. 2011; Brody and Hanks 2016). They are expressed as evolving motor plans (Gold and Shadlen 2007; Donner et al. 2009; de Lange et al. 2013) as well as in action-independent formats (Bennur and Gold 2011; Hebart and Hesselmann 2012; O'Connell et al. 2012; Hebart et al. 2016). Decision variables not only reflect the categorical choice but also the observer's confidence about the correctness of that choice, henceforth referred to as decision confidence (Kiani and Shadlen 2009; Hebart et al. 2016).

These considerations raise a number of questions. Do sequential biases emerge from the accumulation of decision variables in motor or action-independent formats? Does this accumulation help to adapt sequential biases to the statistical structure of the evidence? Does the strength of such adaptation depend on decision confidence? The latter might be particularly important under conditions, in which no unambiguous external feedback about the true stimulus category is available. Here, we addressed these issues in a visual motion discrimination task widely used in the neurobiology of choice (Gold and Shadlen 2007). To this end, we systematically modeled human choice behavior under experimental manipulations tailored to address the above questions.

Materials and Methods

We report results from two experiments (referred to as Experiment 1 and 2) quantifying sequential biases during a random dot motion discrimination (up vs. down) task widely used in monkey (Gold and Shadlen 2007) and human (Donner et al. 2009; de Lange et al. 2013; Kelly and O'Connell 2013; Hebart et al. 2016) work on perceptual choice. Analyses of behavior irrespective of sequential effects from Experiment 1 were previously published (Tsetsos et al. 2015). Here, we re-analyzed these data to quantify the dependence of choice on previous stimuli, choices, and motor responses. The two experiments aimed at manipulating on different aspects of choice behavior, and thus used different variants of the basic random dot motion task. We used variants of the same statistical approach (Fründ et al. 2014) to model sequential biases in both.

Participants

Six healthy human observers (2 males, mean age: 25; range: 22–29 years) participated in Experiment 1, which was approved by the ethics committee of the Department of Psychology of the University of Amsterdam (reference number 2011-OP-1588). 26 healthy observers (15 female, mean age: 26, range: 20 - 36) participated in Experiment 2, which was approved by the local ethical review board (Ärztekammer Hamburg, reference number PV4714). Four participants were excluded from the data analysis: Three of those did not complete all sessions and one exhibited substantially worse performance than the rest of the group (64 percent correct overall, and 63 percent correct for the easiest trials in the three levels with highest motion coherence). 22 participants remained for the data analysis. All observers gave their written informed consent.

Experimental design

Experiment 1

The following description summarizes the aspects of the experimental design that were most important to current paper; a comprehensive description can be found in our previous paper (Tsetsos et al. 2015). Random dot kinematograms (Figure 1A) were composed of 785 (average) white dots on a black screen. The dots were moving within a circle of 9.1° diameter. A red fixation cross of 0.4° x 0.4° was centered in the middle of the circle. The dot density was 12.07 dots/deg^2. The population of dots was split into "signal dots" and "noise dots". The signal dots moved either upwards or downwards with a velocity of 2.6°/s. The noise dots changed position randomly from frame to frame. The percentage of signal dots defined the motion coherence, a measure of motion strength. On each trial, three different sequences of dot motion (at the same coherence and direction) were presented in an interleaved fashion, making the effective speed of signal dots 0.87 °/s. Further, one of six different levels of motion coherence, (0.05, 1.26, 3.15, 7.92, 19.91 and 50%) and one of six different viewing durations (150, 300, 600, 1200, 2400, and 4800 ms) were chosen randomly, under the constraint that they occurred equally often within a block of 144 trials.

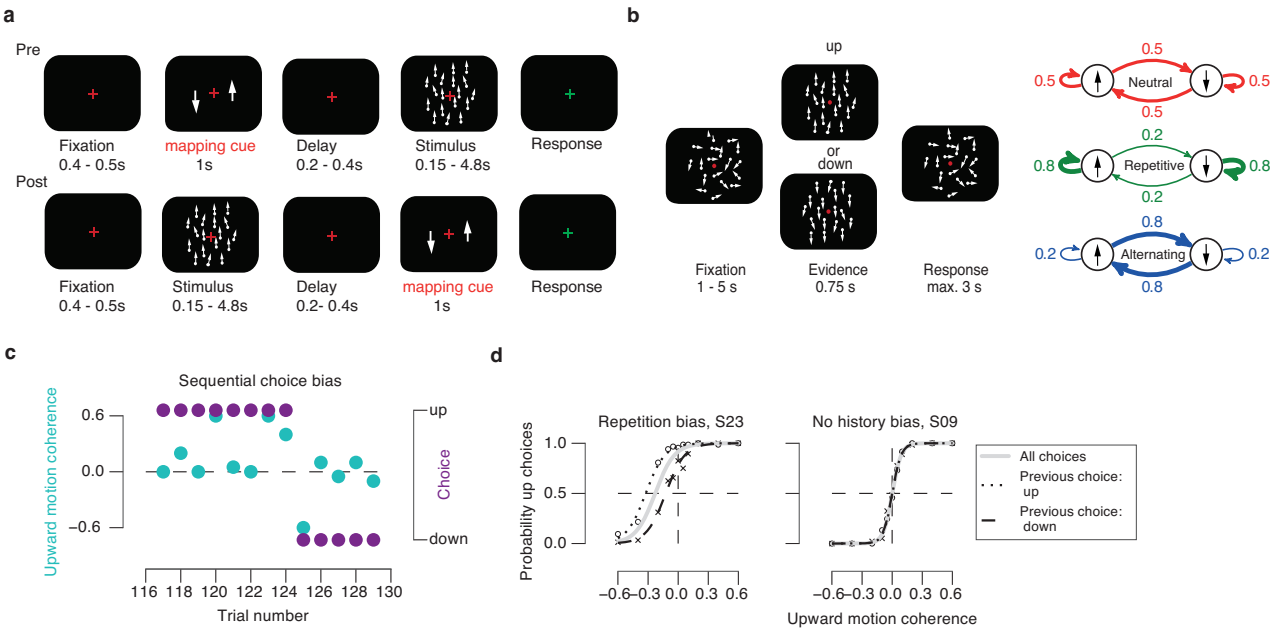

Figure 1. Quantifying sequential choice bias and behavioral task. (**a**), (**b**) Behavioral task protocols. In all experiments presented here, the net direction and motion coherence of a dynamic random dot pattern varied from trial to trial, and observers were asked to judge the net direction ("up" vs "down"). (**a**) Experiment 1. Random dot motion task with variable decision-response mapping. At the beginning of each trial a red fixation cross was shown for 0.4 – 0.5s. A decision-response cue was shown for 1s before the presentation of the stimulus in the Pre condition and after stimulus presentation in the Post condition. Six different durations of the stimulus presentation ranging from 0.15s to 4.8s were used. In both conditions the presentations of the mapping cue (instructing the mapping between choice and motor response) and of the stimulus were separated by a delay of 0.2 – 0.4s, in which the fixation cross was shown. Participants were instructed to respond after stimulus presentation in the Pre condition and after the presentation of the mapping cue in the Post condition. Auditory feedback was given after incorrect responses. (**b**) Experiment 2. Top: Random dot motion and a red fixation cross were shown throughout the task. Each trial started with a fixation interval of 1-5s. A beep indicated the onset of the decision interval, in which a coherent motion signal of dots moving either upwards or downwards was shown for 0.75s. A second beep indicated the offset of the evidence and the onset of the response interval, which ended as soon as the participant pressed one of two buttons to state their decision or after 3s in case no response was given. Bottom: Three different manipulations of the repetition probability between stimulus directions were used across different sessions, resulting in a Neutral condition (random stimulus sequence with repetition probability of 0.5), a Repetitive condition (stimulus repetition probability of 0.8) and an Alternating condition (stimulus repetition probability of 0.2). (**c**) Local streaks of repeated choices in the face of ambiguous or weak evidence of opposite sign (example observer from Experiment 2). (**d**) Psychometric functions of two example participants from Experiment 2 conditioned on previous choice. Positive values of stimulus intensity correspond to upward motion and negative ones to downward motion. The dotted line referred to previous 'up 'choices (i.e., choice repetitions) and the dashed curve to previous 'down' choices (i.e., choice alternations). Left: Leftward shift from dashed to dotted line, indicating a bias to repeat the previous choice. Middle: Overlap of dashed and dotted curves, indicating absence of sequential bias. Gray line, psychometric function computed based on all trials, irrespective of previous choice.

Stimuli were presented on a 22-inch CRT monitor with a resolution of 800 x 600 pixel and a frame rate of 100 Hz at a viewing distance of 68 cm. The participants were instructed to maintain their gaze on the red cross throughout the trial and judge the net motion direction. The motion viewing interval was followed by a variable delay (uniform distribution ranging from 200 to 400 ms), after which the observers had to report their choice by pressing one of two buttons on a computer keyboard, with either the left or the right index finger. Participants received auditory feedback after incorrect responses (a 1000 Hz tone of 100 ms). Perceptual choices ('up' vs. 'down' motion direction) were decoupled from motor responses (left vs.

right button press) by varying their mapping from trial to trial. This mapping was instructed before motion viewing in one condition ('Pre' condition) and after motion viewing in the other ('Post'), by means of a visual cue that mapped each of the two possible net motion directions (presented as arrows) onto the left or right side. This mapping was randomly selected on each trial. Conditions alternated across blocks. Observer 1-5 participated in both conditions. Participant 6 participated only in the Post condition. We pooled the data from participants 1-5 across both conditions.

Experiment 2

To test for the adaptability of sequential choice biases, we manipulated the sequential stimulus statistics between experimental sessions, to make people perform the task in 'Repetitive', 'Neutral' (no sequential dependence), or 'Alternating' environments. Stimuli, task, and procedure for Experiment 2 were identical to Experiment 1, with the following exceptions (Figure 1B). The circle within which the dots were moving had an outer radius of 12° and an inner radius of 2°. The density of dots was 1.7 dots/deg^2 and each dot had a diameter of 0.2°. The dots moved with a velocity of 11.5°/s and each dot had a lifetime of 6 frames. We used the following coherence levels: 0, 5, 10, 20, 40 and 60% (equally many trials per coherence level). A red bulls-eye fixation target (Thaler et al. 2013) at the center of the screen as well as randomly moving dots (0% coherence) were presented throughout each block. The first trial of each block started with a baseline interval of 5 s, with random dot motion. A beep (duration: 50 ms, sine wave at 440 Hz) indicated the onset of the motion-viewing interval with a fixed duration of 0.75 s and variable coherence levels and directions (see above). A second beep indicated the offset of the motion-viewing interval and prompted the observers' response. Observers reported their perceptual choices by pressing one of two keyboard buttons, with the index finger of the left or right hand. After button press (or a response deadline of 3 s), the inter-trial interval started. Inter-trial intervals were uniformly distributed between 1 and 5 s. Observers received auditory feedback during the training sessions, but no feedback during the subsequent six sessions of the main experiment. The motion viewing duration of 0.75 s was selected because a previous human study found little integration of motion information beyond that duration (Tsetsos et al. 2015), similar to what had been observed in monkeys (Kiani et al. 2008). We used a fixed mapping between choices and motor responses, whereby the two possible mappings (right-hand button for up, left-hand for down, or vice versa) were counterbalanced across observers. Experiment 2 consisted of seven sessions per participant (one for training and six main sessions), whereby each session was divided into 10 blocks of 60 trials. Critically, the transition probabilities (over trials) between the two alternative motion directions (i.e. up vs. down) were manipulated across experimental sessions (Figure 1B, right). Specifically, the probability of a repetition of motion directions (regardless of their identity) was defined as

$$\begin{aligned} &\mathrm{P}(stimulus\ repetition) = 1 - \mathrm{P}(stimulus\ alternation) \\ &= \mathrm{P}(\mathrm{stimulus}_n = \mathrm{up} \mid \mathrm{stimulus}_{n-1} = \mathrm{up}) \\ &= \mathrm{P}(\mathrm{stimulus}_n = down \mid \mathrm{stimulus}_{n-1} = down) \end{aligned} \tag{1}$$

whereby n indexes trials. This repetition probability was held constant within each session, but varied across the main experimental sessions between the following values: 0.5 in the 'Neutral' condition, 0.8 in the 'Repetitive' condition, and 0.2 in the 'Alternating' condition. The Neutral condition allowed for quantifying observers' intrinsic sequential choice bias, which we used as a baseline for quantifying their adaptation to the biased sequential statistics of the Repetitive and Alternating conditions.

During the training session, the motion direction on each trial was chosen randomly and independently, and all participants started with the 'Neutral' condition in session 1 of the main experiment (which was repeated in session 4). Half of the participants then performed the Repetitive condition in sessions 2 and 5, and the Alternating condition in sessions 3 and 6. For the other half of participants, the order of the Repetitive and Alternating conditions was flipped.

Observers were instructed to maintain stable fixation and perform the motion discrimination task as accurately as possible. They were informed that the sequential statistics of the stimulus identities would change from session to session, but stay constant within each session. Specifically, they were informed that the stimulus sequences could be 'like produced by a coin flip' (Neutral condition), 'more likely repeating than alternating' (Repetitive condition), or 'more likely alternating than repeating' (Alternating condition). Observers were neither informed about the order of these conditions, nor about the exact transition probabilities, nor about the use of this information for their choice behavior.

Modeling sequential choice bias

We used logistic regression (Fründ et al. 2014; Urai et al. 2017) to model observers' sequential choice biases under the different experimental conditions assessed here. The basic approach consisted of adding a linear combination of different components of trial history (which depended on the experiment), as a bias term to a logistic function model of the choice probability. Another variant of this approach estimates the effect of previous choices, contingent on their success or failure (Busse et al. 2011; Abrahamyan et al. 2016). We here used a variant that quantified the relative contributions of previous stimuli, choices, and (for Experiment 1) motor responses. We also assessed the effect of previous correct vs. incorrect choices, by re-combining the weights for previous stimuli and choices (Fründ et al. 2014), and the modulation of choice weights by reaction time (for Experiment 2).

Standard psychometric function fit

The probability of making one of the two choices $r_t = 1$ ($r_t = 1$ for 'choice up', $r_t = -1$ for 'choice down') on trial t, given the signed stimulus intensity $\tilde{s_t}$ (i.e., motion coherence times direction) was described by:

$$P(r_t = 1|\tilde{s_t}) = \gamma + (1 - \gamma - \lambda)\ g(\delta + \alpha\tilde{s_t}). \tag{2}$$

where γ and λ were the lapse rates for the choices $r_t = 1$ and $r_t = -1$, and $g(x) = \frac{1}{1+e^{-x}}$ was the logistic function. The bias term δ, the offset of the psychometric function, described the overall bias for one specific

choice. α was the slope of the stimulus-dependent part of the psychometric function, describing perceptual sensitivity.

For visualizing the effect of previous on current choice (Figure 1D), we separated the data from the Neutral condition into two groups, conditioned on the choice from the previous trial, and fitted the psychometric function, i.e. the probability of upward choices, separately for each of the two conditions. Results from two example observers are shown in Figure 1D. The first example observer (left panel) showed a sequential bias towards repeating choices, evident as a rightward shift of the dashed compared to the dotted line. The second observer (right panel) exhibited no sequential choice bias.

Psychometric function fit with history contributions for stimulus and choice

We estimated the contribution of the recent seven stimuli and choices by adding a history-dependent bias term to the argument of the logistic function (Fründ et al. 2014):

$$P\left(r_t = 1 | \tilde{s_t}, \mathbf{h_t}\right) = \gamma + \left(1 \ \gamma - \lambda\right) \ g\left(\delta(\mathbf{h_t}) + \ \alpha \tilde{s_t}\right) \tag{3}$$

with

$$\delta\left(\mathbf{h_t}\right) = \ \delta^{'} + \ \delta_{\text{hist}}\left(\mathbf{h_t}\right) = \ \delta^{'} + \sum_{k=1}^{14} \omega_k h_{kt} \tag{4}$$

The history bias $\delta_{\text{hist}}\left(\mathbf{h_t}\right) = \sum_{k=1}^{14} \omega_k h_{kt}$ consisted of the sum of the preceding seven responses r_{t-1} to r_{t-7} and the preceding seven stimulus identities (i.e., motion directions) z_{t-1} to z_{t-7}, each multiplied with a weighting factor ω_k. The vector $\mathbf{h}_t$ was written as:

$$\mathbf{h}_t = \left(r_{t-1}, \ r_{t-2}, \ r_{t-3}, \ r_{t-4}, \ r_{t-5}, \ r_{t-6}, \ r_{t-7} \ , \ z_{t-1}, \ z_{t-2}, \ z_{t-3}, \ z_{t-4}, \ z_{t-5}, \ z_{t-6}, \ z_{t-7}\right)$$

All terms in $\mathbf{h}_t$were coded as -1 or 1, with the exception of terms coding for stimuli with zero coherence, which were set to zero. The weighting factors ω_k thus modeled the influence of each of the seven preceding responses and stimulus identities on the current choice. All weights were estimated with the same procedure as described previously (Fründ et al. 2014; Urai et al. 2017). Specifically, every set of seven past trials was convolved with three orthogonalised, exponentially decaying basis functions (with decay constants 0, 0.5 and 0.25), reducing the number of estimated parameters from 7 to 3 for each history feature. The estimated coefficients were then multiplied with their basis functions, resulting in the full 7-lag kernels shown.

Positive values of ω_k described a bias to repeat the preceding response or stimulus identity at the corresponding lag, and negative values of ω_k described a bias to alternate the preceding response or stimulus identity. Response and stimulus weights added up to the weights for preceding correct responses and the difference between response and stimulus weights resulted in the weights for preceding incorrect responses.

Psychometric function fit with history contributions for stimulus, choice, and motor response

In Experiment 1, perceptual choices and motor responses were decoupled through a mapping that varied from trial to trial. Thus, we could independently estimate the relative contribution of previous choices and

motor responses to the current choice bias. We added the last seven choices c_{t-1}, c_{t-2}, c_{t-3}, c_{t-4}, c_{t-5}, c_{t-6}, c_{t-7}, each one multiplied with a separate set of history weights ω'_k, to the history bias term $\delta_{hist}(\mathbf{h_t}, \mathbf{c_t})$.

$$\delta(\mathbf{h_t}, \mathbf{c_t}) = \delta' + \delta_{hist}(\mathbf{h_t}, \mathbf{c_t}) = \delta' + \sum_{k=1}^{14} \omega_k h_{kt} + \sum_{k=1}^{7} \omega'_k c_{kt} \quad (5)$$

We performed this analysis in two stages. First, we fitted the psychometric functions separately for each of the six different motion-viewing durations (on the current trial) and compared the resulting weights within each observer. The viewing duration had only negligible impact on the history weights (data not shown), indicating that the history contributions were invariant across viewing durations. Second, we fitted the psychometric functions to the data from all trials irrespective of viewing duration. Group level statistical tests were computed across the six observers, based on the mean weights irrespective of viewing durations. The corresponding weights are presented in Results.

Psychometric function fit with modulation of sequential bias by reaction time

In order to investigate the impact of decision confidence on sequential bias, we added a modulation by reaction time (Urai et al. 2017). Specifically, we extended the model from equation 3 by a term describing the interaction between previous choices and stimuli with previous reaction times $\sum_{k=1}^{14} \omega'_k h_t rt_{kt}$ and a nuisance covariate $\sum_{k=1}^{7} \omega''_k rt_{kt}$:

$$\delta(\mathbf{h_t}, rt_t) = \delta' + \delta_{hist}(\mathbf{h_t}, rt_t) = \delta' + \sum_{k=1}^{14} \omega_k h_{kt} + \omega'_k h_{kt} rt_{kt} + \sum_{k=1}^{7} \omega''_k rt_{kt} \quad (6)$$

with $rt_{kt} = (rt_{t-1}, rt_{t-2}, rt_{t-3}, rt_{t-4}, rt_{t-5}, rt_{t-6}, rt_{t-7})$ the vector of the log-transformed and then z-scored (within participant), reaction times of the preceding seven trials.

In all cases, the parameters α, γ, λ, δ', ν and the history weights ω_k, ω'_k and ω''_k were determined by maximizing the log-likelihood $L = \sum_t \log P(r_t = 1 | \tilde{s_t}, \mathbf{h_t})$ using an expectation maximization algorithm (Fründ et al. 2014).

Statistical analysis

We used t-tests and permutation tests (Efron and Tibshirani 1986) for all statistical comparisons reported in this paper. In particular, we used the following procedure to test the statistical significance of the history weights as function of previous trials (lags). First, a simple t-test was used to determine if the weights at each lag were significantly different from zero and a paired t-test was used to determine if choice and response weights (Experiment 1), or weights for the two biased conditions (Experiment 2), respectively, were significantly different. Second, false discovery rate correction (Benjamini and Hochberg 1995) was applied to correct for multiple comparisons across the 7 lags.

The following comparisons were performed using a paired permutation test (N = 10,000 permutations): (i) individual vectors (i.e., vector lengths, and orientations) pointing from the Neutral to each of the two

biased conditions; (ii) correlations between individual 'stimulus kernels' from each of the biased conditions and the two 'history templates' (see below); (iii) previous correct versus error weights; and (iv) previous correct versus previous RT*correct weights.

The following correlations were determined using Pearson correlation: (i) correlation between individual previous choice weights and performance; (ii) correlation between the correlation of stimulus kernels with history template and performance.

Lastly, we used Rayleigh's test to assess the clustering of angles in 'adaptation space' (Figure 3B) in both the Repetitive and Alternating conditions, and a Watson-Williams test to assess the difference in mean directions of adaptation between these two conditions (Berens 2009).

Results

We here use the term 'sequential choice bias' to describe the tendency to repeat, or alternate, previous choices more often than expected by chance. Figure 1C and D illustrate this for example observers from Experiment 2. Figure 1C shows, for one observer, a 'streak' of eight repeats of the same choice, followed by five repeats of the alternative. These streaks occur in the face of trial-to-trial variations of the direction and strength of the random dot motion stimuli. Figure 1D summarizes local choice biases in terms of the psychometric function conditioned on previous choice. The corresponding shifts between the two functions distinguish observers with repetition bias (left panel) and no bias (right). These shifts are only evident under consideration of the sequence of previous choices. Thus, they are distinct from the overall (frequency) bias towards one choice option, irrespective of experimental history (Bogacz et al. 2006; de Lange et al. 2013). The latter are also evident in these observers as leftward (Figure 1D, left) shifts of both psychometric functions.

The current paper pursued two objectives. First, we aimed to disentangle and compare the contribution of decisional and motor signals of sequential biases in perceptual choice. In laboratory tasks, perceptual choice and motor response used for reporting the choice are typically coupled, but can be decoupled with little effect on performance on the current trial (Tsetsos et al. 2015). While there is evidence for either decisional or motor origin of sequential biases (Akaishi et al. 2014; Pape and Siegel 2016; St. John-Saaltink et al. 2016), their relative contributions have not been quantitatively compared. Second, we aimed to quantify the adaptation of sequential biases to the environment, as a function of varying levels of internal decision confidence, in the absence of external feedback.

Disentangling the impact of previous stimuli, choices, and motor responses

In order to quantitatively compare the relative contribution of previous choices or motor responses on current choice, we reanalyzed data from a previously published study (Tsetsos et al. 2015). In this experiment (Experiment 1), observers performed a random dot motion task under trial-to-trial variations in the mapping between choice and motor response. The direction of motion was chosen randomly and independently

on each trial, so that maximizing performance required basing choices solely on the current stimulus and not on its history (i.e., previous stimuli, choices, or motor responses). We used a logistic regression model (Materials and Methods) to quantify the contribution of those three history effects, as well as of the current sensory evidence, on the current choices.

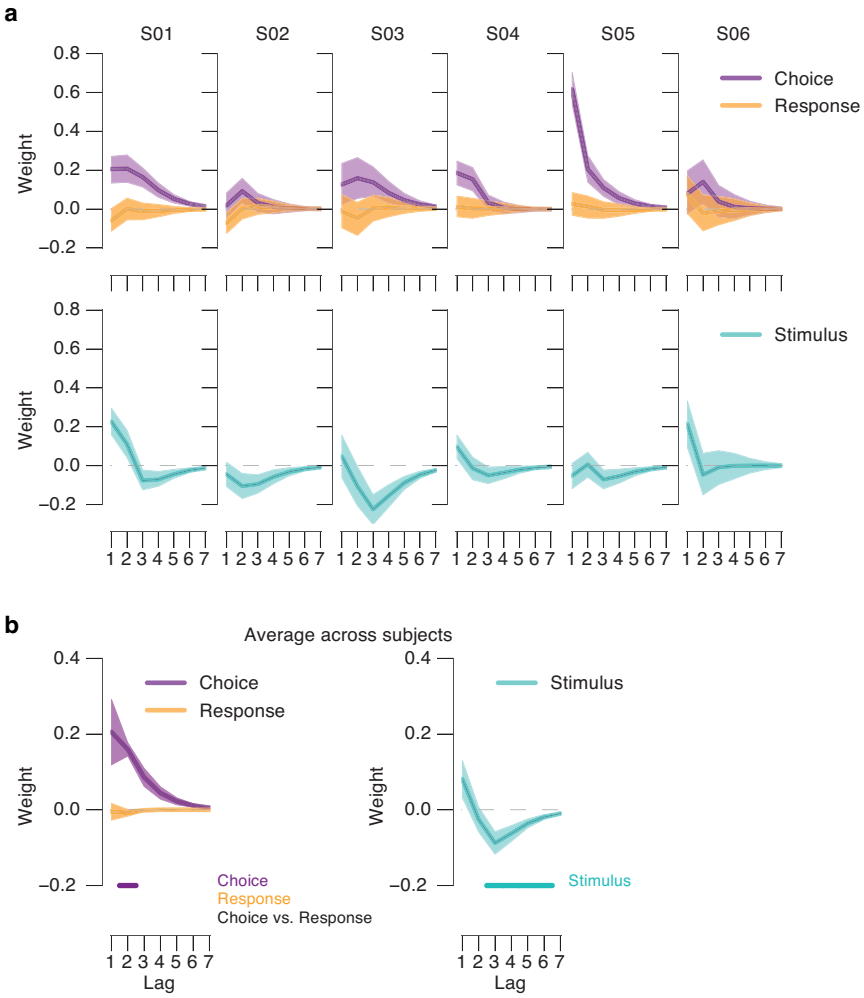

Figure 2. Stronger impact of previous choice than of previous motor response, on current bias Contributions of previous choices, motor responses and stimuli in Experiment 1. (**a**) Mean history contributions as a function of lags. Choice weights (purple), response weights (orange), and stimulus weights (blue; see Materials and Methods). Shading, 95% bootstrap confidence intervals. (**b**) As (**a**), but averaged across participants. S.e.m. and t-tests are computed across participants, $p < 0.05$ (FDR-corrected t-test) of choice, response and stimulus weights and of the difference between choice and response weights.

Observers showed a significant tendency to repeat their previous choices (indicated by positive choice weights), but not their motor responses (Figure 2A, upper row). In all individual participants, the effect of the previous choice on current choice was positive and stronger than the effect of the previous motor response. The same was true, and statistically significant at lag 2, in the group average (Figure 2B, left column). By contrast, preceding stimuli had negative, and statistically significant, weights at longer lags in the group average (Figure 2B, right; compare panel A, bottom row, for individual observers). This indicates

a long-lasting repulsion of choices from the previous stimulus across a time range of about 35 seconds.

The analyses presented here collapsed across two different conditions, in which observers were instructed about the required mapping between choice and response either before or after the presentation of the sensory evidence. We also analyzed the data from the two conditions separately and found no difference between them in the group analysis (data not shown). Thus, the predominance of previous choices over motor responses did not depend on observers' knowledge about the required mapping between choice and motor response during decision formation.

In sum, previous choices had robust and distinct (in terms of sign and time course) effects on current choice, while motor responses had only a minor impact. This indicates that the commonly observed choice repetition biases are specifically due to previous choices and not the motor responses used to the report them, which has implications for their neural sources (see Discussion). We next investigated the adaptability of sequential choice biases under fixed mapping.

Confidence-dependent adaptation of sequential choice biases to the environment

In laboratory tasks used to study perceptual choice, it is common to generate random sequences of the two alternative stimulus categories. But the states of natural environments, and hence the sensory signals generated by them, often exhibit correlations across time, so that it might be beneficial for decision-makes to adapt their sequential choice biases to that correlation structure (Yu and Cohen 2008). In Experiment 2, we tested for such adaptation, by systematically manipulating the repetition probabilities between the two possible motion directions across three conditions blocked by experimental session: Repetitive, Alternating, and Neutral (two sessions per condition; see Materials and Methods). Importantly, observers received no external trial-by-trial feedback about the correctness of their choices, which enabled us to also study the impact of decision confidence on the adaptation of their sequential biases to changing environmental statistics.

Also in this experiment, participants exhibited 'intrinsic' sequential patterns in their choice behavior (i.e., in the absence of biased stimulus sequences) during the Neutral condition (Figure 3A). Choice and stimulus weights exhibited a similar temporal profile as the weights observed in previous experiments using standard choice tasks without manipulation of stimulus sequences (e.g., compare with Figure 2B and Urai et al. 2017 Figure 5C). Choice weights were significantly positive for lag 2 and 3. As in the previous experiment, there was a statistically significant negative impact of previous stimuli on current choice across lags 3-7, indicating long-lasting repulsion of choices from the previous stimulus across a time range of about 30 s.

Adaptation of sequential choice biases to environmental statistics

The manipulation of the sequential statistics of the environment had clear effects on observers' sequential biases (Figure 3B-E). We mapped the individual data into a 'strategy space', defined by previous stimulus weights on the y-axis and previous choice weights on the x-axis (Figure 3B). Absolute x-values larger than

absolute y-values indicated a strategy to repeat (in case of positive values) or alternate (for negative values) previous choices independent of their correctness. Whenever stimulus weights were absolutely larger than choice weights, the resulting strategy depended on the accuracy on the previous trial. Positive stimulus weights (upper triangle) referred to repeating choices that were consistent with previously seen stimuli (i.e. repeating previous correct choices and alternating previous incorrect choices), equivalent to a 'win-stay lose-switch strategy'. Conversely, values in the lower triangle indicated a 'win-switch lose-stay strategy'.

Adaptation of observers' choice patterns to biased stimulus sequences predicts that choice and stimulus weights shift in the corresponding direction. Specifically, we expected shifts towards the top right for the Repetitive condition, indicating a bias to repeat the previous choice and stimulus, and towards the bottom left for the Alternating condition, indicating a tendency to alternate previous choice and stimulus (Figure 3B; top right and bottom left corners).

In line with these predictions, participants' weights for both, previous choice and previous stimulus, changed between conditions (Figure 3B; compare dots of different colors). The mean across observers moved from close to zero bias in the Neutral condition (indicated with a red 'x' in Figure 3B) towards more choice (and stimulus) repetition in the Repetitive condition (green arrow in Figure 3B), and towards more choice (and stimulus) alternation in the Alternating condition (blue arrow in Figure 3B), respectively. These shifts in history bias with conditions were statistically significant across the group: The vector angles indicating the shift with respect to Neutral were significantly different from uniform ($p = 0.0002$ in Repetitive and $p = 0.0008$ in the Alternating condition, Rayleigh's test) and the mean angles of these vectors (i.e., direction of the shift) were significantly different between Repetitive and Alternating conditions ($p < 0.0001$, Watson-Williams test).

The changes of history biases between conditions were not simply 'inherited from' the correlations evident in the stimulus sequences. We simulated the performance of synthetic observers, which were constructed individually for each participant, such as to yield the same perceptual sensitivity as that participant (within each of the three conditions), but without any memory for the preceding stimulus and choice sequence. As expected, the choice and stimulus weights obtained for these models varied slightly between models and lags (due to statistical fluctuations for finite amounts of data), but consistently approached zero for increasing numbers of simulated choices (data not shown). Thus, the deviations from zero of the experimentally measured weights reflected an active adjustment in observers' sequential biases to the statistics of the environment.

Bias adaptation tracks environmental statistics and predicts performance

The statistical structure of both biased conditions implied characteristic, and distinct, time courses of the probability of stimulus repetition as a function of lag (Figure 3C). The repetition probability was most different from 0.5 for lag 1 and approached 0.5 for larger lags in both conditions. But the Repetitive condition exhibited a monotonic decay towards 0.5, whereas the Alternating condition exhibited a damped oscillation

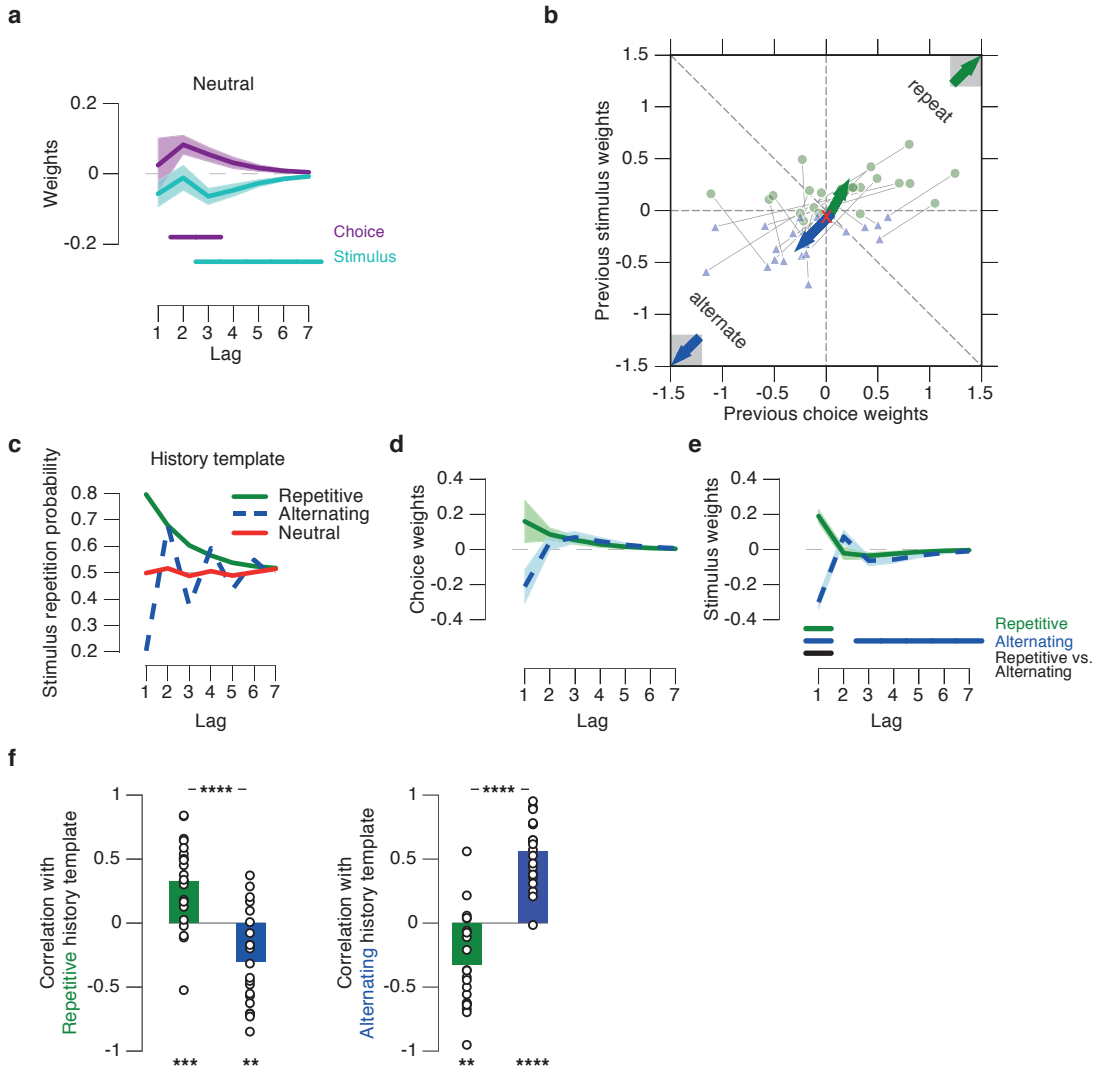

Figure 3. Adaptation of sequential choice biases to environmental statistics. (**a**) Stimulus and choice weights for the Neutral condition. (**b**) Impact of previous stimuli and choices on current choice in Experiment 2. Weight for immediately preceding trial (i.e. lag 1). Dots, single observers, arrows changes of group mean weights from Neutral (black crosses) during Repetitive and Alternating conditions, respectively. Green and blue arrows and underlying grey areas in upper right and lower left in the left panels indicate direction and range of expected effects (see main text for explanation). (**c**) Stimulus repetition probabilities for the Repetitive, Alternating and Neutral conditions. In the Repetitive condition, repetition of the motion direction from two trials back could occur due to a sequence of two repetitions or due to a sequence of two alternations resulting in a probability of 0.8 x 0.8 + 0.2 x 0.2 = 0.68, and so on for further lags. Likewise, for the Alternating condition the probability of repetition of the same direction oscillated around 0.5 as a function of lags, but again with decreasing deviations from 0.5. (**d**) Choice weights as functions of lags in Repetitive and Alternating conditions. (**e**). As (**d**), but for stimulus weights. Shaded areas, s.e.m.; bars, $p < 0.05$ (FDR-corrected t-test) across participants. (**f**) Correlation of stimulus kernels with history templates in the Repetitive (left) and the Alternating (right) condition.

around 0.5 (Figure 3C). In what follows, we refer to these two time courses characterizing the statistical structure of the environment as 'history templates'.

Adaptation to the structure of environment predicted that participants' history weights as a function of lag should exhibit similar profiles. We refer to the latter time courses as 'history kernels'. Indeed, the history

kernels for previous choices and stimuli showed similarly decaying profiles (Figure 3D, E). The adaptation effect was pronounced, and statistically significant, for the stimulus kernels for lag 1 and tracked the shape of the history templates beyond that lag (Figure 3E). We quantified the similarity between the entire history templates and participants' stimulus kernels by means of temporal correlation (Figure 3F). For the Alternating condition, the correlations should be regarded as lower bound of the true similarity between history kernels and history template because the estimation of the former through three independent basis functions precluded sign reversals from lag 4 onwards (see Materials and Methods). The stimulus kernels of both conditions, Repetitive and Alternating, were significantly correlated with their corresponding history template (kernel for Repetitive to template for Repetitive: $p = 0.0005$; kernel for Alternating to template for Alternating: $p < 0.0001$). Furthermore stimulus kernels of both conditions were more similar to their corresponding history template (i.e., kernel for Repetitive to template for Repetitive) than to the non-corresponding history template (Repetitive: $p < 0.0001$; Alternating: $p < 0.0001$; permutation tests; Figure 3F). Thus, participants adapted their history biases to the statistical structure of their environments with a time course matched to the time course of the task statistics.

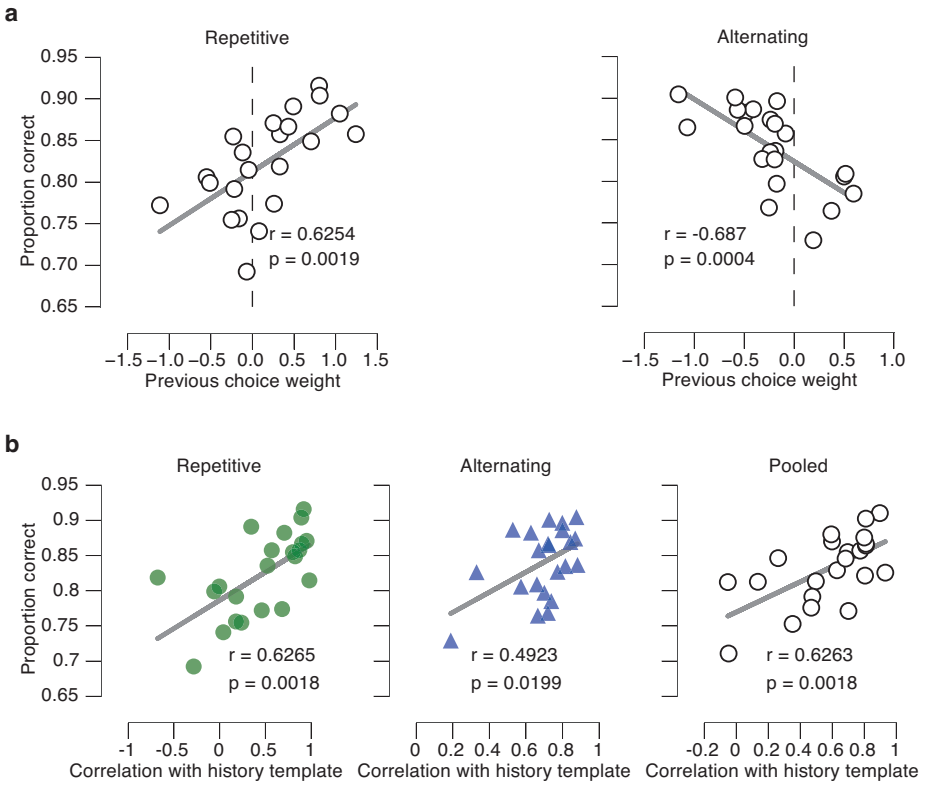

Figure 4. Correlation between adaptation and proportion of correct trials. Assessing the benefit of adaptation of sequential bias for overall performance. (**a**) Left, correlation between previous choice weight and the proportion of correct choices in the Repetitive condition. Right as left panel, but for Alternating condition. R and p: Pearson correlation coefficient and p-value. (**b**) Left: correlation between correlation of stimulus kernels with history template and the proportion of correct choices in the Repetitive condition. Middle: as left panel, but for Alternating condition. Right: pooled across Repetitive and Alternating conditions.

While the bias adaptation was highly consistent across participants, individuals differed in the extent to

which they shifted their history biases between conditions. These individual differences were strongly predictive of participants' overall task performance (Figure 4). Choice weights were significantly correlated with performance, for both the Repetitive (Figure 4A, left) and Alternating (Figure 4A, right) conditions, but with opposite sign, reflecting the opposite directions of the adaptation of choice weights in the two conditions. An analogous effect was evident for the similarity of their stimulus kernels to the history templates of the two conditions: Those participants with a larger similarity to the history templates also performed best, in both Repetitive and Alternating conditions (Figure 4B, left and middle) and averaged across both biased conditions (Figure 4B, right). Thus, the more strongly observers adapted their choice biases to the biased stimulus sequences, the more successful they were on the task. Our results supported the idea that participants accumulated signals from the past previous trials into biases for the current choice. This helped adapt their behavior to the statistics of their sensory environment, improving performance. Our interpretation is in line with current models of sequential effects, which postulate an active accumulation (with leaky memory) for previous stimuli and choices into their biases for upcoming choices (Yu and Cohen 2008; Glaze et al. 2015; Meyniel et al. 2016). Our final series of analyses, presented next, revealed that this accumulation was modulated by trial-to-trial variations in decision confidence.

Modulation of sequential bias adaptation by decision confidence

Previous accounts of sequential effects in behavior postulate the accumulation of binary external variables, such as stimulus repetitions (Yu and Cohen 2008; Meyniel et al. 2016), performance feedback (Abrahamyan et al. 2016), or reward (Sugrue et al. 2004). Our experimental conditions precluded any of the above: Observers performed under generally high uncertainty about the actual stimulus identities, and they did not receive external feedback about choice outcomes. We reasoned that, under these conditions, observers may have accumulated internal signals related to the decision process. Specifically, they might have accumulated the decision variable underlying choice (Glaze et al. 2015), neural representations of which not only encode the categorical choice, but also the graded confidence associated with that choice (Kepecs et al. 2008; Kiani and Shadlen 2009; Hebart et al. 2016). Consequently, the impact of previous choices on current bias should be modulated by the confidence associated with previous choices. We here use the term 'decision confidence' in a statistical sense, to refer to the observer's internal estimate of the posterior probability of having made a correct choice, given the evidence (Kepecs et al. 2008; Pouget et al. 2016; Sanders et al. 2016; Urai et al. 2017). This operational definition of confidence is agnostic about the link to subjective confidence reports (but see Sanders et al. 2016). It is formalized by a model based on signal detection theory (Kepecs et al. 2008; Sanders et al. 2016; Urai et al. 2017), the key features of which are reproduced in Figure 5A. We here used two proxies for decision confidence that were consistent with this model: accuracy (Figure 5B) and reaction time (Figure 5D).

The model shows that correct choices are overall associated with higher confidence (Figure 5A, top). To assess the effects of the correctness of previous choices on current bias, we recombined the response

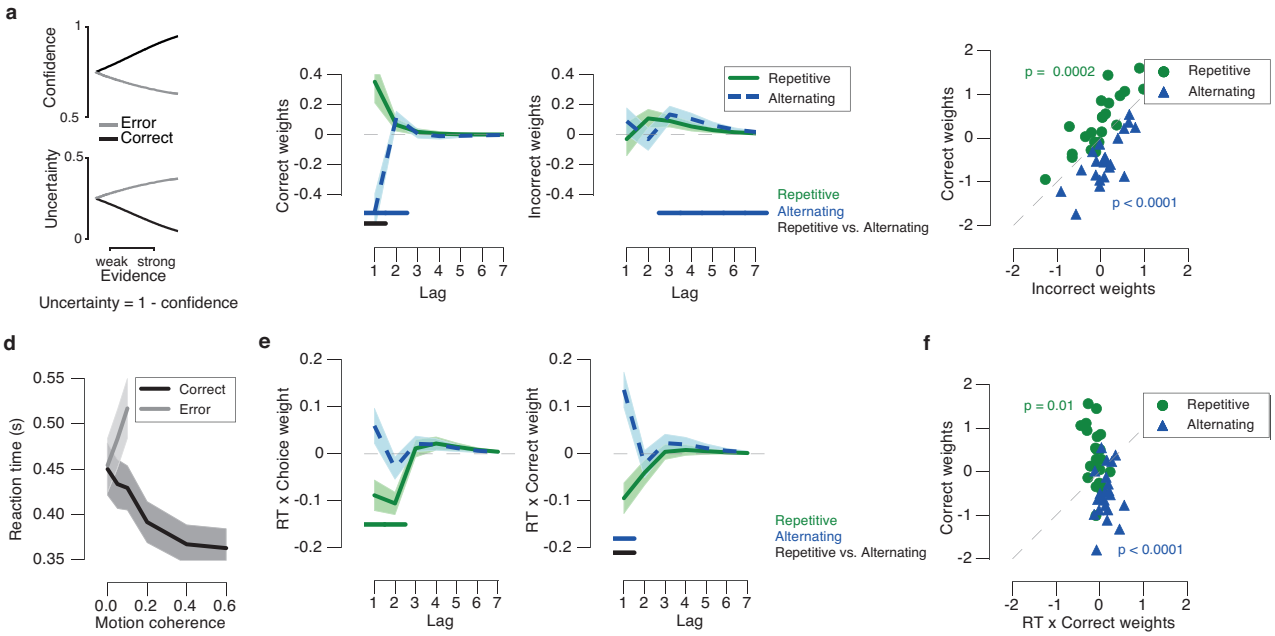

Figure 5. Modulation of adaptation by decision confidence. (**a**) Scaling of a model-based confidence and uncertainty with evidence strength in correct and error trials. Adapted from Urai et al. 2017 under a CC-BY 4.0 license. (**b**) Left, weights of correct preceding choices as function of lags. Right, weights of incorrect choices as function of lags. Shaded areas, s.e.m.; bars, p < 0.05 (FDR-corrected t-test). (**c**) comparison between correct and incorrect weights for previous trial (lag 1). (**d**) Reaction time as a function of motion coherence, sorted by correctness, collapsed across Repetitive and Alternating condition. (**e**) Left, multiplicative modulation of the weights of previous choice by reaction time (RT). Right, modulation of the weights of previous correct choices by RT. (**f**) Comparison between previous correct and previous RT x correct weights.

and stimulus weights for the previous trials into weights for correct and incorrect choices (Materials and Methods). As expected, the impact of previous choices on current bias was larger when these choices were correct (Figure 5B). In both biased conditions, participants' weights for previous correct choices deviated from zero at lag 1 and then decayed for further lags (Figure 5B, left), an effect not evident for previous incorrect choices (Figure 5B, right). At lag 1, the weights were significantly larger for correct than incorrect previous choices in the Repetitive condition, and the other way around for the Alternating condition (Figure 5C). This pattern of results was consistent with the idea that the accumulation of evidence across choices was boosted by decision confidence. In the present data, the scaling of reaction times with the absolute strength of sensory evidence (i.e., motion coherence) for correct and incorrect choices also replicated a characteristic signature of decision uncertainty (i.e., the complement of confidence) from previous studies (Sanders et al. 2016; Urai et al. 2017): reaction times (i.e., uncertainty) decreased with evidence strength for correct choices, but increased for error choices (Figure 5D, compare to Figure 5A). Correspondingly, linear regression revealed opposite-signed relationship between motion coherence and reaction times, separately for correct (β = -0.150, s.e.m. = 0.027, p = 0.005) and error trials (β = 0.628, s.e.m. = 0.025, p = 0.025). This pattern of evidence- and accuracy-dependent changes in reaction time is consistent with the model of decision confidence (Kepecs et al. 2008), as well as with previous psychophysical data (Sanders et al. 2016; Urai et al. 2017). To assess the modulatory effect of reaction times on sequential choice bias, we built on an

extension of our statistical model by multiplicative (interaction) terms, which quantified the degree to which the impact of previous choices on current bias was modulated by previous reaction times (see *Psychometric function fit with modulation of sequential bias by reaction time* in Materials and Methods for details). Longer reaction times (i.e. lower decision confidence) reduced the impact of the previous choices on the current bias (Figure 5E). Specifically, the interaction weights in the Repetitive condition were significantly negative at lag 1 and 2, indicating a confidence-linked reduction of the choice repetition bias in that condition (Figure 5E, left). Conversely, in the Alternating condition the interaction weight for the previous choice was positive, indicating a (albeit not significant) reduction of the choice alternation bias (Figure 5E, left). This modulatory effect of reaction time on the impact of previous choices on current bias was also evident when assessed for previous correct choices in isolation: here the previous interaction weight was statistically significant (Figure 5E, right), for the Alternating condition and significantly larger than the corresponding interaction weight from the Repetitive condition. As expected for a negative modulatory effect of previous decision uncertainty on current choice bias, most weights of the interaction terms (RT x correct) had opposite sign to those of the weights for the corresponding (i.e. correct) previous choices (compare Figure 5E, right and 5B, left). At lag 1, the interaction weights were consistently (and statistically significantly) smaller than the main effect weights for (correct) choices in the Repetitive condition, and the other way around in the Alternating condition (Figure 5F). Thus, even for correct choices only, the varying uncertainty (indexed by reaction time) associated with these choices counteracted their impact on current choice. Taken together, two independent proxies of decision confidence, choice accuracy and reaction time, both supported the conclusion that decision confidence modulated (boosted) the adaption of sequential choice bias to the statistical structure of the environment.

Discussion

Sequential biases are a long-known, pervasive phenomenon in perceptual choice (Fernberger 1920; Fründ et al. 2014), which likely reflect the interplay of multiple factors. Here, we systematically quantified the contributions of a number of important candidate factors: preceding choices versus motor responses, the sequential statistics of the sensory evidence, and decision confidence in the absence of external feedback. We showed that the contribution of previous choices to sequential biases dominated over the contribution of previous motor responses. We further found that, in the absence of explicit feedback about choice outcome, observers could still adapt their sequential choice biases to biased (repetitive or alternating) sequences of sensory evidence. The strength of this adaptation was modulated by observers' confidence about their previous choices (with stronger adaptation after more certain choices), as indicated by previous choice accuracy and reaction time. Finally, we established that the adaptation of sequential bias was beneficial, improving the observer's overall performance. Our report provides a comprehensive characterization of these various different history contributions, within a common statistical modeling framework. Further, our current approach also identified novel effects, in particular the confidence dependent adaptation of se-

quential choice biases. This yielded a number of new insights into the computations underlying sequential biases in perceptual choice, with direct implications for their neural bases.

Multiple factors contributing to history biases in perceptual choice

Our analyses revealed that the contributions of previous stimuli, perceptual choices, and motor responses were dissociable in terms of their strength, sign, and time course (as function of previous trials). For two experiments, we found that previous stimuli exerted weak, but consistently repulsive biases on current choice, across a long timescale (significant several tens of seconds into the past). The fact that these long-lasting, repulsive effects of previous stimuli on current choice were present in both, Experiment 1 and 2, indicate that they are irrespective of the presence of trial-by-trial feedback. Analogous results have been obtained fine discrimination tasks (Fischer and Whitney 2014; Fritsche et al. 2017). These repulsive effects may reflect the impact of bottom-up sensory adaptation mechanisms (Kohn 2007) on sensory-guided behavior.

The dominant and consistent bias in uncorrelated stimulus sequences (irrespective of the presence of trial-by-trial feedback, compare Figures 2 and 3A) was to repeat the preceding one or two choices. Similar observations have been made in a number of previous studies of perceptual choice (de Lange et al. 2013; Akaishi et al. 2014; St. John-Saaltink et al. 2016; Fritsche et al. 2017; Urai et al. 2017). Importantly, we here established that this repetition bias was due to the previous perceptual choices, rather than motor responses.

Two recent studies similarly decoupled perceptual choice and motor response (Akaishi et al. 2014; Pape and Siegel 2016). One of them showed that a bias to alternate response hands from trial to trial systematically contributed to sequential effects, due to activity dynamics within motor cortex (Pape and Siegel 2016). Effector-selective beta-band activity in motor cortex is suppressed during decision formation (Donner et al. 2009; O'Connell et al. 2012; de Lange et al. 2013), followed by a rebound. The stronger this 'beta-rebound', the stronger the tendency to alternate response hands (Pape and Siegel 2016). This motor response alternation bias was superimposed onto a choice repetition bias in their study, but they did not compare the magnitude and time course of these two effects directly. Here, we performed such a direct comparison across seven previous trials. While the lack of a robust response alternation bias in our sample does not rule out the alternation effect, our comparison revealed that the contribution of previous choices to current bias was more consistent across observers, significantly stronger, and more prolonged in time (most robust for a lag of two trials).

This predominance of the impact previous choices over previous motor responses is consistent with the results from Akaishi et al. 2014, who also focused on the contribution of the immediately preceding trial. This predominance implies that sequential biases in perceptual decision-making are governed by neural signals encoding choice in an action-independent format. Such signals exist in associative brain regions, such as posterior parietal or prefrontal cortex (Bennur and Gold 2011; Hebart and Hesselmann 2012; Hebart et al. 2016), which also exhibit the short-term memory dynamics necessary for the persistence of biases

in the decision-making machinery (Wang 2002; Bonaiuto et al. 2016; Morcos and Harvey 2016). Under the assumption that the same principle holds when the mapping between choices and motor responses is fixed, the adaptation of sequential biases to the different environmental statistics observed in our second experiment were also dominated by action-independent decision variables, rather than by motor signals.

Previous evidence for adaptable history biases

Another recent study also revealed adaptable history biases in perceptual choice (Abrahamyan et al. 2016). However, different from us, Abrahamyan et al. provided unambiguous external feedback about the outcome of each choice, and their manipulation of the stimulus sequence depended on the participants' success or failure. Consequently, the process adapting history biases likely depended on the combination of choices and their outcome. By contrast, in our task, participants remained uncertain about their choice outcomes, and could only base their history biases on internal signals (see below). Thus, the adaptation effects in their and our study likely resulted from different mechanisms, and our current results complement those from Abrahamyan et al. Indeed, a direct comparison of sequential effects in saccadic choice tasks with and without feedback points to an important influence of feedback on the profile of stimulus-dependent history biases (Kim et al. 2017).

Computational basis of adaptable sequential biases

The main novel contribution of the present study is the identification of a confidence-dependent adaptation of sequential choice biases. This adaptation likely reflected the ongoing accumulation of internal (likely action-independent) decision variables across trials, in line with recent computational models (Glaze et al. 2015; Bonaiuto et al. 2016). This idea is related to earlier proposals on the origins of sequential effects in cognition and choice behavior (Yu and Cohen 2008; Meyniel et al. 2016), which however postulate the accumulation of external observables rather than internal signals. Integration of information over time is a widely established concept in current models of perceptual choice (Bogacz et al. 2006; Gold and Shadlen 2007; Ratcliff and McKoon 2008; Wang 2008; Ossmy et al. 2013), but, in these models, refers to within-trial accumulation of incoming sensory evidence across hundreds of milliseconds. The notion of accumulation of rewards across trials is widely established in the theory of reinforcement learning and value-guided choice, but there it critically depends on the experience of choice outcomes (i.e., rewards) and spans time scales of tens of seconds or even longer (Sutton and Barto 1998; Sugrue et al. 2004; Glimcher 2011).

Our interpretation assumes an accumulation mechanism with timescales situated in between these two classes of models, which operates on internal decision variables that may, in turn, be computed through within-trial accumulation of sensory evidence. This 'higher-order' across-trial accumulation has been described by recent belief updating models for perceptual choice in changing environments (Glaze et al. 2015; Kim et al. 2017). It will be important to explore such mechanisms with biophysical models of inter-choice

neural dynamics (Bonaiuto et al. 2016), as well as recordings of neural activity during similar tasks as the ones used here.

The framework outlined above also provides a natural interpretation of the modulatory effect of decision confidence on the adaptation of sequential choice biases that we observed. Representations of decision variables in the brain do not only reflect categorical choices, but also the graded certainty about these choices (Kepecs et al. 2008; Kiani and Shadlen 2009; Hebart et al. 2016). In the face of noisy sensory evidence and in the absence of single-trial feedback, such internal decision variables are a good proxy for the state of the environment available to the decision-maker. Uncertain choices are based on a smaller decision variable, they are likely to be incorrect and accompanied by slow responses (Sanders et al. 2016). Consequently, accumulation of internal decision variables predicts that incorrect choices, or choices with long reaction time, have little impact on sequential choice bias. Accumulation of confidence-dependent decision variables into biases for upcoming choice may also account for the confidence-dependent (as measured by reaction times) modulation of 'intrinsic' sequential choice biases measured under conditions of random stimulus sequences (Urai et al. 2017). In sum, one possibility is that the 'confidence-weighting' of the impact of previous choices on current bias observed here and in Urai et al. 2017 may directly result from the accumulation of graded decision variables.

Conclusion

We conclude that, in the absence of certainty about stimulus identities and choice outcome, the brain accumulates confidence-dependent (and likely action-independent) decision variables towards choice biases in a way that promotes adaptation to correlated environments. Our findings also put strong constraints on the candidate neural sources of sequential choice biases, setting the stage for future neurophysiological work into the underlying mechanisms.

Acknowledgements

We thank Ainhoa Hermoso, Alexandre Hyafil, Jaime de la Rocha, and Florent Meyniel for comments on the manuscript. This research was supported by the German Academic Exchange Service (DAAD, to A.E.U.) and the German Research Foundation (DFG), SFB 936/A7, SFB 936/Z1, DO 1240/2-1 and DO 1240/3-1 (to T.H.D.).

2.3 | Choices bias the rate of evidence accumulation in the next trial

Urai AE, Donner TH.

Abstract

Perceptual decision-making under uncertainty is biased by previous choices, a phenomenon referred to as serial choice bias. We investigated how this bias affects decision dynamics, as described by the accumulation of noisy sensory evidence towards opposing bounds for each choice option. We used drift diffusion modeling to disentangle two possible biasing mechanisms: an offset of the accumulation starting point, versus a bias in the drift (accumulation rate) towards one of the choice options. Across six different psychophysical datasets, individual serial choice biases were captured by a history-dependent change in drift bias. Additionally, when stimulus sequences were biased towards repetition or alternation, observers' adapted serial choice patterns also correlated to history-dependent changes in starting point. We conclude that the history of choices primarily biases the rate of evidence accumulation towards a particular choice option.

Author contributions: A.E.U. and T.H.D. designed research; A.E.U. collected and analyzed the data; A.E.U. wrote the manuscript; T.H.D. supervised research.

Introduction

Serial patterns in choice sequences have been described in the psychophysical literature for almost a century (Fernberger 1920), and have recently seen a renewed surge of interest in neuroscience. Both humans (Fründ et al. 2014) and mice (Busse et al. 2011) show strong and variable serial biases when stimulus sequences are uncorrelated, as is the case in most psychophysical experiments. These serial choice patterns are most strongly driven by previous decisions, rather than the associated motor response or sensory stimulus (Braun et al. 2017; Akaishi et al. 2014; St. John-Saaltink et al. 2016; but see Pape and Siegel 2016). They may result from the accumulation of decision-relevant evidence across trials, which improves performance when stimulus sequences are structured (Yu and Cohen 2008; Glaze et al. 2015; Braun et al. 2017; Kim et al. 2017). Indeed, when presented with structured stimulus sequences, humans can adapt their serial choice biases (Abrahamyan et al. 2016; Braun et al. 2017; Kim et al. 2017). Serial choice biases are modulated by decision confidence in both neutral (Urai et al. 2017) and biased environments (Braun et al. 2017), suggesting that confidence-dependent accumulation of information across trials gives rise to sequential choice patterns.

To quantify serial choice biases, most recent studies have used observers' choices to compute choice repetition probabilities or shifts in decision criterion (de Lange et al. 2013; Akaishi et al. 2014) or fit history-dependent psychometric functions (Busse et al. 2011; Fründ et al. 2014; Abrahamyan et al. 2016; Bonaiuto et al. 2016; St. John-Saaltink et al. 2016; Braun et al. 2017; Urai et al. 2017).

One crucial factor not captured by this approach is the time needed to choose. Extending signal detection theory to include a temporal dimension, sequential sampling models posit the integration of evidence as unfolding over time. The drift diffusion model (DDM) is a popular variant of this class of models (Bogacz et al. 2006; Ratcliff and McKoon 2008). The DDM posits that external sensory evidence, corrupted by Gaussian noise, is accumulated over time. When the integrated *decision variable* reaches one of two pre-set bounds, corresponding to the two choice options, a response is triggered (Figure 1, black lines and Methods). Each reaction time thus arises from the dynamics of this accumulation process, as well as time needed to encode the stimulus and execute the motor response.

Importantly, the DDM posits that choice bias can have two sources: (i) a change in the starting point of accumulation (or, equivalently, asymmetric decision bounds), and (ii) a bias in the rate with which evidence for one or the other choice is accumulated. The latter bias, also termed 'drift criterion', is an evidence-independent constant added to the decision variable per unit time. Throughout this paper, we will refer to this bias as 'drift bias'.

The difference between the two biasing mechanisms, and their effect on RT distributions, is shown in Figure 1. A shift in starting point mostly affects the leading edge of the RT distribution, shifting its mode. The effect of a biased drift rate instead grows with time, and affects the trailing edge of the distribution without changing its mode (Ratcliff and McKoon 2008; Leite and Ratcliff 2011; White and Poldrack 2014). Importantly, these two mechanisms can give rise to the exact same choice frequency (and hence the same choice bias). Previous work thus leaves open the question which of these two mechanisms better explains

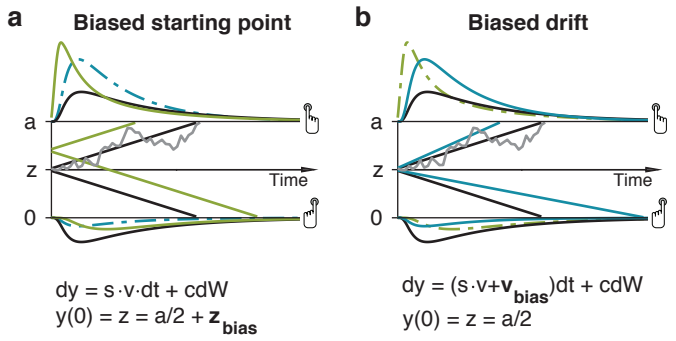

Figure 1. Bias in the DDM. The effect of biased (**a**) starting point or (**b**) drift on choices and RT distributions. Noisy sensory evidence is accumulated over time, until the integrated decision variable reaches one of two thresholds, triggering a response. Note that the two bounds 0 and a correspond to the two choice options, rather than correct and error responses. Repeating this process with identical stimuli over many trials gives distributions of reaction times for the two choice options. From the shape of these RT distributions, the underlying parameters of the decision process can be inferred. Specifically, a bias in drift towards one option (blue) changes the trailing edge, but not the mode, of the RT distributions. A bias in starting point (green), on the other hand, affects the leading edge and the mode of the RT distributions. The RT distributions were simulated according to the equations described by Ratcliff and Tuerlinckx 2002, using code provided by Murphy et al. 2016a.

serial choice behaviour.

Here, we fit the DDM to six psychophysical datasets, conditioning both drift bias and starting point bias on the previous choice. We find that previous choices affect both starting point and drift bias, with history-dependent drift bias best able to capture individual serial choice patterns.

Materials and Methods

Bias in the DDM

Unbiased accumulation of evidence towards one of two choice bounds (black lines in Figure 1) is described by a drift diffusion process

$$dy = s \cdot v \cdot dt + c dW, \quad y(0) = z = a/2 \tag{1}$$

where s is the signed stimulus identity $[-1, 1]$, v is the drift rate, and $c dW$ is Gaussian distributed white noise with mean 0 and variance $c^2 dt$ (Bogacz et al. 2006). Accumulation starts at the midpoint z, which in a unbiased case is situated midway between the two decision bounds 0 and a.

Bias towards one of the two choice options can be implemented in this model in two ways. First, the starting point of accumulation can be offset from its unbiased midpoint between the two bounds (Figure 1a):

$$dy = s \cdot v \cdot dt + c dW, \quad y(0) = z = a/2 + z_{bias} \tag{2}$$

where z_{bias} reflects the biased starting point (Link and Heath 1975; Ratcliff and McKoon 2008).

Second, the accumulation of evidence per unit time can be biased towards one of the two choices (Figure 1b):

$$dy = (s \cdot v + v_{bias})dt + c dW, \quad y(0) = z = a/2 \tag{3}$$

where v_{bias} reflects drift bias (Ashby 1983; Ratcliff and McKoon 2008). Importantly, these two biasing mechanisms can result in an identical pattern of biased choices. They are distinguished by their effects on the shape of reaction time distributions, with a nonzero z_{bias} affecting the leading edge and the mode, and v_{bias} affecting the tail of the distribution (Figure 1).

Tasks and datasets

2AFC RT

Twenty-two observers (12 men and 10 women, aged 18-25 years) performed five sessions of 500 trials (divided into ten blocks) of a classical up-down motion discrimination task (Figure 2a). Moving noise-dots were continually displayed on the screen, starting with a fixation period of 0.6-0.8 seconds. Stimulus onset (the signal dots starting to move up or down) was indicated with a short beep. The coherently moving dots were on the screen until the subject pressed a button, or after 3 seconds. After a variable interval of 1.5-2.5s, a tone was played as feedback. After 2-2.5s, the dots stopped moving for an ISI of 1s, after which the next trial started.

Psychometric functions were used to determine individual perceptual thresholds before the start of the experiment. To obtain the psychometric function, 600 stimuli of different difficulties were randomly interleaved (0, 2.5, 5, 10, 20 and 40% motion coherence), and were shown for 1s. To determine each individual psychometric threshold, we fit a cumulative Weibull as a function of difficulty c, defined as

$$\phi(c) = \delta + (1 - \delta - \gamma)\left(1 - e^{(\frac{-c}{\alpha})^{\beta}}\right) \tag{4}$$

where δ is the guess rate (chance performance), γ is the lapse rate, and α and β are the threshold and slope of the psychometric Weibull function, respectively (Wichmann and Hill 2001a). While keeping the guess rate δ bound at 50% correct, we fit the parameters γ, α and β using a maximum likelihood procedure implemented by minimizing the logarithm of the likelihood function. This was done using a Nelder-Mead simplex optimization algorithm as implemented in Matlab's *fminsearch* function. The individual threshold was taken as the motion coherence corresponding to a 65% correct fit of the cumulative Weibull.

Dynamical random dot stimuli were presented in a central circle (outer radius 14°, inner radius 2°) around fixation. The annulus was defined by a field of dots with a density of 1.7 dots/degrees^2. Dots were 0.2° in diameter and were white, at 100% contrast from the black screen background. 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 replotted in a random location after being on the screen for four consecutive frames.

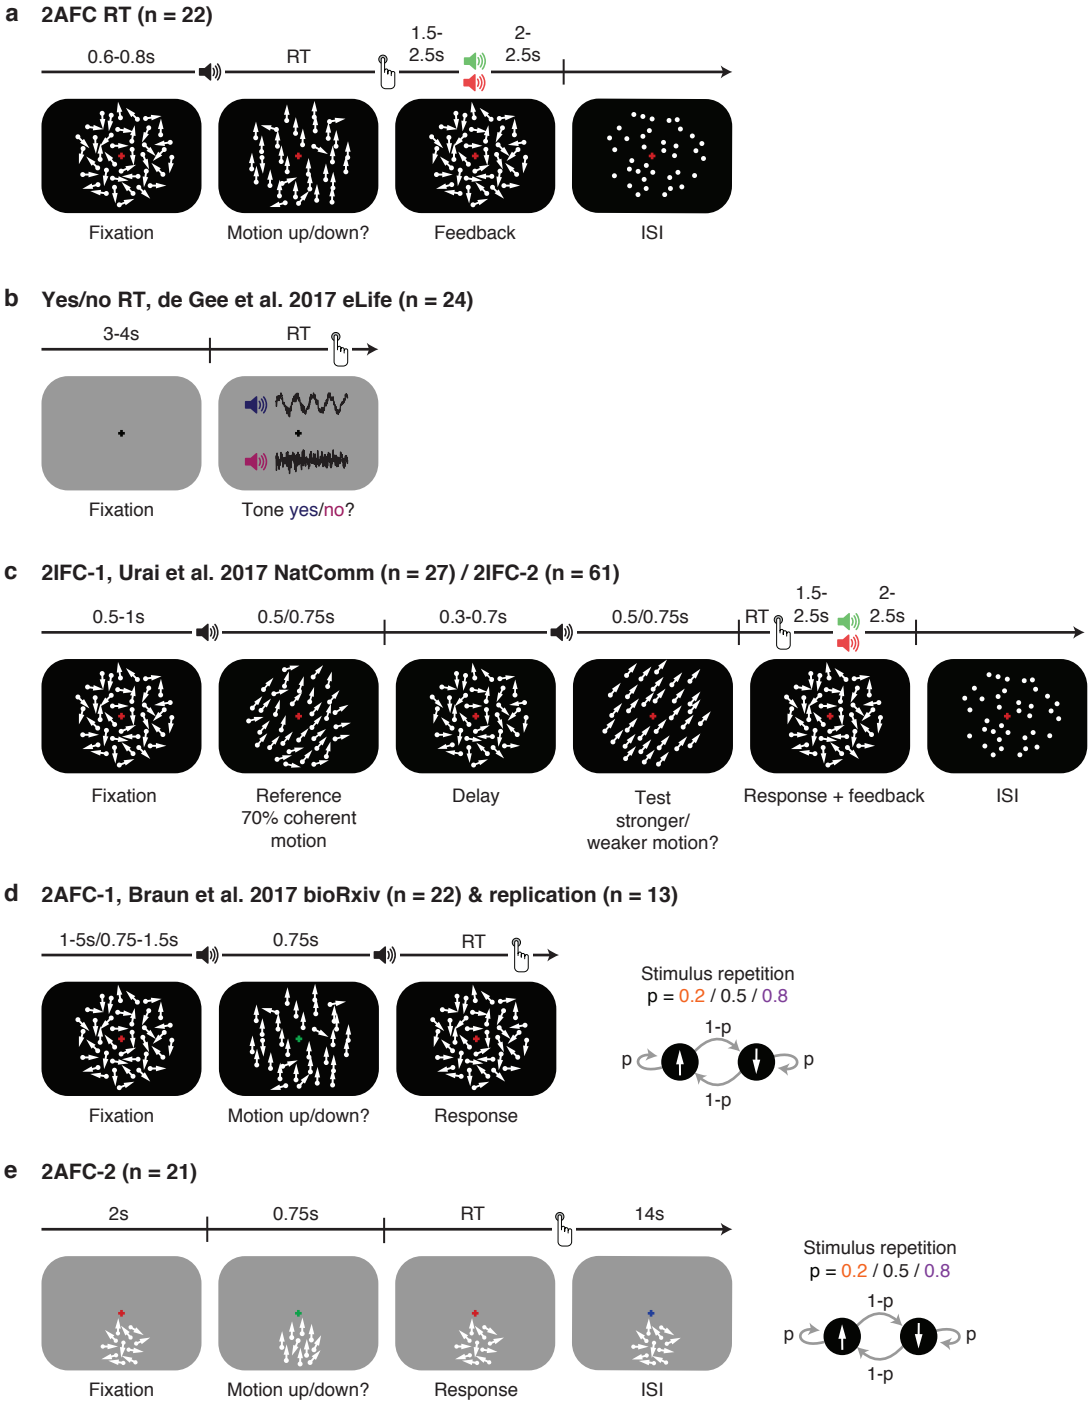

Figure 2. Tasks and datasets. (**a**) 2AFC (up/down) reaction time task. Motion coherence was constant, and titrated to 65% correct individual perceptual thresholds. (**b**) Yes/no auditory reaction time task (de Gee et al. 2017). Within a stream of auditory noise (TORCs), observers detected whether a pure sine wave (2 KHz) was present. Task difficulty was titrated using an adaptive staircase procedure, resulting in an average accuracy of 74%. (**c**) 2IFC-1 coherence discrimination, with variable evidence strength (Urai et al. 2017). In two consecutive intervals, dots moved in a constant direction with varying coherence. Motion coherence in the first interval was always 70%, and observers decided whether motion in the second interval was stronger or weaker than in the first (0.625-30% coherence difference). The same task was used in the 2IFC-2 replication dataset, but with a fixed level of difficulty titrated to individual thresholds. The timing of the task differed slightly between the two studies, as indicated in the figure and Methods section. (**d**) 2AFC-1 (up/down) task (Braun et al. 2017). Motion coherence varied from trial to trial (0-81%). Between sessions, transition probabilities between the two stimulus identities were 0.2, 0.5 or 0.8. In an additional replication, observers performed the task in the MEG scanner, where they additionally received auditory feedback. (**e**) 2AFC-2 (up/down) task. Motion coherence varied from trial to trial between two individually titrated difficulty levels (65% and 80%). Between experimental runs, transition probabilities between the two stimulus identities were 0.2, 0.5 or 0.8.

Additionally, to avoid subjects tracking individual signal dots as they move through the annulus, three independent motion sequences were interleaved on subsequent frames, making the effective motion speed in the display 3.8 11.5°/second. Noise dots were assigned a random location within the annulus on each frame, resulting in 'random position' noise with a 'different' rule (Scase et al. 1996). Throughout the experiment, a red 'bulls-eye' fixation target of 1.5° diameter (Thaler et al. 2013) was present in the centre of the screen. These stimulus characteristics were used in all experiments using random dot displays unless indicated otherwise.

Yes/no RT

These data were previously published by de Gee et al. 2017, and are available at `https://doi.org/10.6084/m9.figshare.4806562`. Twenty-four observers (20 women and 4 men, aged 19–23) performed an auditory tone detection task, in which they detected the presence of a pure sine wave (2 KHz) embedded in auditory noise (TORCs, McGinley et al. 2015b). After an inter-trial interval of 3-4 seconds, either noise or noise with a superimposed signal was played until the observer's button press or 2.5s. Each individual's signal volume was determined by an adaptive staircase procedure before the start of the experiment. Observers performed between 1320 and 1560 trials each, divided over two experimental sessions.

2IFC task

These data were previously reported in Urai et al. 2017, and are available at `http://dx.doi.org/10.6084/m9.figshare.4300043`. Twenty-seven observers performed approximately 2500 trials over five sessions of a two-interval random dot motion discrimination task. Observers saw two consecutive intervals of coherent dot motion (in a constant direction towards one of the four diagonals, counterbalanced across participants) and judged whether the second 'test' interval contained motion that was stronger or weaker than the first 'reference' interval. Stimulus difficulty, defined as the difference in motion coherence between the two intervals, varied from trial to trial (0.625, 1.25, 2.5, 5, 10, 20, 30%). They received auditory feedback on their choice after a variable delay.

Additionally, we collected a larger replication sample (n = 61, 19 men and 42 women, aged 19-35 years) with the same task. The experiment differed from the description in Urai et al. 2017 in the following ways. Observers performed 5 sessions, of which the first and the last took place in the MEG scanner (600 trials divided over 10 blocks per session) and the three sessions in between took place in a behavioural lab (1500 trials divided over 15 blocks per session). In an initial screening session, we determined each individual's motion discrimination threshold (at 70% of the cumulative Weibull as described above), which was then kept constant throughout the experiment. The duration of each interval was 750 ms. In the MEG sessions, auditory feedback was presented 1.5-3s after response, and an ISI with stationary dots started 2-3s after feedback. Participants initiated the next trial with a button press. In the training sessions, auditory feedback

was presented immediately after the response. This was followed by an ISI of 1s, after which the next trial started.

Participants in the MEG study were assigned to one of three pharmacological groups. At the start of each experimental session, they orally took 40 mg atomoxetine (Strattera®), 5 mg donepezil (Aricept®), or placebo. Since the time of peak plasma concentration is 3 hours for donepezil (Rogers and Friedhoff 1998) and 1-2 hours for atomoxetine (Sauer et al. 2005), we used a placebo-controlled, double-blind, double-dummy design, entailing an identical number of pills at the same times before every session for all participants. Participants in the donepezil group took 5 mg of donepezil 3 hours, and a placebo 1.5 hours before starting the experimental session. Participants in the atomoxetine group took a placebo 3 hours, and 40 mg of atomoxetine 1.5 hours before the experimental session. Those in the placebo group took identical-looking sugar capsules both 3 and 1.5 hours before starting the session. This ensured that either drug reached its peak plasma concentration at the start of the experimental training. The drug doses were based previous studies with healthy participants (Chamberlain et al. 2009; Rokem and Silver 2010). Blood pressure and heart rate were measured and registered before subjects took their first and second pill. In the three hours before any MEG or training session, participants waited in a quiet room. In total, 19 people in the placebo, 22 in the atomoxetine, and 20 in the donepezil group completed the full study. We did not observe differences in serial choice biases between these groups (Figure S1), and therefore pooled all observers for the analyses presented here.

2AFC-1 fixed duration task

These data were previously reported in Braun et al. 2017. Twenty-two observers (15 women and 7 men, aged 20 - 36) performed a fixed duration version of the random dot motion discrimination (up/down) task, with stimuli (0, 5, 10, 20, 40 or 60% motion coherence) displayed for 750 ms. Between sessions of the experiment, we varied statistics of the trial sequence by setting the stimulus transition probability to 0.2 (alternating), 0.5 (neutral) or 0.8 (repeating). Per condition, observers did 1200 trials. They did not receive trial-by-trial performance feedback.

Additionally, we added data from fifteen observers who performed the same task in the MEG scanner. This version of the experiment was different in the following ways. First, observers performed 1782 trials over 3 sessions, in which the stimulus transition probability varied between blocks of 99 trials. Motion coherence varied from trial to trial from (0, 3, 9, 27, 81%). Observers received auditory feedback 1.5-2.5s after their response, and the ISI started 2-2.5s after feedback. The stimuli were confined to a circle of 2.5° radius, which was placed in the lower half of the visual field at 3.5° from the fixation cross in order to best record responses from visual cortex.

In both versions of this experiment, observers were informed that the stimulus sequences could be "as produced by a coin flip", "more likely repeating than alternating", or "more likely alternating than repeating", but were not told the exact transition probabilities, the order in which these conditions were presented,

nor what to do with this information to perform the task (Braun et al. 2017).

2AFC-2 fixed duration task

Twenty-one observers performed an up/down motion direction discrimination task in the fMRI scanner. In total, each observer did 720 trials over 3 separate days. Coherence levels were titrated to individual 65% and 80% accuracy before the start of the experiment, using a method of constant stimuli and a Weibull fit (see above). After a fixation baseline of 2 s, an up/down random dot stimulus appeared for 750 ms. After stimulus offset, observers had 1250 ms to respond with their left or right hand (stimulus-response mapping counterbalanced across blocks within a session), followed by an ITI of 1.25 s - RT plus 14 s (7 TR). During the fixation baseline and ITI, 0% coherent random dot motion was continually present, while the fixation cross changed color to indicate the pre-stimulus baseline and response window (Figure 2e).

The dots moved within a circle of 2.5° radius, which was placed in the lower half of the visual field at 3.5° from the fixation cross. The dots where white on a grey background, and signal moved at 11.5 °/s, with a maximum lifetime of 10 frames. Additionally, three independent sequences of the moving dot stimuli were interleaved, making the effective speed of motion 3.8°/second. The cloud of dots had a density of 6 dots/degrees^2, with each dot 0.06° in radius.

Stimulus sequences were generated based on repetition probabilities [0.2, 0.5, 0.8]. Within each session of the experiment, observers did two blocks of 40 trials per block, presented in random order. They were not instructed about these transition probabilities, and only received cumulative feedback on their average performance at the end of each block.

Model-free analysis of behaviour

We computed d', which measures perceptual sensitivity in a bias-free way (Green and Swets 1966), separately for each individual and coherence level:

$$\mathrm{d}' = \Phi^{-1}(\mathrm{H}) - \Phi^{-1}(\mathrm{FA}) \tag{5}$$

where Φ is the normal cumulative distribution function, H is the fraction of hits and FA is the fraction of false alarms. In the 2AFC and 2IFC datasets, one of the choice options was arbitrarily treated as the stimulus-absent category. Both H and FA were bounded between 0.001 and 0.999 to allow for computation of d' in case of near-perfect performance (Stanislaw and Todorov 1999).

We quantified serial choice bias by the repetition probability $\mathrm{P}(\mathrm{repeat})$, computed as the fraction of choices that were a repetition of the previously made choice, regardless of the stimulus sequence.

HDDM model fits

We fit hierarchical drift diffusion models to each dataset, using the HDDM toolbox (Wiecki et al. 2013). Choices were coded as one of the two response identities (termed 'stimulus coding' in the toolbox), rather

than as correct and error responses, which allowed us to estimate bias towards one of two choice options.

For all datasets, we estimated overall drift rate, boundary separation, non-decision time, starting point and drift bias. Across-trial variability in drift rate was estimated at the group-level only (Ratcliff and Childers 2015). For datasets in which multiple levels of sensory evidence (motion coherence or the difference in motion coherence between two stimulus intervals) were presented, drift rate was estimated separately for each level. We specified 5% of responses to be contaminants, meaning they arise from a process other than the accumulation of evidence - for example, a lapse in attention.

For each model, we ran 15 separate chains with 10,000 samples each. Of those, half were discarded as burn-in and every second sample was discarded for thinning, reducing autocorrelation in the chains. This left 2,500 samples per chain, which were concatenated across chains. Individual parameter estimates were then estimated from each observer's posterior distribution across the resulting 37,500 samples. All group-level chains were visually inspected to ensure convergence. Additionally, we computed the Gelman-Rubin $\widehat{\mathrm{R}}$ statistic (which compares within-chain and between-chain variance) and checked that all group-level parameters had a $\widehat{\mathrm{R}}$ between 0.98-1.02 (Wiecki et al. 2013).

To fit serial choice patterns we then allowed starting point, drift bias, or both to vary with the preceding choices. In the datasets with different stimulus repetition probabilities 2AFC-1 and 2AFC-2, these biases were estimated separately for the three conditions (neutral, alternating and repetitive). We then computed the *history shift* as the difference in bias between the two previous choice identities, such that a positive history shift reflects repetition, and a negative history shift reflects alternation. Model comparison based on the Deviance Information Criterion (Spiegelhalter et al. 2002), computed over the appended model chains, was used for model selection.

Statistical tests

We quantified between-subject correlations using Pearson's correlation coefficient. The difference between two correlation coefficients that share a common variable, and its associated p-value, was computed using Steiger's test (Steiger 1980). Even though individual subject parameter estimates are not independent due to the hierarchical nature of the HDDM fit, between-subject variance in parameter point estimates can reliably be correlated to an external variable (in our case, $\mathrm{P}(\mathrm{repeat})$) without inflation of the false positive rate (Katahira 2016). To test for the significance HDDM parameters, we derived p-values directly from the group posterior (Kruschke 2013; Wiecki et al. 2013).

Results

We fit a hierarchical DDM (Wiecki et al. 2013) to six psychophysical datasets, estimating the following parameters: non-decision time, starting point, boundary separation, mean drift rate and an additive drift bias.

The DDM fit the data well, including those datasets from fixed duration tasks. In all datasets, the estimated individual drift rate correlated with perceptual sensitivity, quantified as d' (Figure 3a). In datasets

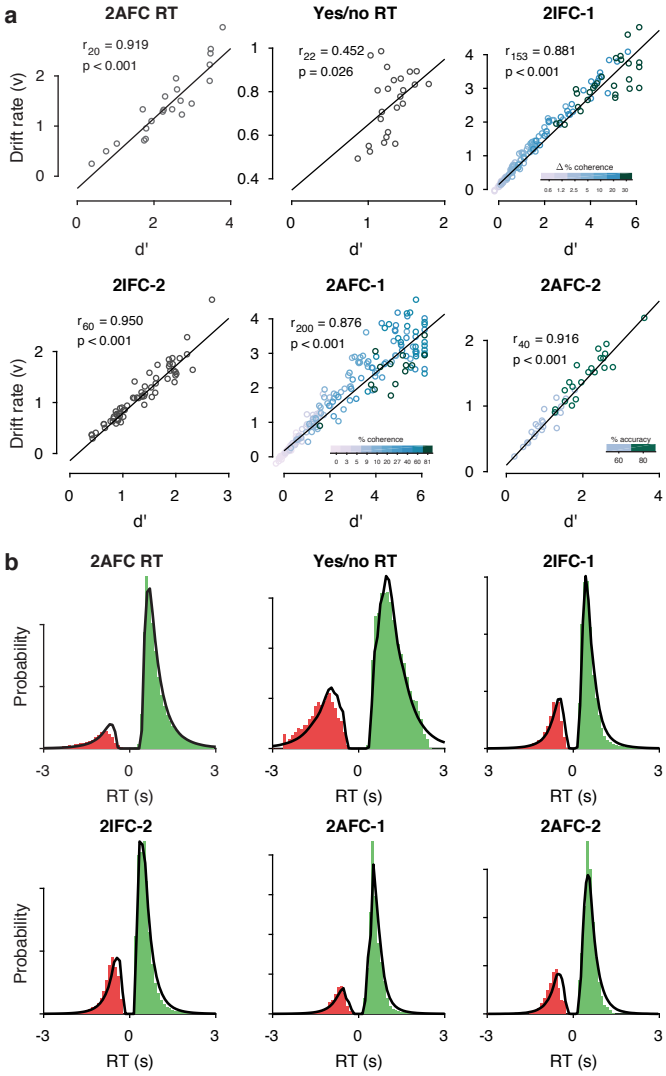

Figure 3. HDDM model fits. (**a**) Correlation between drift rate v and d'. In datasets 2IFC-1, 2AFC-1 and 2AFC-2 drift rates were estimated separately for each level of stimulus difficulty. For these datasets, the correlation coefficient displayed is a partial correlation between v and d', accounting for stimulus difficulty (indicated in the colorbar). (**b**) Posterior predictive fits. Positive RTs (green) indicate correct, and negative RTs (red) error trials. Histograms show the RT distributions of all data collapsed across subjects, with predicted RT distributions (estimated for 100 simulated trials per real trial, and smoothed by a kernel density estimate) overlaid in black.

with multiple levels of evidence strength, drift rate increased for easier stimuli (Figure 3a). Simulated data from the estimated model parameters show a close correspondence between observed and predicted RTs (Figure 3b). We computed the median and quartiles of the difference between predicted and observed RTs for all datasets: 2AFC RT, 0.357 (0.137-1.070); Yes/no RT, 0.647 (0.217-2.013); 2IFC-1, 0.285 (0.098-1.017); 2IFC-2, 0.361 (0.118-1.028); 2AFC-1, 0.180 (0.072-0.662); 2AFC-2, 0.317 (0.123-1.020). There did not seem a consistently worse fit for dataset with fixed duration stimuli. This suggests that the DDM well captures choices and reaction times when the stimulus duration is determined by the observer, as well as the experimenter (Kiani et al. 2008).

We then asked whether serial choice patterns can be explained by a history-dependent shift in starting

point, drift bias, or both. We fit the HDDM while allowing starting point, drift bias, or both to vary as a function of previous choices. We selected the best model based on the Deviance Information Criterion (DIC, Spiegelhalter et al. 2002), computed over the appended model chains.

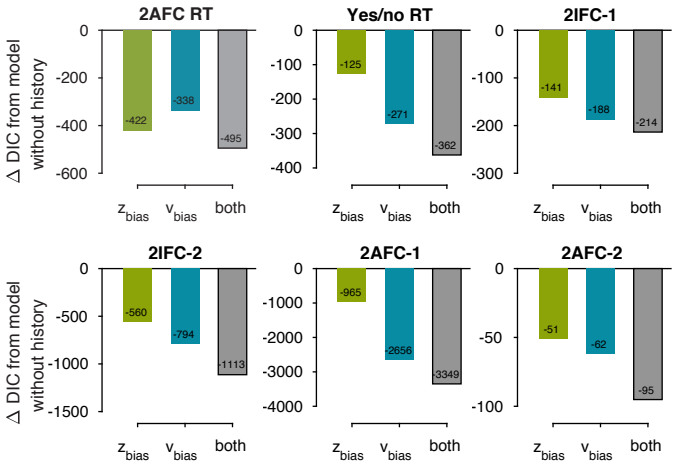

Figure 4. Model comparison. For each dataset, we compared the Deviance Information Criterion (DIC) between models where drift bias, starting point bias or both were allowed to vary as a function of previous choice. The DIC for a model without history dependence was used as a baseline for each dataset. Lower DIC values indicate a model that is better able to explain the data, after taking into account the model complexity. A $\Delta\mathrm{DIC}$ of 10 is generally taken as a threshold for considering one model a sufficiently better fit. The winning model for each dataset is indicated by a dark border.

As compared to a baseline model without history dependence, models with history terms better explained the data (Figure 4). This corroborates previous reports, which have found that serial choice patterns are ubiquitous in a broad range of psychophysical data sets and can capture a significant fraction of behavioural variability (Fründ et al. 2014). We then compared models where drift bias, starting point bias, or both were allowed to vary as a function of previous choice. Across datasets, a model with both components best explained the data, as indicated by the lowest DIC values (Figure 4). Furthermore, a model with only history-dependent drift bias better captured the data than a model with only history-dependent starting point bias (except for the 2AFC RT dataset; blue vs green bars in Figure 4).

We used the winning model, in which both starting point bias and drift bias varied with previous choice, to further investigate the parameter estimates derived from this joint fit. We combined the parameters estimated separately for the two previous choice options into a measure we call 'history shift', quantifying how much each parameter varies as a function of the immediately preceding choice. We then asked to what degree the history shift in these two parameters, drift bias and starting point bias, captured model-free repetition probability. Repetition probability correlated to the history shift in drift bias, and less so with the history shift in starting point, across datasets (Figure 5, significant correlations indicated with solid regression lines).

For each dataset, we computed the correlation between $\mathrm{P}(repeat)$ and the history shift in starting point, and the correlation between $\mathrm{P}(repeat)$ and the history shift in drift bias (Figure 6). Comparing the two correlation coefficients showed that model-free serial choice biases were more closely related to shifts in

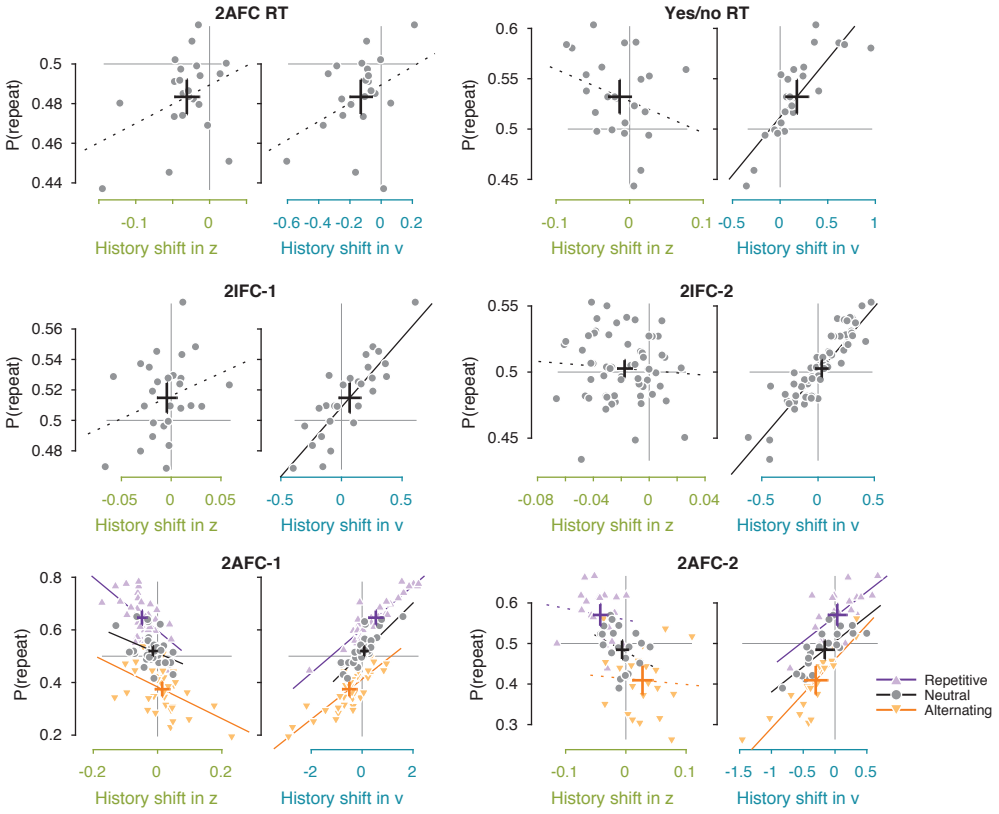

Figure 5. Correlations between repetition behaviour and DDM biases. Correlations between $\mathrm{P(repeat)}$ and (i) history-dependent changes in starting point bias (left, green axis) and (ii) history-dependent changes in drift bias (right, blue axis). The parameter estimates were obtained from a model where both bias terms were allowed to vary with previous choice. For datasets 2AFC-1 and 2AFC-2, model-based and model-free estimates of serial choice bias were quantified separately for runs where the stimulus tended to repeat (purple), alternate (orange) or did not depend on previous stimuli (grey). Horizontal and vertical lines indicate the neutral point of no bias towards either repetition or alternation. Crosses show the group mean ± s.e.m. in both directions. Thick lines show the best fit of a linear regression, dashed if the correlation between two variables was not significant.

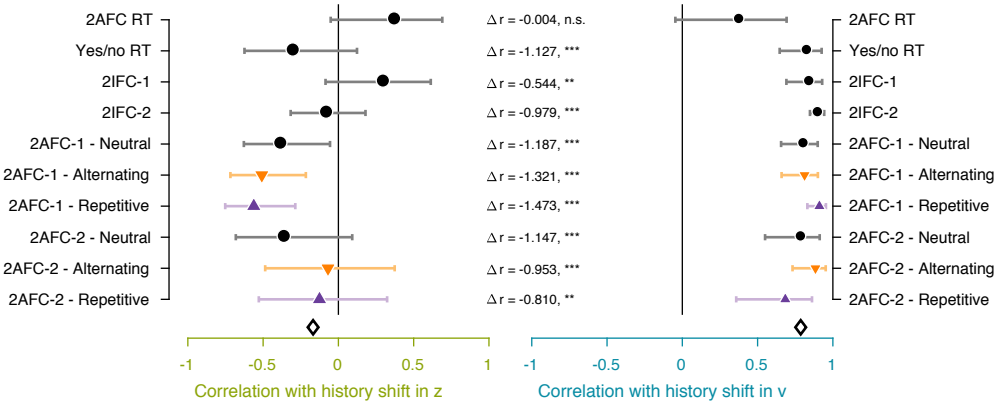

Figure 6. Summary of correlations across datasets. Correlations between $\mathrm{P(repeat)}$ and history-dependent changes in starting point bias (left) and history-dependent changes in drift bias (right). Error bars indicate the 95% confidence interval of the correlation coefficient. Δr quantifies the extend to which the two DDM bias terms are differentially able to predict individual $\mathrm{P(repeat)}$, and the corresponding p-value was computed using Steiger's test. The black diamond indicates the mean correlation coefficient across datasets.

drift bias than shifts in starting point across datasets.

When environmental statistics were biased towards repetition or alternation (datasets 2AFC-1 and 2AFC-2), we additionally observed effects on the history shift in starting point. Surprisingly, these shifts went in the opposite direction from those in drift bias (Figure 7) and correlated negatively with individual repetition probabilities (Figure 6). To ensure that this pattern was not solely a result of the hierarchical Bayesian fit, we also did the analyses using χ^2 optimization (Ratcliff and Tuerlinckx 2002), and obtained qualitatively similar results (data not shown).

The opposite shifts in drift bias and starting point as a function of previous choice can be explained when considering that these parameters not only affect choice bias, but also the associated reaction times. For example, shifts in drift bias and starting point towards the previous choice not only predict choice repetition, but also predict faster reaction times for the repeated choice (examining the shape of the RT distribution is crucial to dissociate between the two; Figure 1). It might be that the biased environments induced serial biases (repetition or alternation) but without changes in the associated reaction times. In that case, the model can converge on a regime where a shift drift bias accounts for the serial bias, while the associated effect on reaction time is neutralized by an opposite shift in starting point.

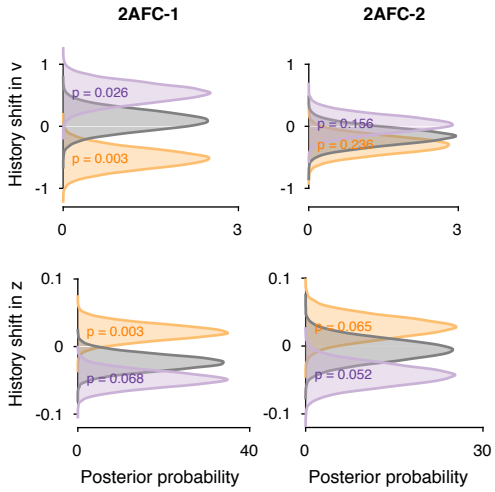

Figure 7. Posterior predictives of history effects in the biased conditions. For each of the datasets with biased stimulus sequences, the posterior predictive distributions are shown for repetitive, neutral and alternating conditions. Top: history shift in drift bias. Bottom: history shift in starting point bias. P-values show the overlap between each biased distribution and the neutral distribution in that dataset (Kruschke 2013). Note that the direction of the history bias is opposite for the two parameters.

Interestingly, the effects of starting point were largest in the 2AFC-1 datasets as compared to the 2AFC-2 dataset. With identical underlying stimulus repetition probabilities, one crucial difference between the two studies were the instructions given to observers. In the 2AFC-1 dataset, observers were informed that there could be structure in the stimulus sequence, and that these sequences could vary between sessions (Braun et al. 2017). In the 2AFC-2 dataset, on the other hand, no mention was made of stimulus transition probabilities. When debriefed after the final session, the majority of observers indicated that they did not perceive specifics blocks as repeating or alternating, and reported that they did not feel as if they could

predict the upcoming stimulus from the previous sequence. This may indicate that a history-dependent shift in starting point relies on observers' strategic adjustments of their response strategy.

Discussion

Observers' choices often depend on choices made before, even in environments where subsequent stimuli are independent. Such serial biases are ubiquitous in perceptual decision-making, but have so far mainly been investigated using the binary outcome of each choice. We here used Drift Diffusion Modeling to further decompose serial choice patterns into two possible sources of bias: a shift in the starting point of evidence accumulation, or a change in the rate with which evidence for one versus the other option is accumulated. These two effects can lead to identical choice patterns, but they can be distinguished based on their effects on reaction time distributions. Most importantly, these two sources of bias reflect dissociable algorithmic components of the decision-process, allowing for an investigation into the specific implementations of bias during perceptual decision-making.

A number of previous studies have used the drift diffusion model to tease apart starting point and drift bias effects due to various manipulations of experimentally induced choice bias. When choice bias is induced by assigning asymmetric prior probabilities to the two available choice options, the starting point of evidence accumulation shifts in the direction of the more likely choice option (Leite and Ratcliff 2011; Mulder et al. 2012; White and Poldrack 2014). Similarly, a higher reward associated with one response option moves the starting point towards that option (Gao et al. 2011; Leite and Ratcliff 2011; Mulder et al. 2012; White and Poldrack 2014). There shifts in starting point correspond to optimal behaviour; within the DDM, performance is maximized by shifting the starting point by a distance proportional to the relative probability or payoff between the two alternatives (Edwards 1965; Bogacz et al. 2006; van Ravenzwaaij et al. 2012).

Biases in the drift rate of evidence accumulation can be induced by instructing observers on the criterion separating stimulus classes - for example, when judging whether there are 'many' or 'few' items in a display based on an arbitrary cutoff (Leite and Ratcliff 2011; White and Poldrack 2014). Drift bias is closely related to stimulus processing, and can either arise from a bias already at the level of sensory encoding or by selective integration of evidence for one or the other choice option into the decision variable. On the other hand, shifts in starting point are thought to be more closely related to response preparation, and do not interact with processing of the stimulus itself. Our findings suggest that serial choice biases most likely arise from a change in stimulus processing, in a process perhaps akin to history-dependent feature-based attention.

Several previous studies have investigated the effect of previous choices on only the starting point of evidence accumulation (Gold et al. 2008; Bode et al. 2012; Zhang et al. 2014). To our knowledge, however, we show for the first time that drift bias captures a major component of behavioural variability resulting from previous choices.

We additionally observed a history-dependent shift in the starting point of evidence accumulation when stimulus sequences were biased towards repetition or alternation. Using a different sequential sampling

model (the LATER model, Carpenter and Williams 1995), Kim et al. 2017 recently found that the starting point of evidence accumulation tracks the prior of observers making saccades in an environment where transition probabilities between two stimulus options changed unpredictably throughout the experiment. Notably, this result was interpreted in the context of a normative model, where observers' beliefs about the upcoming stimulus are continually updated based on expectations about the volatility of the environment (Glaze et al. 2015).

It is interesting to speculate that serial choice patterns in biased environments may reflect the superposition of two underlying processes. First, an automatic history-dependent shift in attention may bias the processing of subsequent stimuli, and is reflected in drift bias estimates in experiments without serial dependencies between stimuli. Second, a history-dependent shift in starting point may reflect adjustments in response preparation, present when biased stimulus sequences lead observers to strategically update the starting point of accumulation. Future work is needed to investigate not only under what conditions history-dependent shifts in starting point and drift bias co-occur, but also why they shift in opposite directions in the datasets with biased environments presented here.

In conclusion, we found that serial choice biases in neutral environments are most closely related to history-dependent shifts in drift bias; previous choices affect the way stimuli are processed or integrated into a decision. When stimulus sequences are biased towards consistent repetition or alternation, an additional shift in the starting point of evidence accumulation was observed. These findings suggest that serial choice patterns, which are ubiquitous in psychophysical experiments, may arise from a biased evidence accumulation process that depends on experimental history.

Code availability

Code to fit the HDDM models and reproduce all figures is available at `https://github.com/anne-urai/serialDDM`.

Acknowledgements

We thank Bharath Chandra Talluri, Anke Braun and Jan Willem de Gee for kindly sharing their behavioural datasets, Gilles de Hollander for discussions about HDDM model fits, Peter Murphy for sharing the code to generate analytical RT distributions, and Jan Willem de Gee for discussions and comments on the manuscript. Elisa Battistoni and Fleur Duyser assisted with data collection for the 2AFC RT dataset. Christiane Reissmann, Karin Deazle, Samara Green and Lina Zakarauskaite helped with participant recruitment and data acquisition for the 2IFC-2 MEG study. We acknowledge computing resources provided by NWO Physical Sciences. This research was supported by the German Academic Exchange Service (DAAD) (to A.E.U.) and the German Research Foundation (DFG), SFB 936/A7, SFB 936/Z1, DO 1240/2-1 and DO 1240/3-1 (to T.H.D.).

Supplementary Figures

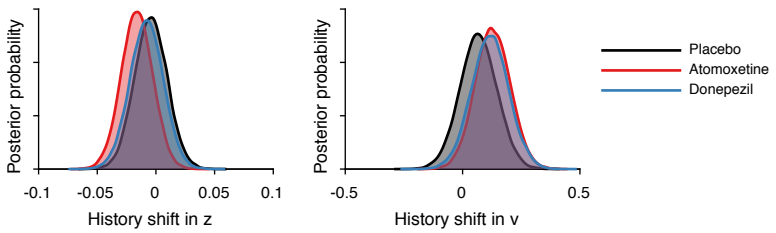

Figure S1. Group-level posterior estimates of serial choice bias for each pharmacological intervention group. (**a**) Posterior probability of the history shift in starting point, separately for each group. There was no significant difference between atomoxetine and placebo ($p = 0.267$), between donepezil and placebo ($p = 0.430$), or between atomoxetine and donepezil ($p = 0.338$). (**b**) Posterior probability of the history shift in drift bias, separately for each group. There was no significant difference between atomoxetine and placebo ($p = 0.285$), between donepezil and placebo ($p = 0.329$), or between atomoxetine and donepezil ($p = 0.458$).

2.4 | Neural bases of serial choice bias

Urai AE, Engel AK, Donner TH.

Abstract

When making perceptual decisions about sequences of independent stimuli, observers are consistently biased towards or away from their previously chosen response. We previously found that such *serial choice biases* arise mostly from preceding choices, rather than the associated motor responses (Braun et al. 2017) and that they might reflect the across-trial accumulation of uncertainty-weighted decision variables (Braun et al. 2017; Urai et al. 2017).

Here, we used MEG to tease apart the contributions of neural signatures of sensory encoding, evidence accumulation, and motor preparation in driving serial choice patterns. The goals of this paper are two-fold. First, we found that post-choice pupil responses and reaction times both reflected decision uncertainty. These two measures reduced observers' tendency to repeat their responses from one trial to the next, replicating our previous findings with a larger sample (Urai et al. 2017). Second, we found that previous choices affect sensory representations as well as motor planning and evidence integration. Neural signatures of both sensory and motor representations could be used to predict observers' upcoming choices, over and above the identity of the current stimulus and previous choice. This suggests that previous choices can affect neural activity at the level of sensory encoding and response preparation, which may then give rise to serial choice patterns in behaviour.

Author contributions: A.E.U., A.K.E. and T.H.D. designed research; A.E.U. collected and analyzed the data; A.E.U. wrote the manuscript; T.H.D. supervised research.

Introduction

Serial choice biases, individual tendencies to systematically repeat or alternate previous choices, are ubiquitous in perceptual decision-making. These behavioural patterns may be similar across species, and might arise from the interplay of different mechanistic and algorithmic effects (Fernberger 1920; Gold et al. 2008; Yu and Cohen 2008; Busse et al. 2011; de Lange et al. 2013; Akaishi et al. 2014). When quantifying behaviour globally, these biases can easily go unnoticed; when conditioning on the immediately preceding choice, however, strong patterns of individual choice patterns emerge.

A range of behavioural studies have started to comprehensively characterized serial choice biases (Akaishi et al. 2014; Fischer and Whitney 2014; Fritsche et al. 2017; Urai et al. 2017), showing that serial biases are modulated by decision confidence (Braun et al. 2017; Urai et al. 2017) and can be adapted to the statistical structure of stimulus sequences (Abrahamyan et al. 2016; Braun et al. 2017). However, the neural underpinnings of serial choice biases are unclear. Specifically, neural correlates of serial choice biases have been reported in visual (St. John-Saaltink et al. 2016), motor (Pape and Siegel 2016) and frontal cortices (Akaishi et al. 2014).

Two studies have used MEG and found that low-frequency lateralization signals reflect the last trial's motor response (de Lange et al. 2013; Pape and Siegel 2016), which may bias observers to alternate their responses (Pape and Siegel 2016). In monkeys performing a perceptual choice task, previous choices bias the starting point of firing rate levels in oculomotor, but not in sensory areas (Gold et al. 2008). An fMRI study found that the orientation of previously chosen stimuli can be decoded from early visual cortex (St. John-Saaltink et al. 2016), possibly through decision-related feedback signals from higher areas. It is thus an intriguing possibility that previous choices exert an effect on decision processing similar to selective attention, influencing decision-making and motor circuits as well as early sensory circuits through feedback from higher areas (Nienborg and Cumming 2009; Buffalo et al. 2010; Wimmer et al. 2015).

It remains an open question whether serial choice biases result from single neural basis, or instead reflect the superposition of neural mechanisms that may differentially contribute to behaviour. Here, we used MEG to address this question. Sixty-one participants performed a perceptual decision-making task in which subsequent stimuli were independent. We then extracted oscillatory signatures of sensory encoding, evidence integration and response preparation, and found that all three neural signatures were affected by previous choices. Gamma-band responses over visual cortex, a measure of sensory encoding, reflected not just the current stimulus but also the previous response. Beta-band lateralization over motor cortex showed a response-specific 'rebound' that carried over to the next trial (de Lange et al. 2013; Pape and Siegel 2016). During evidence accumulation, the slope of lateralized beta-band buildup was also steeper for decisions that were a repetition of the previous choice. Importantly, single-trial neural signatures at the sensory and motor level - measured before stimulus onset - predicted observers' upcoming choice. This suggests that previous choices affect neural activity at multiple processing stages during perceptual decision-making, biasing choice behaviour.

Materials and Methods

Participants

Sixty-four participants (aged 19-35 years, 43 women and 21 men) participated in the study after screening for psychiatric, neurological or medical conditions. All subjects had normal or corrected to normal vision, were non-smokers, and gave their informed consent before the start of the study. The experiment was approved by the ethical review board of the University Medical Center Hamburg-Eppendorf (reference PV4648). Two participants stopped after the first session, and were excluded from all analyses. One participant was additionally excluded from all analyses because of excessive head motion during MEG measurements. The results are thus based on the data from 61 participants.

Task

Participants compared the strength of random-dot motion stimuli in two intervals (Figure 1b, Urai et al. 2017). Throughout the experiment, a red 'bulls-eye' fixation target of 1.5° diameter (Thaler et al. 2013) was present in the centre of the screen. Each trial started with a pre-stimulus interval of 500 - 1000 ms of randomly moving dots. A beep (50 ms, 440 Hz) indicated the onset of a reference stimulus, consisting of 70% coherence motion, that lasted for 750 ms. This was followed by a variable (300 – 700 ms) interval of dynamic random motion. An identical beep indicated the onset of the second test stimulus, which had a variable difficulty (motion coherence deviation from the baseline of 70% coherent motion). We applied a counterbalancing scheme to ensure that each stimulus type (weaker or stronger second interval) is followed by itself and the other stimulus equally often (Brooks 2012). Dots moved in one of the four diagonal directions, counterbalanced across and constant within participants. After the offset of the second stimulus, observers indicated whether they perceived the test stimulus as stronger or weaker compared to the reference stimulus, with a maximum response time of 3 seconds. The hands used to indicate weaker and stronger stimuli were counterbalanced across subjects. Feedback (correct/incorrect) was then indicated by a tone of 150 ms (880 Hz or 200 Hz, feedback-tone mapping counterbalanced between subjects).

Stimuli

Dynamical random dot stimuli were presented in a central annulus (outer radius 14°, inner radius 2°) around fixation. The annulus was defined by a field of dots with a density of 1.7 dots/degrees^2. Dots were 0.2° in diameter and had 100% contrast from the black screen background. Signal dots were randomly selected on each frame, and moved with 11.5°/second in one of four diagonal directions (counterbalanced across, and constant within participants). Signal dots that left the annulus wrapped around and reappeared on the other side. Moreover, signal dots had a limited 'lifetime', and were replotted 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' noise with a 'different' rule (Scase et al. 1996). Additionally, to

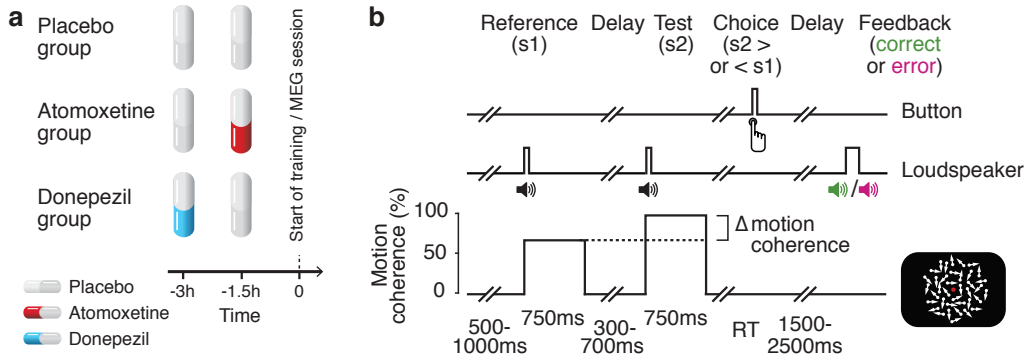

Figure 1. Pharmacology and task. (**a**) Schedule of pharmacological intervention. Reproduced with permission from Pfeffer et al. 2017. (**b**) 2-interval motion strength discrimination task. Dynamic random dot patterns were displayed throughout each trial. In two successive intervals (onset cued by beeps), the dots moved in one of the four diagonal directions (fixed per observer): A first 'reference' interval $s1$ with always 70% motion coherence, and a second 'test' interval $s2$ with motion coherence larger or smaller than the reference. The motion coherence difference between the two intervals was titrated at each individual's 70% accuracy threshold. Observers reported whether the test stimulus contained stronger or weaker motion than the reference, by pressing one of two buttons. They received auditory feedback after a variable delay. Adapted from Urai et al. 2017.

avoid subjects tracking individual signal dots as they move through the annulus, three independent motion sequences were interleaved on subsequent frames (Roitman and Shadlen 2002).

Thresholding

Before their first MEG session, participants received instructions and then did a short thresholding session in which a total of 600 stimuli of 1.25, 2.5, 5, 10, 20 and 30% motion difference (from the 70% coherence baseline) were randomly interleaved. The inter-stimulus interval was 1 second, and subjects took a short break after each set of 125 trials. They did not receive feedback. Stimuli were presented on an LCD screen at 1920x1080 resolution and 60 Hz refresh rate, 60 cm away from the subjects' eyes. To determine each individual's psychometric threshold, we fit a cumulative Weibull as a function of difficulty c, defined as

$$\Phi(c) = \delta + (1 - \delta - \gamma)\left(1 - e^{\left(\frac{-c}{\alpha}\right)^{\beta}}\right) \quad (1)$$

where δ is the guess rate (chance performance), γ is the lapse rate, and α and β are the threshold and slope of the psychometric Weibull function, respectively (Wichmann and Hill 2001b). While keeping the guess rate δ bound at 50% correct, we fit the parameters α, β and γ using a maximum likelihood procedure implemented by minimizing the logarithm of the likelihood function. This was done using a Nelder-Mead simplex optimization algorithm as implemented in Matlab's *fminsearch* function. A parametric bootstrap was used to obtain 95% confidence intervals around the Weibull parameter fits (Wichmann and Hill 2001a). The individual threshold was taken as the stimulus difficulty corresponding to a 70% correct fit of the cumulative Weibull.

Participants additionally went through a short staircase procedure, using a 2-up 1-down method. The coherence difference between the two stimuli started at their 70% threshold as obtained from the Weibull fit, and was increased by 0.1% coherence on making an error, and decreased by 0.1% on giving two consecutive

correct answers. Participants did 100 trials, over which we computed the average threshold. This procedure ensured that any early learning effects or strategy adjustments during the thresholding session would not lead to artificially high threshold values (and possible ceiling effects in behavioural performance). Thresholds ranged from 3.3% to 13.4% (mean 6.9%) motion coherence difference.

Training

Participants performed the task at their individual motion threshold for 600 trials during the first MEG session, three times for 1500 trials outside the scanner (on separate days), and for a last MEG session of 600 trials. During MEG sessions, when the pupil was recorded, random dots were displayed 1.5-3 seconds before, and 2-3 seconds after feedback. These 'pupil-rebound' intervals served to optimally measure the pre-feedback pupil response by (i) allowing the pupil to return to baseline before presenting the feedback sound and (ii) encouraging the participants to keep their eyes open and fixated after the feedback for clear measurement of pupil size. After this second 'rebound' interval, subjects could blink during a break indicated by stationary dots on the screen, and continued to the next trial by pressing a button.

In the behavioural training sessions, we did not record the pupil, and presented feedback immediately after the subjects' response. An ISI of 1s was observed before continuing to the next trial. Subjects completed training on 4500 trials, over 3 separate sessions, between the two MEG recordings.

Pharmacology

At the start of each experimental session, participants orally took 40 mg atomoxetine (Strattera®), 5 mg donepezil (Aricept®), or placebo. Since the time of peak plasma concentration is 3 hours for donepezil (Rogers and Friedhoff 1998) and 1-2 hours for atomoxetine (Sauer et al. 2005), we used a placebo-controlled, double-blind, double-dummy design, entailing an identical number of pills at the same times before every session for all participants (Figure 1a). Participants in the donepezil group took 5 mg of donepezil 3 hours, and a placebo 1.5 hours before starting the experimental session. Participants in the atomoxetine group took a placebo 3 hours, and 40 mg of atomoxetine 1.5 hours before the experimental session. Those in the placebo group took identical-looking sugar capsules both 3 and 1.5 hours before starting the session. This ensured that either drug reached its peak plasma concentration at the start of the experimental training. The drug doses were based previous studies with healthy participants (Chamberlain et al. 2009; Rokem and Silver 2010). Blood pressure and heart rate were measured and registered before subjects took their first and second pill. In the three hours before any MEG or training session, participants waited in a quiet room. To avoid the accumulation of drugs, we ensured a minimum of 2 days between subsequent experimental sessions. In total, 19 participants in the placebo, 22 in the atomoxetine, and 20 in the donepezil group completed the full study.

Personality questionnaires

In the waiting time before their last MEG session, participants filled out three personality questionnaires in German or English. The behavioural inhibition/activation (BIS/BAS) questionnaire, measuring behaviour activation and inhibition, was scored according to (Carver and White 1994). The Autism Quotient (AQ) questionnaire, which measures autistic traits (Baron-Cohen et al. 2001), was scored according to (Freitag et al. 2007). The Penn State Worry Questionnaire (PSWQ-D), which measures worrying behaviour, was scored according to (Meyer et al. 1990). See Figure S2 for an overview of personality scores.

Eye-tracking and pupillometry

During MEG sessions, horizontal and vertical gaze position and pupil diameter were recorded at 1000 Hz using an MEG-compatible EyeLink 1000 on a long-range mount (SR Research) at 60 cm from the subject's eye. The eye tracker was calibrated before each block of training. To minimize any effect of light on pupil diameter, the overall luminance of the display was held constant throughout the experiment by always displaying a cloud of white dots on the screen.

Missing data and blinks, as detected by the EyeLink software, were padded by 150 ms and linearly interpolated. Additional blinks were found using peak detection on the velocity of the pupil signal and linearly interpolated. We used a deconvolution approach to estimate the effect of blinks and saccades on the pupil signal and removed these artefacts from the data using linear regression (Knapen and de Gee 2016). The residual pupil time series were band-pass filtered using a 0.01 to 10 Hz second-order Butterworth filter, z-scored per run of 60 trials, and resampled to 400 Hz.

Behavioural analyses

To test whether RT reflected decision uncertainty, we used data from both MEG sessions and did a paired t-test on the median RT for correct and error trials for each subject. On the pupil timecourses, we performed a cluster-corrected permutation test on the mean pupil response between correct and error trials, which takes into account the temporal autocorrelation in the signal to correct for multiple comparisons (Maris and Oostenveld 2007).

We then tested whether these two proxies for decision uncertainty affected serial choice patterns, i.e. the tendency to repeat the last choice. To visualize the effect, we split the data into three equal-sized bins of RT or pupil, and computed the probability that the next choice was a repetition of the current choice (Figure 2c,d). Because we only used one fixed level of task difficulty, we could not fit the full logistic regression model previously used (Fründ et al. 2014; Urai et al. 2017). Instead, we used generalized linear mixed models (GLMMs) to predict each choice as a repetition or alternation from the previous choice. We included an intercept (capturing group-level repetition behaviour), a vector of across-trial pupil responses or response times (normalized by z-scoring the log-transformed RT for each subject and session), and random

intercepts for each observer. For each GLMM, we report the estimated fixed effects coefficient β, its lower and upper 95% confidence intervals, and its p-value.

MEG acquisition

MEG was recorded using a 275-channel CTF system in a shielded room. Horizontal, vertical and radial EOG, bipolar ECG, and an electrode at location POz (about 4 cm above the inion) were recorded simultaneously. See Figure S1 for an analysis of ECG data. All signals were low-pass filtered online (cut-off: 300 Hz) and recorded with a sampling rate of 1200 Hz. To minimize the displacement of the subject's head with respect to the MEG sensors, we used online head-localization (Stolk et al. 2013) to show the head position to the subject inside the MEG chamber before each block. Subjects were then asked to move themselves back into their original position, correcting slow drift of their head position during the experiment. Moreover, between the two recording days, the original head position from day one was used as a template for day two.

Stimuli were projected into the MEG chamber using a beamer with a resolution of 1024 x 768 pixels and a refresh rate of 60 Hz. The screen was positioned 65 cm away from their eyes.

MEG preprocessing

Data were analyzed in Matlab using the Fieldtrip Toolbox (Oostenveld et al. 2011) and custom scripts. MEG data were first resampled to 400 Hz, matched with the corresponding EyeLink timecourses, and epoched into single trials from fixation to 2s after feedback. Trials with squid jumps were detected by fitting a line to each single-trial log-transformed Fourier spectrum, and rejecting trials where the intercept was detected as an outlier based on Grubb's test. To remove the effect of line noise on the data (which were not always restricted to 50 Hz and its harmonics due to technical issues), we computed the cross-spectrum of the data at 50 Hz, resulting in a complex matrix of size n by n, where n is the number of channels. We applied singular value decomposition to this cross-spectrum, and took the first eigenvector (corresponding to the largest singular value) as the spatial topography reflecting line noise. The two-dimensional space spanned by the real and imaginary parts of this eigenvector was then projected out of the data, effectively suppressing any signal that co-varies with activity at 50 Hz. Line noise around 50, 100 and 150 Hz was then removed by a bandstop filter, and each trial was demeaned and detrended.

We also removed trials where (i) no response was recorded in the 3s after stimulus offset, (ii) head motion of any of the 3 fiducial coils (nasion, right and left pre-auricular point) exceeded a translation of 6 mm from the first trial of the recording (Stolk et al. 2013), (iii) the min-max data range was larger than 7.5 pT, usually caused by cars driving by close to the building, (iv) muscle bursts, eye blinks or saccades were detected in the time before buttonpress. Blinks were detected by applying a 1-15Hz bandpass filter to the horizontal and vertical EOG channels, z-scoring the data, and marking any epochs with a z-score larger than 5 as a blink event. Muscle bursts were detected by applying a 9th-order 110-140 Hz Butterworth filter to all MEG

channels, z-scoring the data, and marking any z-score above 20 as a muscle burst. We then epoched the data on each trial around the reference stimulus s1 (-0.5 to 1.5s), test stimulus s2 (-0.5 to 1.5s), buttonpress (-1 to 1s) and feedback (-0.5 to 1.5s).

MEG time-frequency analysis

We computed time-frequency representations for each of the four time windows of interests. High and low frequency ranges were analyzed separately. For the low frequencies (3-35 Hz in steps of 1 Hz), we used a Hanning window with a length of 400 ms in steps of 50 ms and a frequency smoothing of 2.5 Hz. For high frequencies (36-120 Hz in steps of 2 Hz), we used the multitaper technique with five discrete proloid slepian tapers (Mitra and Pesaran 1999), a window length of 400 ms in steps of 50 ms, and 8 Hz frequency smoothing. We decomposed the data into two planar gradients before estimating power, and combined these afterwards to simplify the topographical representation of evoked power.

We converted the time-frequency representations into % signal change from the pre-trial baseline. This baseline was defined as the across-trial average in the -300 to -200 ms before reference onset across, separately for each sensor and frequency.

On the time-frequency spectra, we used cluster-based permutation testing (Maris and Oostenveld 2007) to find clusters across time and frequency that differed for our contrast of interest.

Frequency selection

We based our choice of frequency bands on previous work: for motor lateralisation in the beta-band, we used the range 12-36 Hz (Donner et al. 2009) and for visual gamma-band responses we used 65-95 Hz (Siegel et al. 2006). Note that while the original paper by Siegel et al. (2006) found the frequency range from 60-100 Hz to have the strongest response to coherent motion, we here excluded 60Hz to avoid the evoked steady-state visual evoked potential (Figure S3). Similarly, we excluded 100 Hz to avoid including power-line noise into our estimates of gamma-band responses.

Sensor selection

To define sensors/voxels corresponding to visual motion encoding, we contrasted trials with strong vs. weak motion coherence in the second interval, and computed power in the gamma-range (65-95 Hz) from 150 to 750 ms after stimulus onset. We then selected the 25 most active sensors in the first and second session separately.

Similarly, for sensors/voxels corresponding to response preparation, we contrasted trials in which the left vs. the right hand was used to respond. We computed power in the beta range (12-36 Hz) in the 500 ms before button press (Donner et al. 2009), and used the same split-half approach to define the 25 most active sensors for the contrast left vs. right, as well as the 25 most active sensors the for opposite contrast, to extract left and right motor regions respectively.

For each session, we then extracted single-trial values at each point in time by collapsing over the selected frequency range and the sensors defined on the other session.

MEG source reconstruction

Structural T1-weighted magnetization prepared gradient-echo images (TR = 2300 ms, TE = 2.98 ms, FoV = 256 mm, 1 mm slice thickness, TI = 1100 ms, 9° flip angle) with 1x1x1 mm^3 voxel resolution were obtained on a 3T Siemens Magnetom Trio MRI scanner (Siemens Medical Systems, Erlangen, Germany). Fiducials (nasion, left and right intra-aural point) were marked on the MRI. A single-sphere volume conduction model (Nolte 2003) was created for each individual, and individual source models were warped to the Colin27 brain (Holmes et al. 1998) using a nonlinear transformation for group analyses and visualization. We were not able to obtain an individual MRI for one of the 61 observers; this participant was included in sensor-level, but not in source-level analyses.

We extracted virtual sensors based on source-reconstructed voxel timecourses from a sliding window DICS beamformer (Van Veen et al. 1997; Gross et al. 2001). Within each frequency range (beta/gamma), we computed a common filter based on the cross-spectral density matrix separately for each stable epoch. For both the reference and test stimulus epoch separately, those were the baseline (-300 to 0 ms) and the evoked response (150 to 750 ms). For the response epoch, this was the pre-buttonpress window (-400 to 100 ms) and for the feedback, those were the pre-tone (-400 to 0) and post-tone (200 to 700) windows. We then applied the beamformer in a sliding window of 250 ms, with steps of 50 ms, keeping single-trial source estimates, using the appropriate common filter for each epoch. There were then converted into % signal change from the pre-trial baseline, making the units of MEG oscillatory responses comparable across sensor- and source-level analyses.

ROI definition

To define visually selective voxels, we used the atlas by Wang and Munoz (2015). We first performed source reconstruction for each voxel, then averaged all voxels within each atlas region. Lastly, we grouped the individual regions into visual field maps (Wandell et al. 2007): medial occipital (ventral and dorsal V1, V2 and V3); lateral occipital (LO1, LO2 and hMT); ventral occipital (hV4, VOI and VO2); dorsal occipital (V3A, V3B); and posterior parietal (IPS0 to IPS4). See Figure S4 for a representation of these visual regions on the template brain, and Figure S5 for data for each visual field map separately). To increase the signal strength of visually evoked source-level activity, we averaged all visual areas together for the main analyses reported.

To define motor regions, we manually defined three coordinate locations on the Colin27 template brain, and averaged together those three motor regions. In MNI coordinates, we located the 'handknob' of the precentral gyrus, the intersection of the precentral and superior frontal sulcus, and lastly the postcentral aIPS (see Table S1 for MNI coordinates).

We computed the difference between each left and right region into an index of lateralization, and subsequently averaged the three motor regions. This lateralization was flipped for half of the observers who responded 'stronger' with their left and 'weaker' with their right hand, enabling us to express motor lateralization relative to the stimulus identity for all observers.

MEG source statistics

We extracted source-level single-trial estimates of sensory encoding (visual gamma-band responses from 150 to 750 ms after the onset of the test stimulus) and motor preparation (beta-band lateralization during the 500 ms pre-trial fixation baseline).

We used GLMMs (see also Behavioural analyses) to test for the effect of previous choices on these neural measures (using a linear link function), and for the predictive value on these neural measures on current choices (using a logistic link function), including a random effects intercept for each observer.

To assess the effect of previous choices on the build-up of beta-band lateralization (Donner et al. 2009), we first averaged all trials within observers and conditions (previous and current choice) and then computed a linear regression on the time window from 150 ms to 750 after the onset of the test stimulus. This approach allowed for robust estimate of the build-up slope, as compared to fitting regressions on single-trial timecourses. Across observers, we use a repeated measures ANOVA to test for the effect of current and previous choice on the slope of this beta-band lateralization.

Choice probability

We computed ROC indices to quantify the degree of separability between distributions of single-trial gamma-band responses when observers responded 'stronger' vs 'weaker' (Green and Swets 1966; Britten et al. 1996). This measure quantifies, for each observer, the degree to which gamma-band responses during the reference interval predict the upcoming choice. A value of 0.5 indicates completely overlapping distributions, whereas values of 0 or 1 indicate completely separate classes. Across the group, we used a two-tailed t-test and compared the ROC indices to 0.5.

Results

Pupil-linked arousal reflects decision uncertainty and modulates serial choice bias

We first replicated the two main findings from Urai et al. (2017). First, both pupil responses and response times scaled with decision uncertainty. Since we used only one level of stimulus difficulty in this version of the experiment, the statistical model of decision uncertainty (Kepecs et al. 2008) makes the prediction that average uncertainty signals should be higher on error than on correct trials. This was indeed the case for both response times (Figure 2a) and pupil responses in the interval between response and feedback (Figure 2b).

Intriguingly, we also observed a significant cluster in the pupil timecourse already during stimulus presentation (Figure 2b). Here, the pupil was larger on correct trials, possibly reflecting decision confidence rather than uncertainty. A similar pattern (albeit small and not significant) was observed in Urai et al. (2017), Figure 3b. Notably, in both datasets, this effect was already present during viewing of the reference stimulus - before any decision-relevant information was presented to the observer. In monkeys performing a disparity discrimination task with long stimulus presentation times, the pupil was also found to scale positively with confidence during stimulus viewing (Kawaguchi et al. (2017) Decision confidence and motivation are differently associated with task-strategy during perceptual decision making. *European Conference on Visual Perception*, Berlin). Future studies could use long stimulus presentation times as well as a long post-response delay to investigate the possibility that pupil responses scale with decision confidence during stimulus viewing, and with decision uncertainty (the complement of confidence) after response.

Response times (Figure 2c, β = -0.0363 [-0.0512, -0.0214], p < 0.0001) and pre-feedback pupil responses (Figure 2d, β = -0.0174 [-0.0325, -0.0023], p = 0.0242) predicted an overall reduction in choice repetition on the next trial. Despite a few differences in experimental design (most importantly, only one level of difficulty) we thus replicated our previous finding that two proxies of decision uncertainty modulate serial choice patterns (Urai et al. 2017).

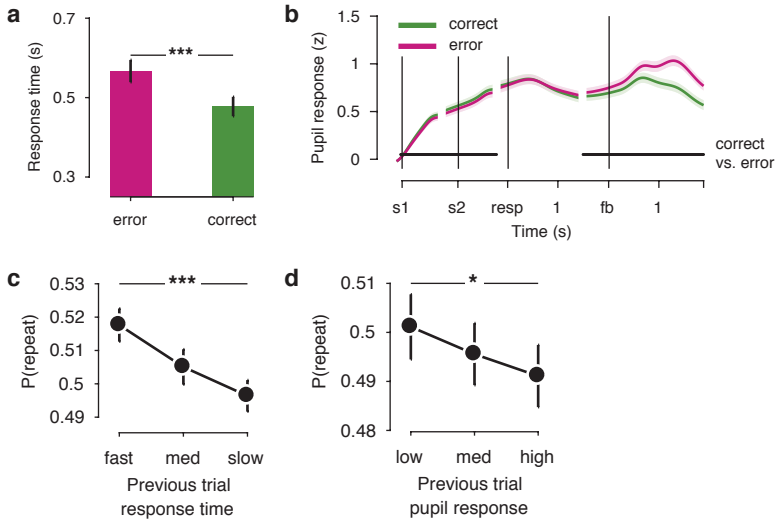

Figure 2. Pupil and RT reflect decision uncertainty and reduce choice repetition. (**a**) Response times (from the offset of the test stimulus, error bars show group mean ± s.e.m. of individual RT medians) are larger on error vs correct trials. (**b**) Pupil responses in the window between response and feedback are larger on error than trials. Black lines indicate significant timepoints at p < 0.05, obtained from a cluster-based permutation test. (**c**) Response times predict a reduction in choice repetition. (**d**) As in (**c**), but for pre-feedback pupil responses (defined as the 250 ms before feedback). Statistics in (**c**-**d**) as obtained from a logistic GLMM. *** p < 0.001, * p < 0.05.

We also investigated the overall group-level patterns of serial choice bias. We quantified serial choice bias using three measures: (i) P(repeat), the average probability of repeating the previous choice regardless of the stimulus sequence, and two measures derived from hierarchical drift diffusion model fits (see chapter 2.3), (ii) the history-shift in drift bias and (iii) the history-shift in starting point. As reported in chapter 2.3,

we observed a strong correlation between individual repetition probability and the history-shift in drift bias, but not in starting point (Figure S2, top left). While stimulus sequences were uncorrelated, average choice repetition patterns were biased towards repetition (average $P(repeat)$ = 0.5167, t_{60} = 4.6569, $p < 0.0001$). This repetition behaviour did not differ between the two MEG sessions ($P(repeat)$ = 0.5202 in the first session and 0.5141 in the last session, t_{60} = 1.2125, p = 0.2301).

Visual gamma-band responses scale with stimulus motion and previous choice

We then used our MEG data to ask which stages of the decision-making process are affected by previous choices. We extracted measures of sensory encoding, response preparation and evidence integration, and tested whether these proxies for different components of decision-making reflect previous choices.

At the sensory stage, we first confirmed that gamma-band responses over occipital cortex scaled with the strength of coherent motion on the screen (Siegel et al. 2006). We selected occipital sensors based on a split-half analyses (see Methods), and computed the full time-frequency response over all trials (Figure 3a) and for the contrast strong vs weak motion (Figure 3b). This shows a narrow-band response at the screen refresh rate of 60 Hz (the steady-state visual evoked potential or SSVEP), as well as a high gamma-band response to the onset of visual motion (Figure 3, see also Figure S3). More importantly, we observed a higher gamma-band response in the test interval for motion that was stronger vs. weaker than the 70% coherence baseline (Siegel et al. 2006), reflecting the decision-relevant stimulus information on the screen.

As we titrated the motion threshold for each observer individually, the physical difference in motion strength (in % coherence) ranged from 6.6% to 26.8% across observers. Across participants, motion coherence did not correlate to the source-level gamma-band response to strong vs. weak motion stimuli (r = 0.1127, p = 0.3912). We did not observe any significant time-frequency clusters for the contrast of previous strong vs. weak choices at the sensor level.

We then used a linear beamformer to obtain source-level estimates of visual gamma-band responses (Hipp and Siegel 2013), extracting retinotopic visual cortical areas as defined by Wang and Munoz (2015). As expected, the resulting time-courses showed a robust effect of stronger vs weaker coherent motion during the test stimulus interval (Figure 4a). From this, we extracted single-trial estimates of sensory evidence, averaging the visual cortical gamma-band response over the 150-750 ms after the onset of the test stimulus, and tested whether these visual responses reflect not just current visual input, but also previous choices.

As expected, cortical gamma-band responses increased with strong current motion (Figure 4b, blue vs. green, β = 4.1022 [2.5592, 5.6453], $p < 0.0001$). Moreover, we observed higher visual cortical gamma-band responses when the observer chose 'stronger' motion on the previous trial (Figure 4b, β = 2.2843 [0.7510, 3.8176], p = 0.0035). This effect was already present during the reference interval (β = 2.5338 [0.9712, 4.0964], p = 0.0015). This suggests that visual responses were boosted in line with the observer's previous choice. The effect of previous choices was also significant when only current correct trials were included (β = 2.4810 [0.7064, 4.2556], p = 0.0061). We repeated this analysis separately for all visual field maps; we

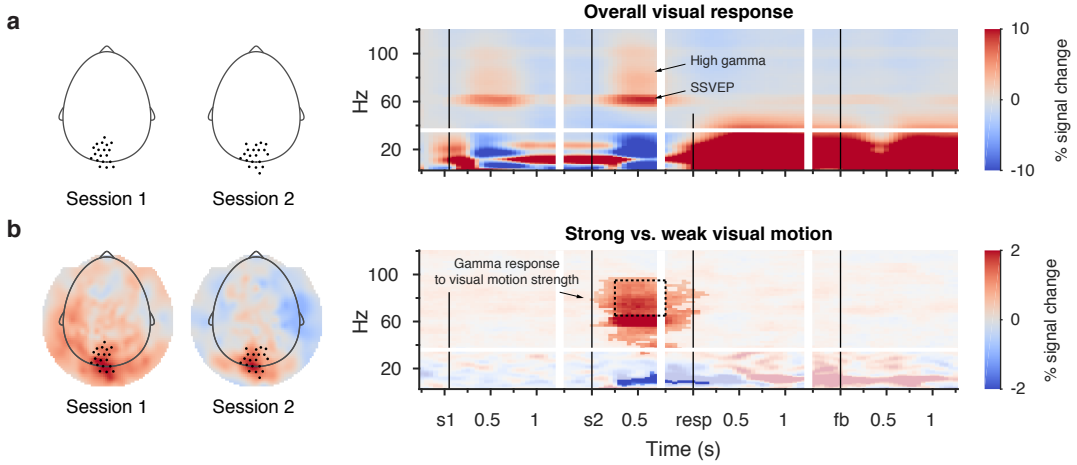

Figure 3. Sensor-level visual gamma-band responses increase with more strongly coherent motion stimuli. (**a**) Left: sensors selected at the group level, separately for both MEG sessions. Right: overall time-frequency responses of occipital sensors. (**b**) Left: sensors selected at the group level, separately for both MEG sessions. Colors indicate the contrast strong vs. weak coherent motion during the test stimulus, in the 65-95 Hz gamma range. Right: TFR contrast for strong vs weak coherent motion, for the occipital sensor groups, expressed as % signal change from baseline. Dashed box indicates the time- and frequency selection for the topographical representations on the left. Saturation indicates significant time- and frequency points from a cluster-based permutation test.

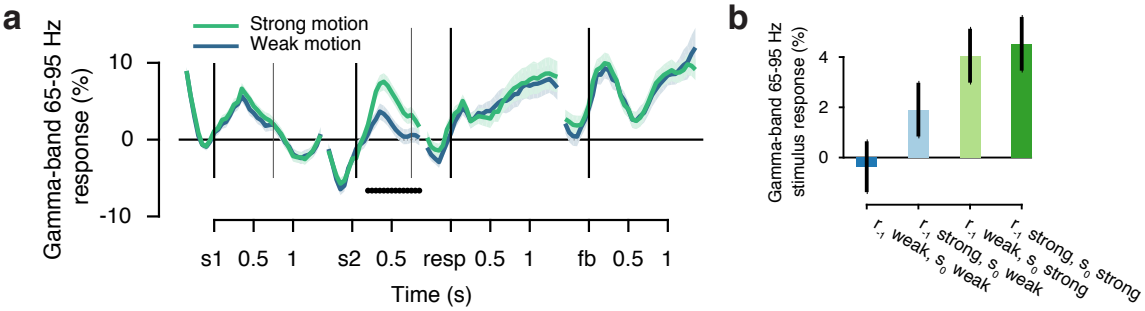

Figure 4. Average gamma-band response of retinotopic visual cortical areas reflects current stimulus and previous choice. (**a**) Gamma-band timecourses in % signal change from the pre-trial baseline, separately for trials in which the test stimulus contained motion coherence that was stronger vs. weaker than the reference stimulus. Vertical lines indicate the onset of the four events in the trial (reference stimulus, test stimulus, response and feedback) as well as the offset of the two stimuli. Gaps in the timecourse indicate the borders of the four events, which were jittered in time. Black dots show a significant difference between the two types of trials, as determined by a cluster-based permutation test. (**b**) Gamma-band responses during stimulus viewing (150-750 ms after s2 onset). Labels: r_{-1} denotes the previous trial's choice, and s_0 denotes the current stimulus. Bars show group mean $\pm$ s.e.m. See also Figure S5 for each separate visual field map cluster.

found significant modulation by the visual stimulus in all, and significant modulation by the previous choice in dorsal occipital, medial occipital, and posterior parietal regions (Figure S5).

Motor lateralization in the beta-band carries over to the next trial

At the motor end, we computed time-frequency responses based on the motor sensors contra- vs. ipsilateral to the trial's button press (Figure 5a). This showed a typical reduction in the beta-band already during stimulus viewing, building up towards the choice (Donner et al. 2009). This beta lateralization also showed

a strong 'rebound', reversing in sign after the button press. We also observed that the beta-rebound in lateralization was long-lasting, and carried over into the beginning of the next trial (Figure 5b, de Lange et al. 2013; Pape and Siegel 2016). In contrast, lateralization in the gamma range emerged around the time of the button press, and was not sustained after the response (Figure 5a, Donner et al. 2009).

We confirmed these results at the source level. As already visible in the sensor-level TFR (Figure 5b), there was a carry-over effect of the 'beta-rebound' from the last trial (Figure 6a). Pre-trial beta lateralization reflected the previous response (β = 7.6802 [5.4352, 9.9252], p < 0.0001) but not the current response (β = -1.7374 [-3.9806,0.5058], p = 0.1290), nor the interaction between the two (β = -1.2489 [-4.3845, 1.8867], p = 0.4350), as visualized in Figure 6c. This replicates the findings by de Lange et al. 2013; Pape and Siegel 2016, who reported that fluctuations of pre-trial beta-rebound lateralization reflected a carry-over from the previous trial. The gamma-band response, which also showed strong lateralization around the time of the choice (Donner et al. 2009), did not affect the next trial (Figure 6b).

Previous work established that the slope of the motor beta build-up can be used as a read-out for an internally accumulating decision variable (Donner et al. 2009). We quantified this buildup using a linear regression fit on the average beta-band timecourse, from 150 ms to 750 ms after stimulus onset. This was done separately for the previous choice r_{-1} and the current choice r_0. As expected, there was a strong effect of the current choice, with beta-band lateralization reaching opposite endpoints depending on the hand used to indicate the response (Figure 6d).

A repeated measures ANOVA showed a significant effect of current response ($F_{1,59}$ = 70.18, p < 0.001) and a significant effect of previous response ($F_{1,59}$ = 12.39, p = 0.001) without an interaction between the two ($F_{1,59}$ = 0.48, p = 0.489). This was also the case when considering only correct current trials (data not shown). Interestingly, we also saw that the previous choice affected this build-up slope, with steeper build-up on trials that were a repetition of the previous choice (Figure 6e). This suggests that decision-relevant evidence is more efficiently integrated into a motor plan when the choice is a repetition of the previously made choice.

Pre-stimulus sensory and motor signals predict the upcoming choice

Using the beamformed source estimates, we extracted two single-trial oscillatory neural signatures that could be defined during the reference interval, before any decision-relevant information was presented to the subject. First, to capture early visual encoding of the stimulus, we used the gamma-band response (65-95 Hz) in visual cortical areas during presentation of the reference stimulus (150-750 ms). Second, we used the residual 'beta-rebound' from the previous trial to capture the carry-over of motor lateralization (Pape and Siegel 2016). Having identified signatures that reflect the previous choice, we then asked if these signatures could be used to predict the participant's upcoming choice. Using a generalized logistic regression model, we were able to predict each trial's choice based on that trial's gamma-band response during the reference, that trial's motor lateralization in the beta band, as well as the previous trial's choice

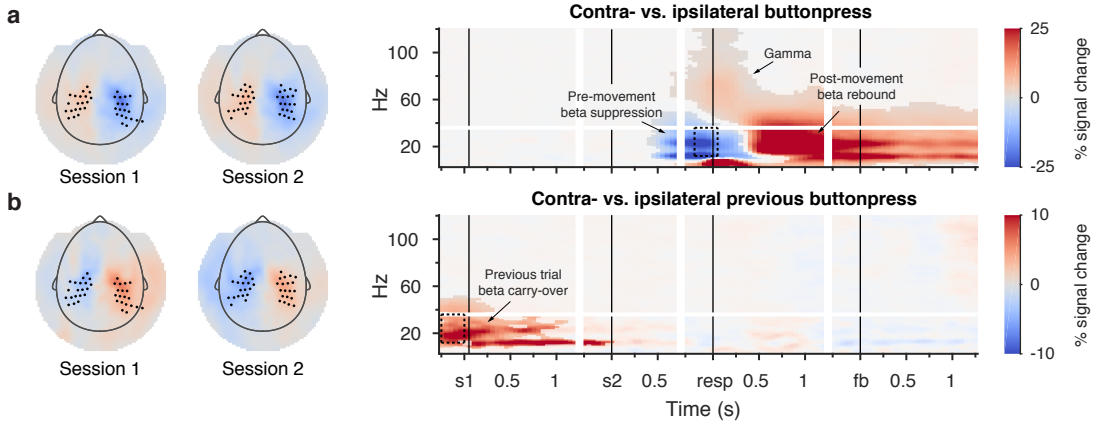

Figure 5. Sensor-level motor lateralization shows known signatures of response preparation and execution. (**a**) Left: sensors selected at the group level, separately for both MEG sessions. Colors indicate the contrast left vs. right button-press in the 250 ms before response, in the 12-36 Hz beta range. Right: overall time-frequency responses of lateralization (contra vs. ipsilateral sensor groups). (**b**) Left: sensors selected at the group level, separately for both MEG sessions. Colors indicate beta-band power for left vs. right button-press in the first 250 ms of the trial. Right: overall time-frequency responses of lateralization with respect to the previous trial's button press (contra vs. ipsilateral sensor groups). In both (**a**-**b**), the dashed box indicates the time- and frequency selection for the topographical representations on the left. Saturation indicates significant time- and frequency points from a cluster-based permutation test.

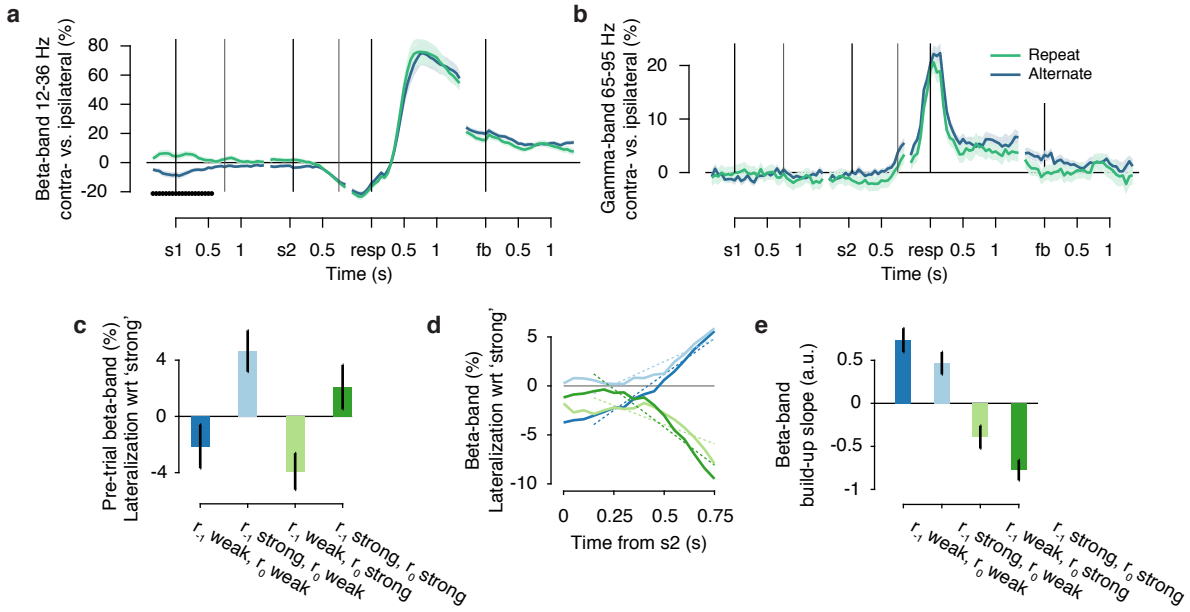

Figure 6. Source-level motor lateralization reflects previous responses. Lateralization timecourses (contra vs. ipsilateral button press on the current choice), for the (**a**) beta (12-36 Hz) and (**b**) gamma (65-95 Hz) range. Timecourses are separately plotted for trials that are a repetition or alternation of the previous choice. Black line shows a significant difference between the two types of trials, as determined by a cluster-based permutation test. Vertical lines indicate the onset of the four events in the trial (reference stimulus, test stimulus, response and feedback) as well as the offset of the two stimuli. (**c**) Beta-band power over left vs. right motor cortex, split by current and previous response. Labels: r_{-1} denotes the previous trial's response, and r_0 denotes the current response. Bars show group mean $\pm$ s.e.m. We used a generalized linear mixed model (see Methods) to test for the effect of previous and current response. (**d**) For each individual, we fit a linear regression on the beta buildup from 150 ms to 750 ms after stimulus onset. The mean fit across observers is shown overlaid on the timecourse. See also Figure S6 for each motor ROI separate. (**e**) Slope of beta-band build-up from 150 to 750 ms after onset of the test stimulus s2, as in (**d**). Labels: r_{-1} denotes the previous trial's response, and r_0 denotes the current response. Bars show group mean $\pm$ s.e.m.

(while allowing for subject-specific intercepts as random factors). We found that, while taking into account both the current trial's stimulus (β = 2.0262 [1.9852, 2.0671], p < 0.0001) and the behavioural response on the previous trial (β = 0.1044 [0.0634, 0.1453], p < 0.0001), both visual gamma-band (β = 0.0004 [0.0001, 0.0007], p = 0.0206) and motor beta lateralization at the start of the trial (β = -0.0004 [-0.0006, -0.0001], p = 0.0022) predicted the subject's upcoming choice. As can be seen from the opposite sign of the two neural measures, higher gamma-band responses preceded more 'stronger' choices, while higher beta-band motor lateralization towards 'stronger' choices precede the opposite 'weaker' choice.

For the visual gamma-band responses we also computed *choice probability*, or the degree to which a signal covaries with the observers' response on that trial. Importantly, we used single-trial gamma-band responses during the reference interval, before decision-relevant information was presented to the observer. Visual gamma-band responses showed small but significant choice probabilities (ROC = 0.5068 $\pm$ 0.0230, t_{60} = 2.3270, p = 0.0234), suggesting that fluctuations in the visual response reflect the observers' upcoming choice even before decision-relevant evidence is presented.

We then asked if these neural signatures of the previous choice also explain the range of individual serial choice biases present across individuals. We computed, for each participant, the effect of previous choices on visual gamma-band responses, motor beta-band lateralization at the start of the trial, and the slope of beta-band build-up. We then correlated each of these measures to individual measures of serial choice bias. There were no significant correlations between individual behavioural measures of serial choice bias and the effect of previous choices on neural signatures of sensory processing, motor preparation and evidence integration (data not shown). One previous study previously reported a correlation between individual repetition behaviour and pre-trial beta-band lateralization (Pape and Siegel 2016). Here, with a larger sample (n = 20 versus n = 60), we did not replicate this correlation with individual serial choice behaviour. We did not observe differences in neural measures of serial choice biases between the three pharmacology groups (all uncorrected p > 0.0899, Mann-Whitney U-test).

Discussion

Here, we used MEG to extract single-trial measure of sensory encoding and motor preparation in sixty-one observers performing a perceptual decision-making task. While subsequent stimuli were independent, observers' choice behaviour was consistently biased by their previously made choices. We found that previous choices affected neural activity at the sensory and motor level, and that both these signatures predicted observers' upcoming decisions.

Firing rates of neurons in visual cortex have shown choice probabilities in the range of 0.5-0.6, increasing up to 0.7 in areas more closely related to the motor response (Britten et al. 1996; Nienborg and Cumming 2006, 2009; reviewed in Crapse and Basso 2015). Here, we observed choice probabilities in visual gamma-band signals in the range of 0.5-0.51, just differing from the 0.5 baseline across our group of observers. This low level of choice probability could have several reasons: single-trial MEG signals contains noise arising

from several non-neural sources, such as physiological noise and head motion. Indeed, we observed significant effects of previous choices on visual gamma-band activity at the source, but not the sensor-level, suggesting that the beamformer improved estimates of the visual response (Hipp and Siegel 2013). Moreover, over and above the already pooled nature of the MEG signal, we combined voxels from a broad set of visual areas (Wandell et al. 2007; Wang and Munoz 2015). Future work could pinpoint cortical regions most predictive of the upcoming choice, increasing the anatomical specificity of our findings. Also, since between-subject variability in gamma-band responses may arise from a number of physiological and and anatomical differences (Shaw et al. 2017) our estimates of sensory encoding could be optimized by determining the gamma-band frequency range for each observer individually. Lastly, it remains an open question to what degree scalp-level electrophysiological signatures allow for the same level of inferential specificity as the firing rate of single neurons in sensory cortex.

Although we did not find a link between algorithmic components of serial choice bias and individual neural signatures between observers, it is tempting to speculate that such a relationship may hold within each participant. Specifically, previous behavioural work suggests a tight link between response preparation and the starting point of evidence accumulation, and between drift bias and the quality of decision-relevant sensory evidence (Leite and Ratcliff 2011; White and Poldrack 2014). This suggests that residual beta-band lateralization from the previous trial may affect the starting point of evidence accumulation, whereas visual gamma-band responses may bias the rate of evidence accumulation towards the previously chosen option. Combined with results from HDDM model fits (Figure S2 and chapter 2.3), which show that drift bias best captures individual choice patterns, this may suggest that the effect of previous choices on sensory cortices may be the dominant factor in driving behavioural serial biases. This would be in line with finding that visual choice probabilities arise at least partly due to feedback connections from higher areas (Nienborg and Cumming 2009; Wimmer et al. 2015); one intriguing avenue for future research is to test to what degree these feedback connections carry information about preceding choices.

A large body of literature has investigated the neural basis of attention and expectations (Summerfield and de Lange 2014). In these experiments, attention and expectations are usually manipulated by adjusting the relative frequency of stimuli, or by cueing observers towards one feature or spatial location. Such cue-induced expectations can activate stimulus-like representational templates in visual cortex (Kok et al. 2013; Kok et al. 2017), suggesting that top-down factors can bias early neural sensory responses. We here observed such choice-predictive activity as a result of the observer's previous choices, rather than experimentally manipulated external variables. This raises the possibility that at least some of the variability observed in neural activity and choice behaviour is not merely noise, but relies on the observer's choice history in systematic ways.

In conclusion, our findings point to previous choices as an important source of variability at the level neural dynamics and the resulting choice behaviour. These results thus contribute to our understanding of exactly how decision processes arise from a rich interplay of sensory information and contextual factors.

Acknowledgements

Christiane Reissmann, Karin Deazle, Samara Green and Lina Zakarauskaite helped with participant recruitment and data acquisition. Roger Zimmermann provided technical assistance, Eline Feenstra scored the personality questionnaires, and Sven Kimenai visually inspected the ECG data. Guido Nolte provided valuable advice and code for the procedure to remove line noise artefacts. This research was supported by the German Academic Exchange Service (DAAD) (to A.E.U.) and the German Research Foundation (DFG), SFB 936/A7, SFB 936/Z1, DO 1240/2-1 and DO 1240/3-1 (to T.H.D.). We acknowledge computing resources provided by NWO Physical Sciences.

Supplementary Figures

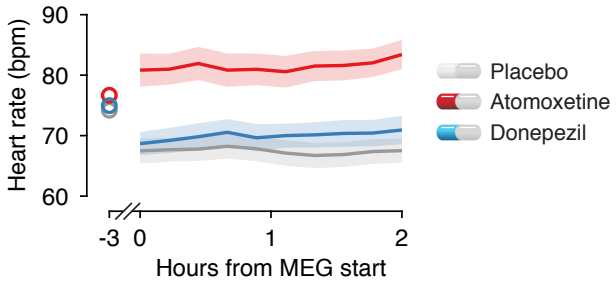

Figure S1. Heart rate increases under atomoxetine, but does not change under donepezil. We used ECG, recorded continuously during the MEG sessions, to test whether pharmacology changed participant's peripheral arousal levels. ECG data were band-pass filtered between 5 and 40 Hz and epoched within each block of trials of approximately 10 minutes. For each participant and block, we then counted the number of peaks in the ECG timeseries, detected based on the constraint that the minimum distance between subsequent peaks was 400 ms (corresponding to a maximum of 150 bpm) and the minimum peak amplitude 0.5 mV. We visually inspected the data to ensure that the algorithm found the R peak of each heartbeat; blocks with insufficient quality for the algorithm to work were discarded. We then computed heart rate in beats per minute (bpm) for each block.

To test whether heart rate changed in response to the pharmacological intervention, we subtracted from the timecourse the participant's heart rate that was measured at the start of the experimental session before receiving their first pill, 3 hours before MEG recording. For statistics, we averaged over all ten MEG blocks. We used a Mann–Whitney U-test to compare the different pharmacological groups, and single-sample Wilcoxon signed rank tests to compare each group's heart rate change from baseline to zero.

Heart rates were stable within the timecourse of the 2-hour MEG session, but the change from baseline (3 hours earlier) depended strongly on pharmacology. The atomoxetine group had a higher heart rate than both the placebo ($p = 0.003$) and donepezil ($p = 0.032$) groups (Wernicke et al. 2003). In contrast to previous reports (Park-Wyllie et al. 2009), we did not observe a significantly different heart rate between the donepezil and placebo groups ($p = 0.590$). Heart rates were lower during MEG scanning than at the beginning of the session in the placebo ($p = 0.0100$) but not the atomoxetine ($p = 0.1311$) or donepezil groups ($p = 0.0840$). This probably results from the timing of the first baseline measurement, which we took immediately after participants arrived in the lab (and therefore after walking or cycling through the hospital campus).

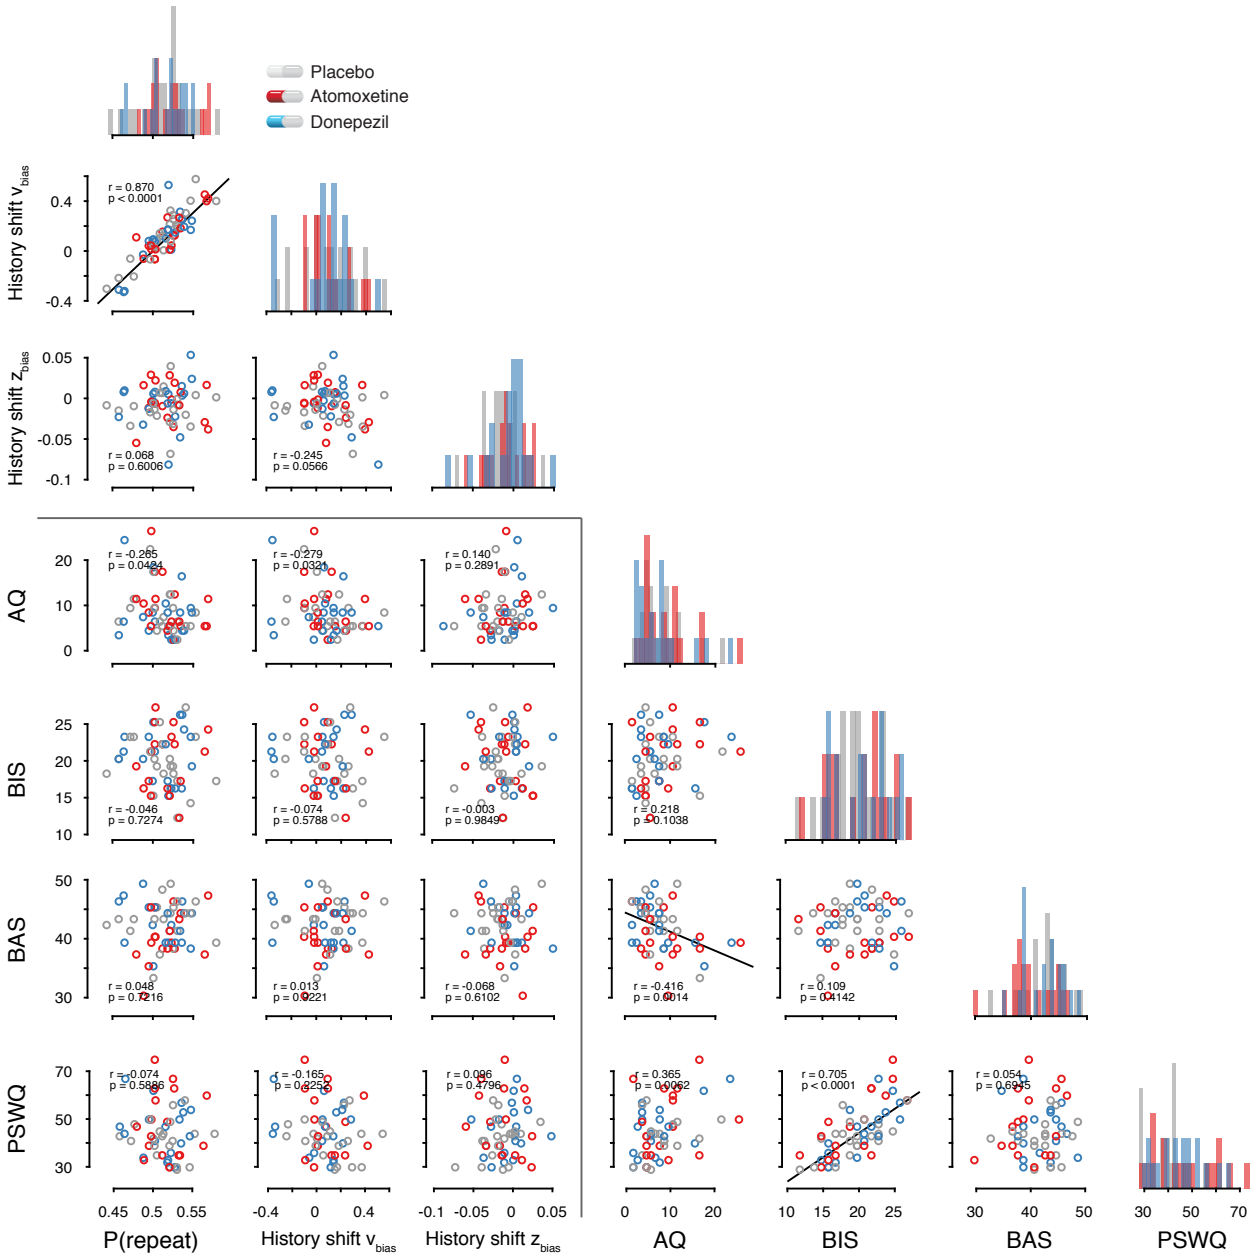

Figure S2. Correlation between serial choice bias, personality scores and pharmacology. Least-squares regression lines are shown for correlations that are significant at the critical FDR-corrected $p = 0.0062$ level (Benjamini and Hochberg 1995). Histograms indicate the distribution of behavioural measures and personality scores for each of the three pharmacological groups. The top-left indicates correlations between behavioural variables, the bottom-left indicates correlations between serial bias and personality scores, and the bottom-right indicates correlations between different personality measures. Colors indicate the pharmacological intervention that each participant received.

Between the four different personality questionnaires, we observed a positive correlation between PSWQ-D (which measures worry) and BIS scores (which measures behavioural inhibition), as previously reported (Muris et al. 1999; Spielberg et al. 2011). We also observed a previously reported correlation between AQ (which measures autistic traits) and BAS scores (which measures behavioural activation) (Gallitto and Leth-Steensen 2015).

None of the personality scores were related to individual measures of serial choice bias. We did not observe significant differences between the pharmacological groups for any of the behavioural or personality variables (Mann–Whitney U test, all uncorrected $p > 0.0703$).

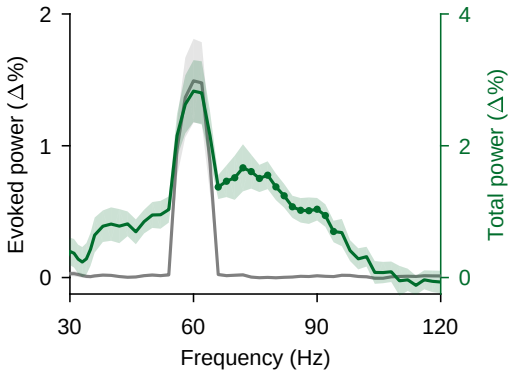

Figure S3. Evoked and total visual responses. High-frequency power spectra of the difference between strong and weak visual motion during the test stimulus. Power was either computed after averaging trials in the time domain (evoked power, grey) or computed on single trials before averaging (total power, green). Power is expressed in % from the average pre-reference baseline. To obtain these sensor-level spectra, we used the occipital electrodes from a split-half sensor definition shown in Figure 3a. Error bars indicate group mean ± s.e.m. Green dots indicate the frequency range of 65-95 Hz used to define the induced gamma-band range.

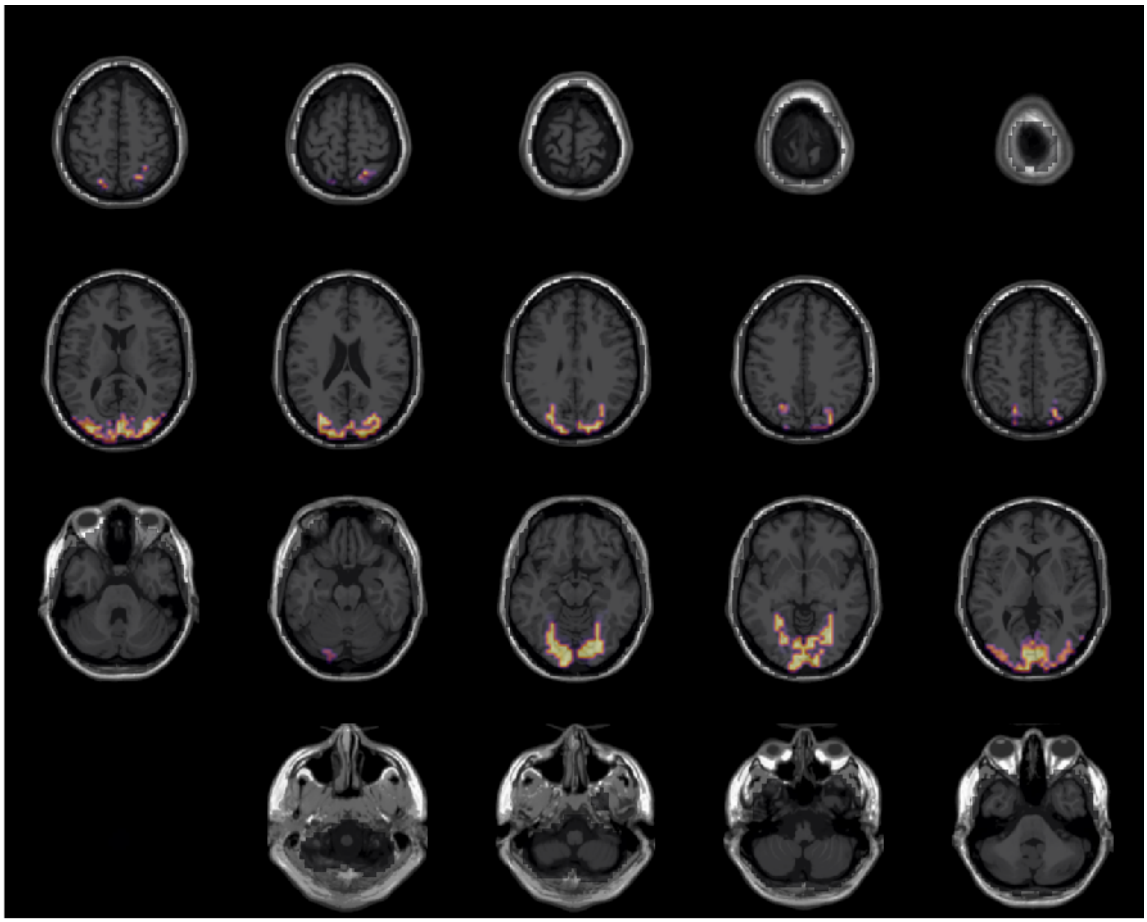

Figure S4. Visual cortical areas used for source reconstruction. We used all visual cortical areas defined in the atlas by Wang and Munoz 2015, and grouped those together into one large visual ROI to maximize power. Shown are the included areas, interpolated onto the Colin27 template brain (Holmes et al. 1998).

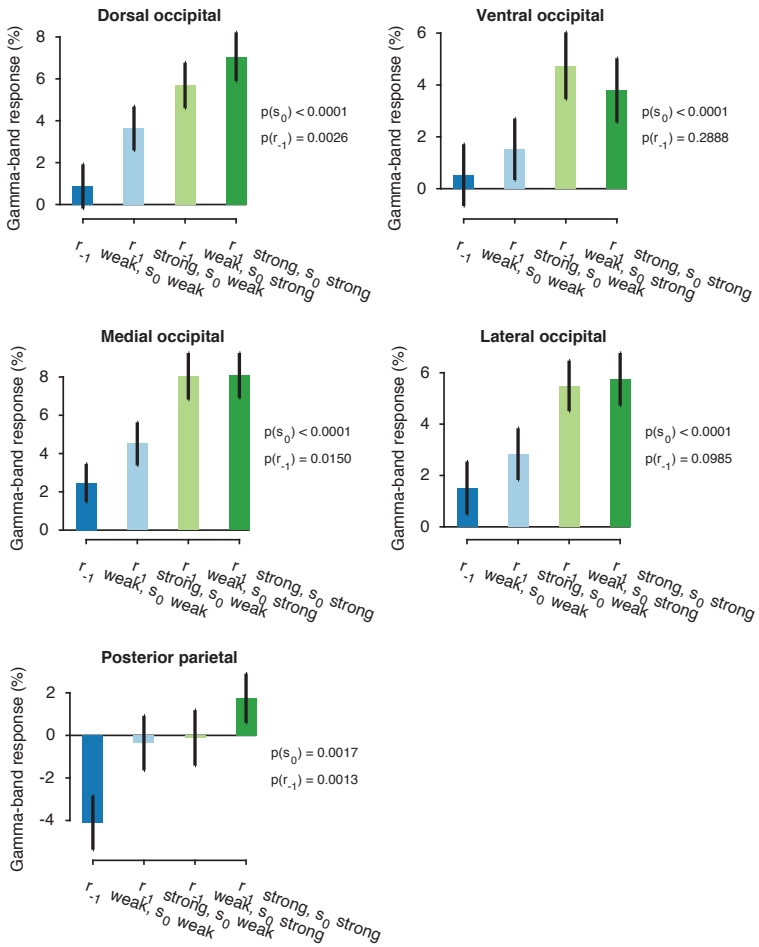

Figure S5. Gamma-band responses for all for all retinotopic visual cortical field areas. We reran the analysis from Figure 4b separately for each visual field map, obtained by combining the areas defined in the atlas by Wang and Munoz 2015 into the clusters specified by Wandell et al. 2007. P-values for the effect of current stimulus s_0 and previous response r_{-1} are printed in each panel.

Region	X	Y	Z
Precentral 'handknob' left	−3.5	−2.6	5.8
Precentral 'handknob' right	3.4	−2.7	5.8
Precentral/superior frontal sulcus left	−3.1	−0.7	4.9
Precentral/superior frontal sulcus right	2.9	−0.7	5.3
Postcentral aIPS left	−3.9	−4.0	5.6
Postcentral aIPS right	4.1	−4.2	5.6

Table S1. MNI coordinates of motor regions. These three regions (both left and right, respectively) were manually defined on the template brain, and interpolated to the nearest MNI coordinate on the voxel grid.

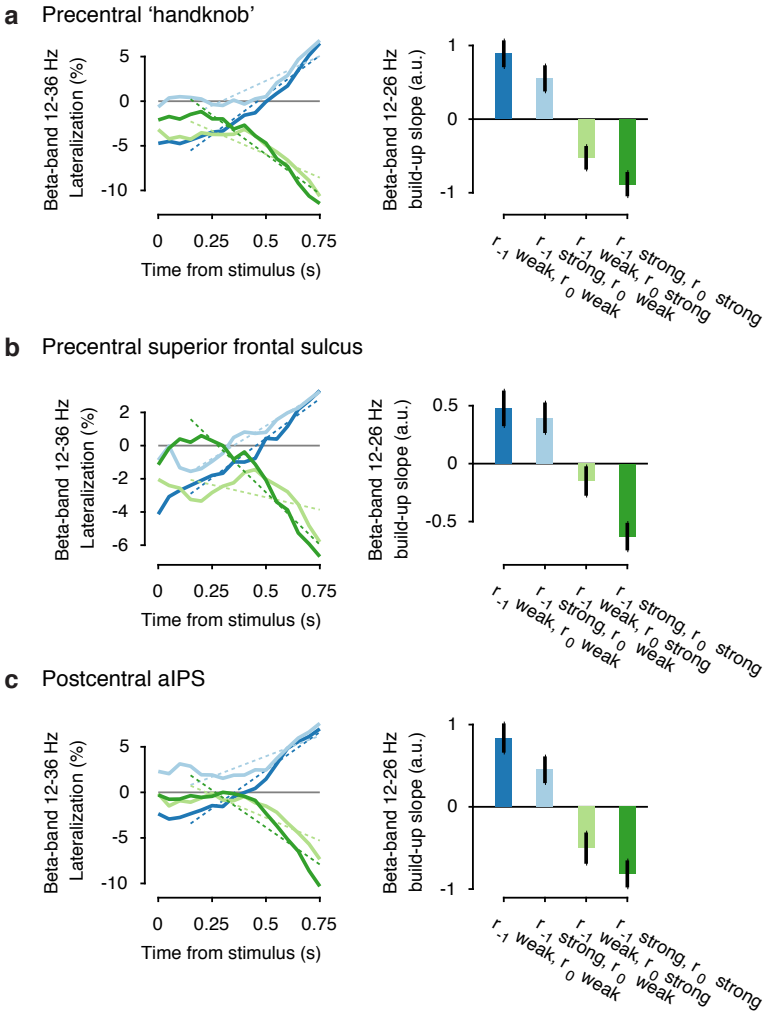

Figure S6. Motor beta build-up slopes separately for each ROI. Dashed lines indicate the average fitted regression line in the interval from 150 to 750 ms after stimulus onset.

2.5 | Choices reduce sensitivity to inconsistent perceptual evidence

Talluri BC*, Urai AE*, Tsetsos K, Usher M, Donner TH.

Abstract

Perceptual inference does not only depend on sensory evidence, but also on top-down factors. One of these is categorical perceptual choice, which might alter subsequent inference in a top-down fashion. Here, we tested the impact of overt perceptual choices on the integration of subsequent sensory evidence into continuous estimation judgments. Observers made a continuous estimation of the average motion direction across two intervals of random dot motion. In addition, an unpredictable auditory cue at the offset of the first interval prompted the observers for a response (button press) halfway through the stream: a fine discrimination judgment based on the first interval of dot motion, or an evidence-independent response. We independently varied the sensory evidence (i.e. motion direction) presented in each interval, and fit statistical models to quantify how observers integrated this sensory evidence into their final estimation of motion direction.

We found that observers decrease their sensitivity to incoming sensory evidence that is inconsistent with their previously made choice. In addition, we found that overt categorical choices reduce overall sensitivity to subsequent evidence (Bronfman et al. 2015). We thus identify two sources of bias in perceptual inference, both of which act to reduce observers' sensitivity to specific incoming sensory evidence.

* Shared first authors. AUTHOR CONTRIBUTIONS: T.H.D. and M.U. formulated the idea for the study and designed the experiment. A.E.U. programmed the task and collected data. B.C.T., A.E.U. and K.T. analyzed data. B.C.T., A.E.U. and T.H.D. wrote the paper. T.H.D. supervised research.

Introduction

Humans and animals make sense of the world around them by performing perceptual inference, combining sensory information with prior knowledge to choose the best possible outcome in a given situation. Traditional models of perceptual decision-making often conceptualize the encoding of sensory information as relying mostly on feed-forward neural mechanisms from low-level sensory areas, and prior knowledge being implemented in the form of top-down modulation using feedback neural mechanisms from higher order areas (Bogacz et al. 2006; Jazayeri and Movshon 2006; Gold and Shadlen 2007; Glaze et al. 2015). Due to this interaction between bottom-up and top-down factors, perceptual inference is influenced by expectations and biases (Rahnev and Denison 2016; Gold and Stocker 2017).

Categorical judgments about sensory evidence are one such biasing factor. In psychophysical experiments, categorical choices bias subsequent decision processes and lead to serial dependencies across trials (Fründ et al. 2014; Braun et al. 2017). Physiologically, categorical choices are more than just the endpoint of the neural inference process: choices can themselves cause feedback signals to sensory areas, thereby altering subsequent perceptual inference (Nienborg and Roelfsema 2015). Specifically, feedback signals from higher choice-related areas contribute to choice-predictive neural activity in early visual areas, a phenomenon termed choice probability (Nienborg and Cumming 2009; Siegel et al. 2015; Wimmer et al. 2015; Goris et al. 2017). When making continuous direction judgments, observers give estimates that are repulsed away from a category boundary used for a binary choice (Jazayeri and Movshon 2007). This behaviour may arise from choice-induced perceptual biases, which are apparent when observers make both a binary choice and a continuous estimation about the same stimulus (Stocker and Simoncelli 2008; Luu and Stocker 2016).

We here set out to quantify the behavioural effects of overt categorical choices on the integration of subsequently presented perceptual evidence. We probed the perceptual inference process by asking observers to report a continuous estimation of perceived motion direction across two stimulus intervals, of which we independently varied the directional evidence (Figure 1A). We then quantified the degree to which each interval's sensory information contributed to the final estimation. After the first interval, observers were prompted to either make a binary judgment about motion direction relative to a reference mark on the screen during the entire duration of a trial (Choice trials), or make a similar motor response that did not require a decision (No-Choice trials). We quantified different effects of making an overt choice on the integration of subsequently presented evidence by fitting statistical models to the continuous motion direction judgements. Moreover, we independently varied the sensory evidence presented in both intervals, and investigated how the consistency between observers' binary choice and the subsequently presented evidence affected their perceptual estimation judgements.

Materials and Methods

Participants

Sixteen participants (six men and ten women, aged 18-29) participated in the study. Two participants did not complete the full experiment, and were discarded from all analyses. The data presented are thus based on the remaining 14 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.

Thresholding

On the first day, observers received initial instructions about the task and performed a thresholding session. Individual observers' motion coherence threshold was determined by a method of constant stimuli on a coarse direction discrimination task (up/down judgment). 600 trials of different difficulties (0, 2.5, 5, 10, 20 and 40% motion coherence), lasting 750 ms, were randomly interleaved. To determine each individual psychometric threshold, we fit a cumulative Weibull as a function of difficulty c, defined as

$$\phi(c) = \delta + (1 - \delta - \gamma)\left(1 - e^{(\frac{-c}{\alpha})^{\beta}}\right) \quad (1)$$

where δ is the guess rate (chance performance), γ is the lapse rate, and α and β are 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 2001a) 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 observers, motion coherence thresholds ranged from 11% to 28% (mean 18%).

Task

After thresholding, observers did a total of 11 experimental sessions across six days. Each session 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 observers acquainted to the task.

Figure 1A shows a schematic of the task. Each trial began with the appearance of a red fixation 'bulls-eye' (Thaler et al. 2013) for 600-800 ms. The first motion interval (X_1) was then displayed for 750 ms, during which the dots moved in one of five directions relative to a reference mark: -20°, -10°, 0° (i.e. aligned with reference mark), 10° or 20°. During display of the motion stimulus, a white line plotted on top of the circle served as a *reference mark*, which changed position from trial to trial. To minimize confusion about the mapping between clockwise/counter clockwise motion and left/right button presses, each subject saw the reference mark in a random position in only the top 0°-180° of the stimulus unit circle, or only in the bottom 180°-360°, and was asked to discriminate motion to the left or to the right of the reference mark.

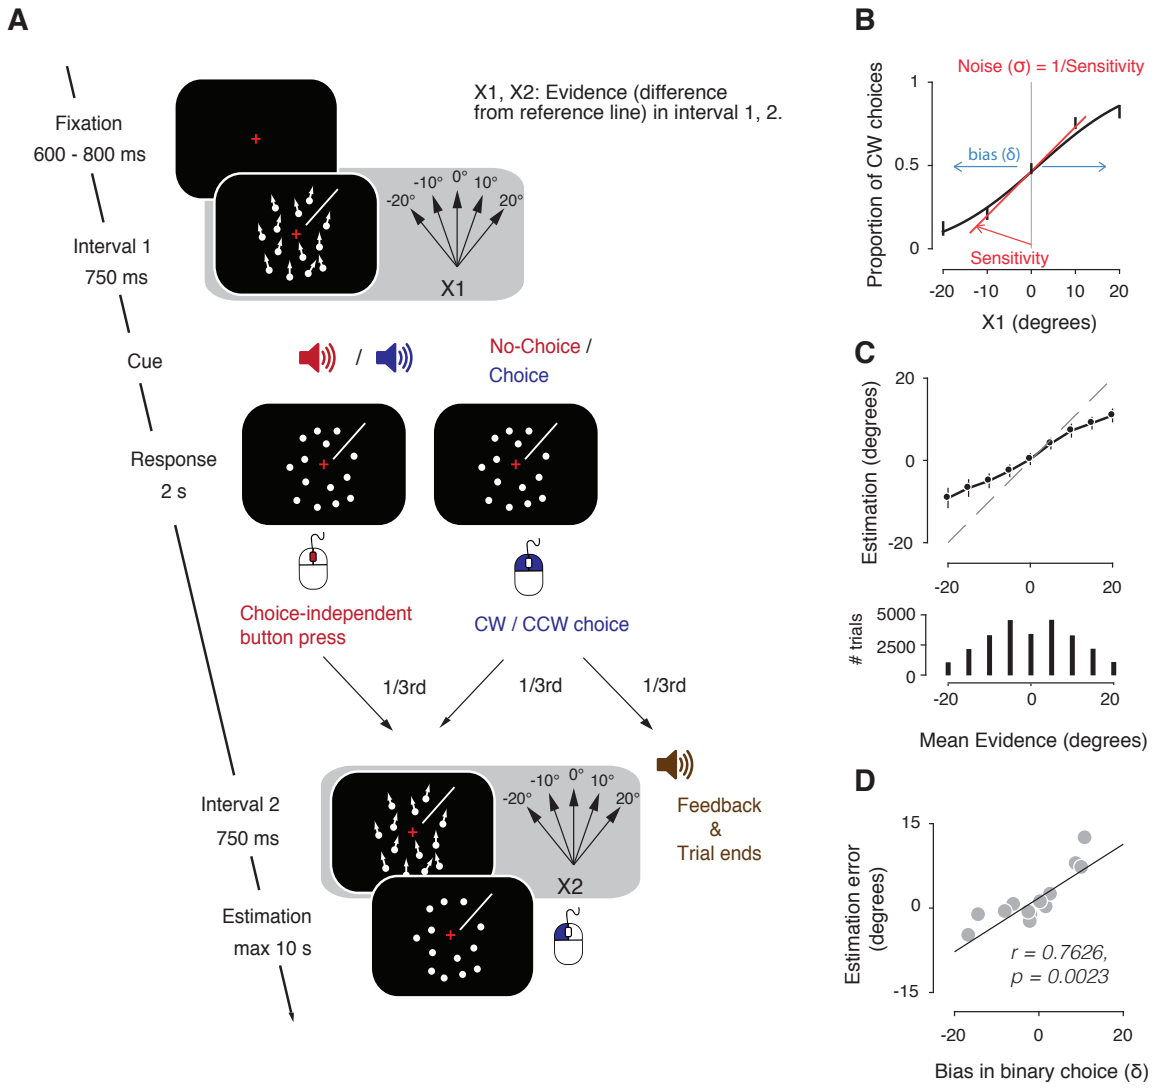

Figure 1. Task and behavioral performance. (**A**) Schematic of the two-interval random dot motion direction estimation task. After the first interval of the motion stimulus X_1, observers received an auditory cue about whether or not to report an overt choice (CW/CCW) about the direction of dots with respect to the reference line. After half of the choice trials (which constituted a third of all trials), auditory feedback was given and the trial terminated. In the other half of choice trials and in No-Choice trials (which together constituted the other two-thirds of all trials), a second interval of moving dots X_2 followed after which observers should estimate the average direction of dots in both the intervals by dragging a line along the direction of average motion. Stimulus directions in X_1 and X_2 were independent of each other and were balanced between Choice and No-Choice trials. In a single block, observers made an overt choice in two-thirds of trials and an estimation in two-thirds of trials. (**B**) Psychometric curve showing the proportion of CW choices for stimulus direction presented in X_1 across observers, along with an illustration of the noise and bias parameters. (**C**) Top: Observers' estimations as a function of mean evidence across the two intervals. The zero-degree direction indicates the reference mark. Deviation from the diagonal shows that observers' estimations showed attraction towards the reference line. Bottom: The distribution of presented mean evidence across all trials and observers. Note that we excluded extremely inconsistent -20°/20°and 20°/-20°trials, leading to a slight decrease in the number of trials with 0° mean (see Figure S1 and Figure S2). (**D**) The mean estimation error across both intervals correlated to bias estimated from the psychometric function on binary choices. N = 14 observers, error bars denote group mean ± s.e.m.

Following X_1 offset, observers received an auditory prompt to either click the central mouse wheel (No-Choice trials, a third of total number of trials) or use the left and right mouse button to indicate whether the dots were moving to the left or to the right of the reference mark (Choice trials, two thirds of the total number of trials). This response window lasted 2 seconds regardless of the reaction time of the subject, ensuring an identical temporal structure between Choice and No-Choice trials.

After half of the choice trials (one-third of all trials), observers received auditory feedback about their binary choice, after which the trial ended. We call these Choice-only trials, and included them to keep observers motivated to perform well on the binary task. If the stimulus angle of motion was at 0° from the reference mark (no angular motion evidence), random feedback was provided. In the other two-thirds of the trials (half of which were Choice, and the other half No-Choice trials), a second motion stimulus interval (X_2) was presented for 750 ms. Following X_2 offset, the reference mark turned red to cue the start of the estimation task. Here, observers estimated the average angle of motion across the first and second intervals by dragging the red line around the circle, and clicking the mouse to confirm their estimate. After the end of estimation, the screen turned black to signify the end of the trial. Observers self-initiated the next trial, to allow enough time to blink their eyes. Feedback about their estimation performance was given at the end of each block as the mean estimation error.

Each trial was defined by the angle of motion (the decision-relevant evidence) in each of the intervals. Twenty-three permutations of the angle pairs for the directions of motion were used with five possible angles (-20°, -10°, 0°, 10°, 20°). The most obviously conflicting direction combinations, -20°/20° and 20°/-20°, were not used. The full experiment presented exactly 90 trials for each set of X_1 and X_2 to the subject, of which 45 in the Choice and 45 in the No-Choice condition.

We used trials where observers made an estimation (Choice and No-Choice trials) in all further analyses. We excluded trials in which observers 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. We also excluded 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 7% of the total trials across observers. The final trial distributions used are shown in Figure S1 and Figure S2.

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. Observers placed their heads on a chinrest 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/degrees^2 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 replotted 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" noise with a "different" rule (Scase et al. 1996). Additionally, to avoid observers 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.

Behavioural performance

We measured each participant's performance on the binary task in Choice and Choice-only trials, by fitting a sigmoidal probit psychometric function (Figure 1B, left panel) on response probabilities as a function of the directional evidence in X_1

$$\mathrm{P}(\mathrm{choice} = CW|\mathrm{X}_1) = \Phi(\delta + \alpha \mathrm{X}_1) \tag{2}$$

where CW was clockwise choice for evidence in X_1, Φ was a cumulative Gaussian, α was the slope of the psychometric function and indicates the observer's perceptual sensitivity, and δ was the horizontal shift of the psychometric function and reflects a perceptual bias towards a specific choice independent of the sensory evidence (see Figure 1B with the slope and bias indicated). The free parameters α and δ were estimated by maximum likelihood estimation. $\sigma = \frac{1}{\alpha}$ gives an estimate of the noise, with which sensory evidence is encoded.

Modeling

For simplicity, we accounted for the effects of bias (δ) and noise (σ) in binary choice using the parameters obtained from the psychometric function in equation 2. Specifically, we replaced the directional evidence (X) in both intervals with $X' = X + \mathcal{N}(\delta, \sigma)$. This accounts for overall perceptual bias (δ) and sensory noise (σ) without adding additional parameters to fit.

Unbiased model

We started with a basic model with three free parameters, two weights accounting for attraction towards the reference and an additional estimation noise. Importantly, this model is unbiased by overt choices or the congruency of sensory evidence. This model was fit so as to obtain a baseline value for the goodness of fit to compare other models.

$$\mathbf{y} = w_1 \mathbf{X_1'} + w_2 \mathbf{X_2'} + \mathcal{N}(0, \xi) \tag{3}$$

where $\mathbf{y}$ is the vector of estimations in all trials, w_1 and w_2 are the weights assigned to corresponding evidence in intervals X_1 and X_2 respectively, $\mathcal{N}(0, \xi)$ is zero-mean Gaussian estimation noise with variance ξ. w_1, w_2 and ξ are the free parameters in the model. w_1 and w_2 capture observers' overall attraction

relative to the reference mark, and ξ captures additional noise in the continuous estimations that cannot be explained by the noise-levels estimated from the psychometric function.

Selective gain model

We modeled a selective change in sensitivity to consistent vs. inconsistent sensory evidence. This selective change in sensitivity can arise in two forms: consistency defined based on the binary choice, or on the sensory evidence presented in both intervals.

Choice-based consistency

Evidence presented following an overt choice can be up-weighted or down-weighted depending on whether the evidence with respect to the reference line is consistent or inconsistent with the overt choice. To model such a selective, consistency related change in sensitivity to evidence following an overt choice, we allowed weights to vary depending on whether evidence in X_2 is consistent (consistent trials) or inconsistent (inconsistent trials) with overt-choice. The estimations are thus given by

$$\mathbf{y} = w_{1cc}\mathbf{X'_{1cc}} + w_{2cc}\mathbf{X'_{2cc}} + \mathcal{N}(0, \xi_{cc}) \qquad sign\,(\mathbf{X_{2cc}}) = \mathbf{D} \tag{4}$$

$$\mathbf{y} = w_{1ic}\mathbf{X'_{1ic}} + w_{2ic}\mathbf{X'_{2ic}} + \mathcal{N}(0, \xi_{ic}) \qquad sign\,(\mathbf{X_{2ic}}) \neq \mathbf{D} \tag{5}$$

where w_{1cc} (w_{2cc}) and w_{1ic} (w_{2ic}) are the weights assigned to noisy evidence in the first (second) interval in consistent trials and inconsistent trials respectively, D is the vector of intermediate binary choice and takes the values $[1, -1]$. Since consistency cannot be defined in trials where evidence in $X_2 = 0°$, we excluded the subset of these trials before fitting the Choice-based consistency model to the Choice trials.

Evidence-based consistency

Consistency can also be defined between evidence in the two intervals. We model estimations as

$$\mathbf{y} = w_{1cc}\mathbf{X'_{1cc}} + w_{2cc}\mathbf{X'_{2cc}} + \mathcal{N}(0, \xi_{cc}) \qquad sign\,(\mathbf{X_{2cc}}) = sign\,(\mathbf{X_{1cc}}) \tag{6}$$

$$\mathbf{y} = w_{1ic}\mathbf{X'_{1ic}} + w_{2ic}\mathbf{X'_{2ic}} + \mathcal{N}(0, \xi_{ic}) \qquad sign\,(\mathbf{X_{2ic}}) \neq sign\,(\mathbf{X_{1ic}}) \tag{7}$$

Since consistency cannot be defined in trials where evidence in either $X_1 = 0°$ or $X_2 = 0°$, we excluded the subset of these trials before fitting the Evidence-based consistency model. The evidence-based consistency model was fit separately on Choice trials and No-Choice trials, and the resulting parameters were averaged to obtain a single consistent and inconsistent measure per subject.

The choice-based consistency model and evidence-based consistency model capture different aspects of a selective change in sensitivity. While evidence-based consistency allows us to investigate changes in sensitivity arising due to an externally induced change in consistency of evidence, choice-based consistency allows us to investigate changes in sensitivity due to internal decision-induced changes in consistency. In addition, evidence-based consistency model can be fit to both Choice and No-Choice trials whereas choice-based consistency model can only be fit to Choice trials.

Global gain model

We also modeled a global, choice-related reduction in sensitivity to evidence following an overt choice, by allowing perceptual weights to vary separately in Choice trials and No-Choice trials (Bronfman et al. 2015). Estimations are given by

$$\mathbf{y} = w_{1c}\mathbf{X_{1c}^{'}} + w_{2c}\mathbf{X_{2c}^{'}} + \mathcal{N}(0, \xi_c) \tag{8}$$

$$\mathbf{y} = w_{1nc}\mathbf{X_{1nc}^{'}} + w_{2nc}\mathbf{X_{2nc}^{'}} + \mathcal{N}(0, \xi_{nc}) \tag{9}$$

where w_1c and w_2c are the weights for the noisy evidence encoded in intervals 1 and 2 in Choice trials. Similarly, w_1nc and w_2nc represent the weights for noisy evidence in the two intervals in No-Choice trials.

Model SNR

The parameters obtained from the model-fitting give information about the weight given to evidence in each interval in the estimations in each trial. However, the weight parameters by themselves need not necessarily measure the sensitivity to evidence in an interval, which is composed of both the weight and noise in the estimation process. To resolve this, we calculated a model signal-to-noise ratio (Model SNR) as a measure of sensitivity to evidence in an interval

$$\mathrm{SNR} = \frac{w}{\sqrt{\sigma^2 + (0.5\xi)^2}} \tag{10}$$

Likelihood computation

We used maximum likelihood estimates to estimate parameters and the goodness of fit of different models. In each trial, we generated Gaussian probability distributions X_1' and X_2' (which incorporated individual bias and noise from the psychometric function, see *Behavioural performance*). These distributions were then multiplied by the respective weight w. In Choice trials, we set the probability of non-chosen side in X_1 to zero and renormalized the probability distribution to make its area 1.

We combined the distributions corresponding to X_1 and X_2 using convolution and renormalized the resulting distribution. To account for estimation noise, we then additionally generated a zero-mean Gaussian probability distribution with variance ξ and combined this distribution with the probability distribution from the previous step, resulting in a distribution of estimations on each trial.

We used this probability distribution to calculate the likelihood of the observer's reported estimation on each 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. This log-likelihood was then optimized to find the best-fitting parameters for each model and observer. Using this analytical method avoids the need to use multiple starting points to obtain the best-fitting parameters.

Fitting procedures

To obtain the best fitting parameters that maximize the likelihood function of each model, we used the subplex algorithm (Rowan 1990; 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 chosen in the interval $[0, 20]$ for ξ and $[0, 1]$ for w_1 and w_2.

Bootstraps

We bootstrapped the computational models of each subject to obtain confidence intervals for the fitted parameters. Specifically, we randomly selected trials with replacement and fit the global gain and selective gain models 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

To ensure that the models we specified could be recovered, we simulated data with different sets of parameters with the number of trials as a typical dataset, and fit the model. This allowed us to confirm that our fitting procedures are able to recover the parameters, with the ground truth of the data known (see Figure S5). Moreover, we verified that all models captured the attraction towards the reference seen in estimations (Figure 1C). To this end, we confirmed that the sum of weights across two intervals as less than 1. Unbiased model: mean across observers = 0.5054, p = 0.00004; Global gain model: mean across conditions and observers = 0.4869, p = 0.0003; Evidence-based selective gain model: mean across conditions and observers = 0.3687, p = 0.00009; Choice-based selective gain model: mean across conditions and observers = 0.4037, p = 0.00005; p values for one-tailed permutation test.

Model selection

We used Bayesian Information Criterion (BIC) to select the model that best explains the data:

$$\mathrm{BIC} = -2\ln(\mathcal{L}) + m \ln(n) \tag{11}$$

where $\mathcal{L}$ is the likelihood value, m is the number of free parameters in the model and n is 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 depend on the number of observations used to fit the model, we fit all the models on the same subset of Choice trials to enable us identify the model that best explains the data.

We calculated BIC values for all individual models in each subject to identify the model that better explains the data for that subject. In addition, we calculated a group BIC value by using the sum of log-likelihood values, the sum of number of trials and the sum of number of parameters for all observers.

ROC analysis for sensitivity to differences in evidence

We assessed the impact of sensory evidence in X_2 on observers' bias-free estimations by means of ROC analyses. The analyses were based on the receiver operating characteristic (Green and Swets 1966), but focused on different effects (the overall effect of choice or consistency). Both quantified the effect of a certain directional difference in evidence on the final estimation distributions. We used the ROC-index (the area under the ROC-curve, Figure 2), to quantify the separability between two distributions of estimations on trials where X_2 differs by 10°. The ROC-index ranges from 0 to 1. An ROC-index of 0.5 implies indistinguishable distributions (i.e., no effect of difference in evidence) and any deviation from 0.5 implies a change in sensitivity of the estimation distributions due to the evidence. Specifically, an ROC-index of 1 or 0 means that two distributions are completely separable, with opposite direction of the shift. By computing ROC indices between sets of trials that differed in their directional evidence, we could assess the sensitivity of the observer in using that input to guide their final estimation reports.

We ran the analyses on all pairs of distributions of directions separated by 10° of evidence in X_2: -20° vs. -10°, -10° vs. 0°, 0° vs. 10°, and 10° vs. 20° (or a subset of those for the selective gain analysis; see below). This gave us four ROC-indices per subject, one index for every pair of distributions compared. We then averaged the resulting ROC-indices across all these comparisons, resulting in a single average ROC index per observer. This average index was finally compared between conditions. To make sure that each trial contributes to the ROC indices only once, we randomly split the trials in distributions which were used more than once (-10°, 0° and 10°) into two halves such that all the pairs of distributions have unique trials. We repeated this process 500 times to obtained confidence intervals for the ROC indices.

Each of the above distributions was made up of trials containing several different directions in X_1. For example, the distribution of estimations produced for $X_2 = -10°$ consisted of trials with all different directions in X_1 (-20°, -10°, 0°, 10°, 20°). We used linear projection to remove the effect of these trial-to-trial variations in X_1 evidence on the final estimation distributions

$$\mathbf{y}' = \mathbf{y} - (\mathbf{y}^{\mathsf{T}}\mathbf{r})\mathbf{r} \tag{12}$$

where $\mathbf{y}$ was the vector of mean-centered estimations, $\mathbf{r}$ was the vector of single-trial evidence directions on X_1 (normalized to unit length), and T denotes matrix transpose. Adding the mean back to the resulting vector $\mathbf{y}'$ yielded the distribution of residual estimations, which was now only affected by the evidence in X_2. Two separate analyses compared the effects of either choice (comparing Choice vs. No-Choice conditions)

or consistency (comparing Consistent vs. Inconsistent conditions) on sensitivity to differences in evidence in X_2. The specifics of these two analyses are described the subsequent sections.

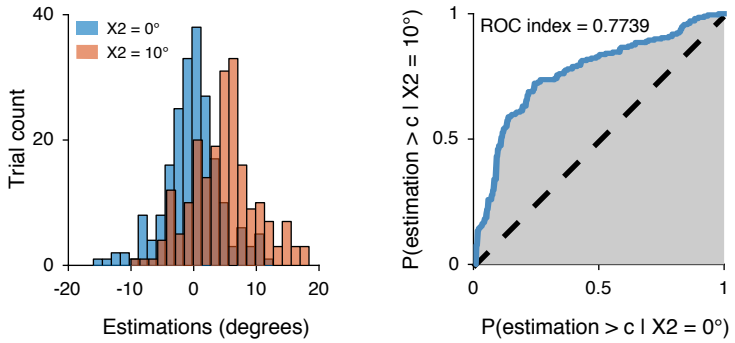

Figure 2. ROC-index. The ROC-index was calculated from the receiver-operating characteristic (ROC) curve. Left: Estimation distributions for two different values of X_2 directions in Choice trials after regressing out the effect of X_1 (equation 12) for an example subject. Right: An ROC curve is constructed by shifting a criterion across both distributions, and for each position of the criterion the fraction of estimations from each distribution that were larger than the criterion were plotted against each other. The area under the resulting curve (grey shaded area in right panel) is the ROC-index (= 0.7739 in this example) that we used to quantify the separability of the two estimation distributions. The dashed line indicates an ROC-index of 0.5, which implies that the distributions are not separable.

To test for an overall reduction in sensitivity due to overt choice, we computed the ROC index separately for Choice and No-Choice trials using the analysis described above. The orthogonal projection (equation 12) was done on both Choice and No-choice trials together. To test for a selective reduction in sensitivity due to consistency, we obtained all Choice trials (or No-Choice trials) and performed the ROC analysis on Consistent and Inconsistent trials separately for X_2 (and X_1, see Figure S3). Here, the orthogonal projection to remove the evidence from one interval on the estimations was done on Choice trials (or No-Choice trials) separately.

For the choice-based selective gain, we defined consistency as between the direction of evidence presented in X_2 and the binary choice. Since consistency in this case could not be defined when the $X_2 = 0°$, we excluded this subset of trials. For the evidence-based selective gain, we defined consistency between the direction of evidence presented in both intervals. Since consistency in this case could not be defined when the directional evidence was 0°, we excluded the subset of trials where the evidence was 0° in either X_1 or X_2. This analysis was done separately for Choice trials and No-Choice trials.

Statistical tests

Non-parametric permutation tests (Efron and Tibshirani 1986) were used to test for group-level significance of individual measures. 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 10,000 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.

Results

Fourteen observers performed a two-interval random dot motion task, consisting of a binary choice or neutral button press halfway through the trial, and a final estimation of the average motion direction at the end (Figure 1A). The trials with choice-independent button press halfway through the trial were used as control trials to compare with the effect of making an overt choice halfway through the trial. On the binary choice task, performance improved for motion directions further away from the reference mark (Figure 1B). Observers were faster responding on Choice trials than on No-Choice trials (Choice RT = 0.68 ± 0.04 s, No-Choice RT = 0.83 ± 0.04 s, difference p = 0.0011; mean of median RT across observers ± s.e.m). This suggests that observers may have needed additional time to press the mouse wheel as they were only prompted to do so in one third of trials.

Observers' mean estimation judgments scaled with the average motion direction, and were biased towards the reference line (Figure 1C). This attraction to the reference mark could have arisen from the structure of our task: across the experiment, observers were presented with fewer trials where the mean evidence was far away from 0°, than trials where the mean evidence was close to 0° (Figure 1C, histogram, see also Figure S1 and Figure S2). Observers may have learned this, and adapted their estimation behaviour accordingly.

Lastly, individual biases in estimation judgments, measured as the mean estimation error in degrees, were correlated with individual bias on the binary choice task, obtained from the psychometric function (Figure 1D). Across observers, these biases were correlated, suggesting they originate from the inference process itself rather than a preference for a specific effector (i.e. index or ring finger used to make the binary choice, compared to just the index finger used to make the estimation response).

Selective down-weighting of inconsistent subsequent evidence

We then asked whether (i) observers change their sensitivity to incoming sensory evidence depending on whether that evidence is consistent or inconsistent with a previously made binary choice, and (ii) observers change their sensitivity to incoming sensory evidence depending on a previously made overt choice, as compared to a choice-neutral motor act. We modeled single-trial estimations as the noisy weighted average of noisy sensory evidence in both intervals by splitting trials into various conditions: Choice trials, No-Choice trials, Consistent trials and Inconsistent trials. We then compared estimated perceptual sensitivity across these conditions to investigate the effect of overt choice or consistency on estimation behaviour.

First, we modeled estimations using an "Unbiased" model that did not include any effects of consistency or choice, and served as a baseline for model comparison. Second, we created an extended "Selective gain" model that allocates different weights to evidence that is consistent or inconsistent with the observers' previously made binary choice. Model comparison shows that across the group, the Selective gain model best explains the data (Figure 3; strong evidence in favor of Selective gain model over Unbiased model across

the group, and in 8 out of 14 individual observers), indicating that consistency of perceptual evidence with a previous choice affects perceptual inference.

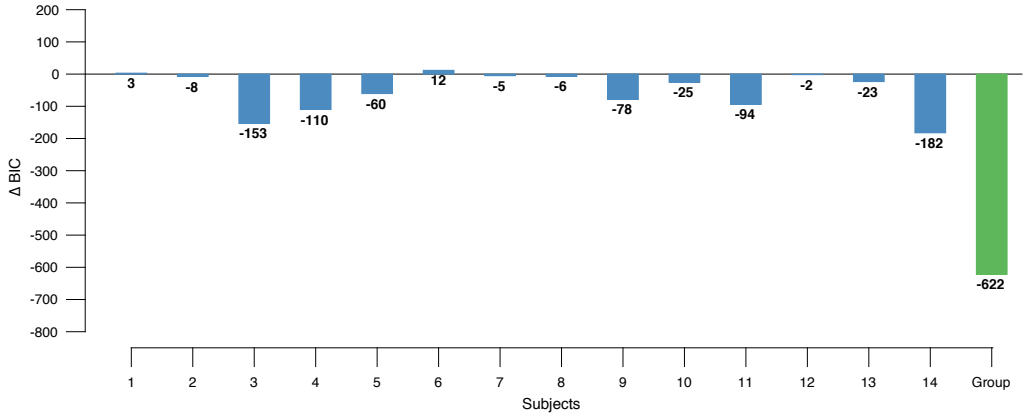

Figure 3. Model comparison. ΔBIC values for Choice-based Selective gain model (6 free parameters) as compared to the Unbiased model (3 free parameters) in Choice trials. We only used trials where $X_2 \neq 0°$, since consistency is undefined for those trials. Negative ΔBIC values indicate evidence in favor of Selective gain model and ΔBIC < -10 indicates strong evidence in favor of the model.

We then compared the estimated weights between consistent and inconsistent trials in the Selective gain model. Observers assigned higher weight to evidence in X_2 that was consistent with their binary choice than to evidence in X_2 inconsistent with their binary choice (Figure 4A, mean difference across observers, X_2 Consistent – X_2 Inconsistent = 0.1693; p = 0.0036).

We designed a similar model, now defining consistency between directional evidence presented in X_1 and X_2. This revealed a similar pattern, showing that inconsistent information was down-weighted (Figure 4A, left panel) compared to consistent information (mean difference across observers, X_2 Consistent – X_2 Inconsistent = 0.0938; p = 0.0166). Importantly, the evidence-based consistency could be fit on both Choice and No-Choice trials (whereas choice-based consistency is undefined on No-Choice trials). These separate fits showed a similar pattern of higher weights to consistent information in both Choice trials and No-Choice trials. Specifically, the weights in No-Choice trials (Figure 4B, right-bottom panel; mean difference across observers, X_2 Consistent – X_2 Inconsistent = 0.1112; p = 0.0115) show that the evidence-based selective effect is present even in the absence of an overt choice. This suggests that subjects may have had an internal estimate of choice in No-Choice trials as they received the prompt about whether or not to make an overt choice after X_1.

Model SNR, a measure of sensitivity to evidence in X_2 (see Model SNR under Modeling in Materials and Methods), also shows a similar pattern as that of model weights in choice-based (Figure 4C) and evidence-based (Figure 4D) selective gain models.

To corroborate these results using a method that does not rely on model fitting, we computed ROC indices that quantify observers' sensitivity to evidence in X_2. ROC indices quantify the degree of separability of two distributions, in this case the estimation distributions whose evidence in X_2 is separated by 10°

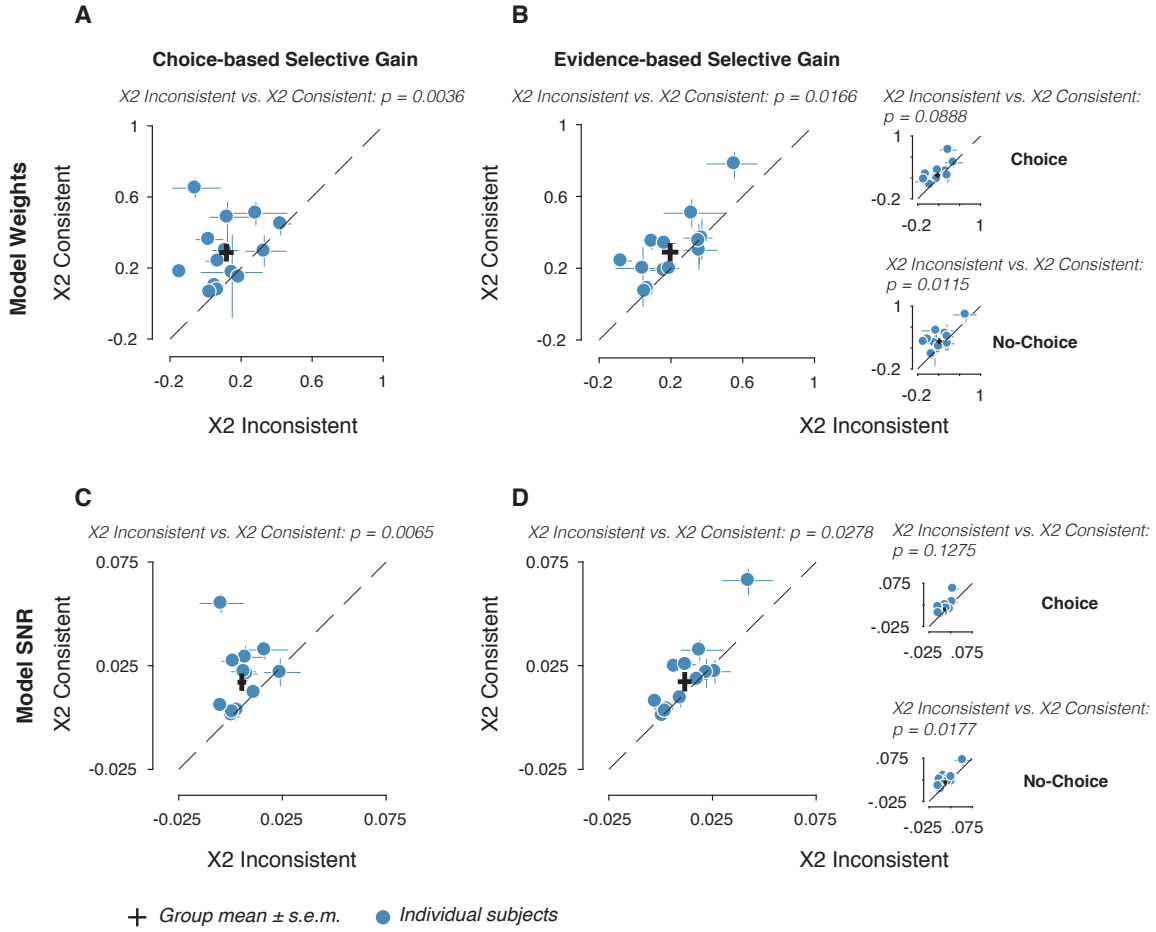

Figure 4. Selective gain model parameters. Model weights (top panel) and signal-to-noise ratio calculated from the model parameters (Model SNR: bottom panel) showing the selective change in sensitivity to evidence in X_2 depending on its consistency to the overt choice (Choice-based selective gain) or to the evidence in X_1 (Evidence-based selective gain). (**A**) Evidence in X_2 consistent with the overt choice is up-weighted compared to evidence inconsistent with the overt choice midway through the trial. (**B**) Evidence in X_2 consistent with the evidence in X_1 is up-weighted compared to the evidence in X_2 inconsistent with the evidence in X_1. Panels on the right shows the parameters computed separately for Choice and No-Choice trials and panel on the left shows the same parameters averaged across Choice and No-Choice trials. (**C-D**) Model SNR corresponding to the Choice-based and Evidence-based selective gain models. Model weights and Model SNR corroborate each other. Error bars show 95% confidence intervals from bootstrapped parameter estimates with 500 iterations in each subject. Dashed line is the identity line. p values were computed using permutation test. N = 14 observers.

after linearly regressing out the effect of evidence in X_1. The ROC indices showed an increased sensitivity in consistent as compared to inconsistent trials, both in Choice-based and Evidence-based selective gain models. This difference however, did not reach statistical significance across the group in Choice-based Selective Gain model. The difference between Consistent and Inconsistent model SNR and ROC indices showed a strong correlation across observers, confirming that these two ways of analyzing the data capture the same pattern (Spearman's correlation = 0.7187, p = 0.0052).

Perhaps surprisingly, observers showed a decrease in sensitivity to evidence that was inconsistent with their previously choice as well as the previously presented evidence; the latter effect was moreover present in

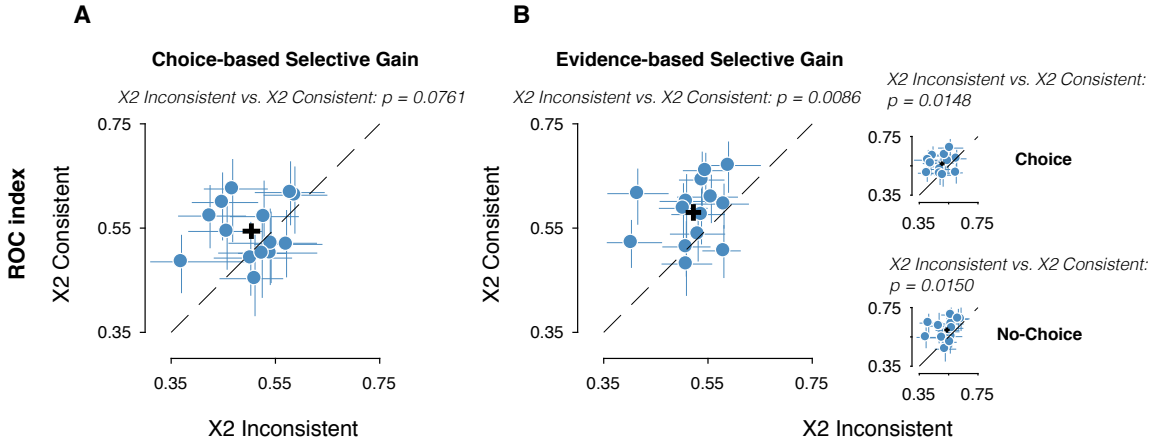

Figure 5. ROC indices for Selective gain model. (**A**) ROC indices showing the selective change in sensitivity to evidence in X_2, depending on its consistency to the binary choice (Choice-based selective gain). Evidence in X_2 consistent with the binary choice has higher sensitivity compared to the evidence in X_2 inconsistent with the binary choice. (**B**) ROC indices showing the selective change in sensitivity to evidence in X_2 depending on its consistency to the evidence in X_1 (Evidence-based selective gain). Evidence in X_2 consistent with the evidence in X_1 has higher sensitivity compared to the evidence in X_2 inconsistent with the evidence in X_1. Panels on the right shows the ROC indices computed separately for Choice and No-Choice trials and panel on the left shows the same ROC indices averaged across Choice and No-Choice trials. Error bars show 95% confidence intervals from bootstrapped ROC indices with 500 bootstraps per subject. Dashed line is the identity line. p values are computed using permutation test. N = 14 observers.

No-Choice trials as well as Choice trials. Do binary choices induce any additional change in the weighting of sensory evidence, over and above that caused by inconsistent evidence? To answer this question, we directly compared the explanatory power of the Choice-based and the Evidence-based Selective gain model, when fit on exactly the same subset of trials (i.e. including those trials were $X_1 \neq 0°$ and $X_2 \neq 0°$). First, we confirmed that on this smaller subset of trials the Choice-based selective gain effect still held (X_2 Consistent - X_2 Inconsistent = 0.1729; p = 0.0058). Second, we compared the BIC values between those two models. The Choice-based Selective gain had a lower BIC for 13 out of 14 observers (and ΔBIC < -10 for 7 out of those 13). The group-level ΔBIC was -324 in favor of the Choice-based over the Evidence-based Selective gain model. Combined with the results on No-Choice trials (Figure 4B, bottom right), this suggests two things. First, inconsistent evidence within a trial leads observers to down-weigh this evidence into their final estimation judgment, regardless of whether they made a Choice or No-Choice judgment halfway through the evidence stream. Second, when making an overt binary choice, consistency with this binary choice more strongly predicts the weighting of subsequent evidence than consistency with perceptual evidence. Choices thus seem to be the dominant, but not only, factor driving the reduction in sensitivity to incoming perceptual evidence.

Down-weighting of subsequent evidence following an overt choice

Earlier work by Bronfman et al. 2015 showed that overt choices reduce sensitivity to subsequent evidence, regardless of its consistency in a numerical integration task. To test this prediction in our data, we first

averaged the ROC indices across Consistent and Inconsistent trials from the Evidence-based selective gain model, separately in Choice and No-Choice trials (Figure 6A).

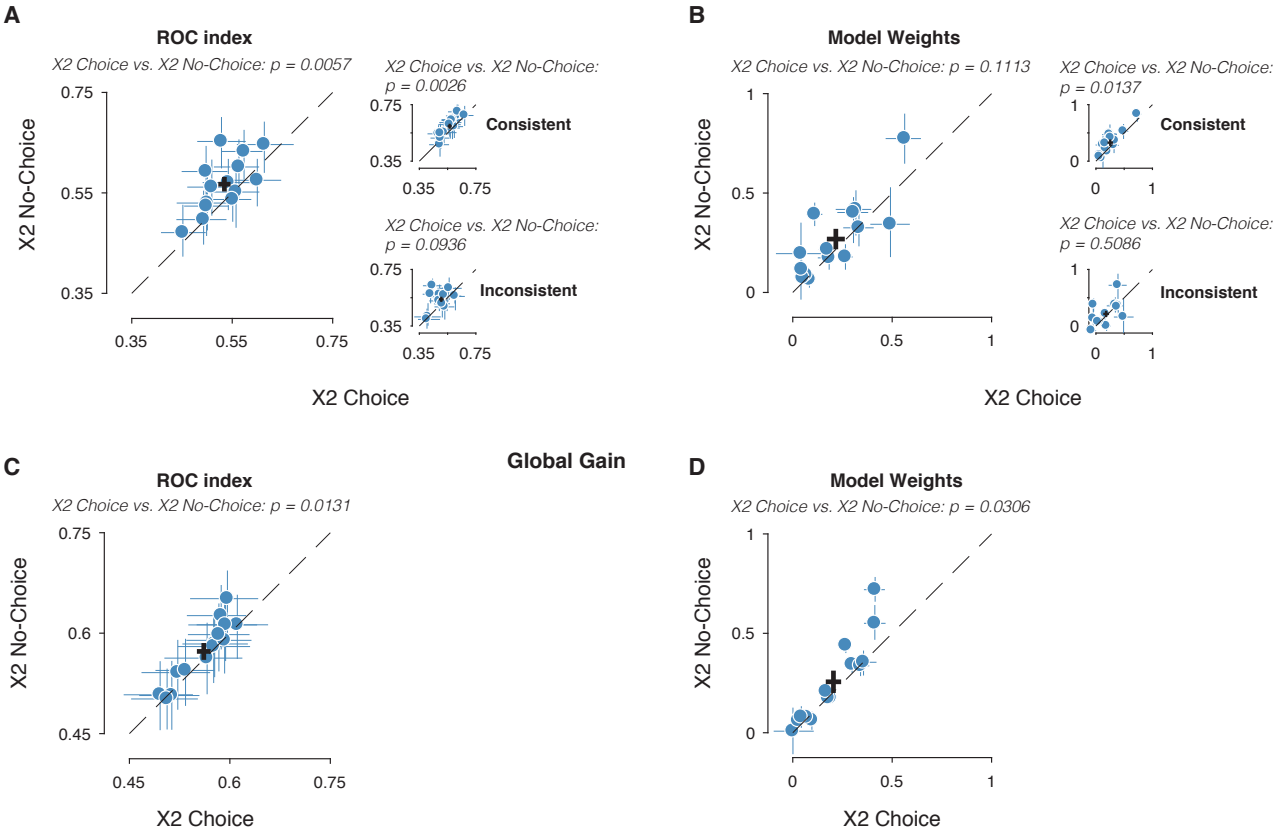

Figure 6. Measures quantifying the effect of overt choice on subsequent behaviour. (**A**) Model-free ROC indices comparing the values from evidence-based selective gain model across Choice and No-Choice trials. Evidence in X_2 following an overt-choice had reduced sensitivity compared to evidence in X_2 following a choice-independent button press. Panels on the right show the ROC indices computed separately for Consistent and Inconsistent trials and panel on the left shows the same ROC indices averaged across Consistent and Inconsistent trials. (**B**) Model weights estimated from the evidence-based selective gain model comparing the Choice and No-Choice conditions. Model weights show a similar trend as that of the ROC indices in (A) but do not reach statistical significance across the group. Panels on the right show the Model weights separately for Consistent and Inconsistent trials and panel on the left shows the same values averaged across Consistent and Inconsistent trials. (**C-D**) Model-free ROC indices and Model weights respectively, quantifying the global gain model. Both measures corroborate each other and show that overt choices reduce overall sensitivity to subsequent evidence. Error bars show 95% confidence intervals from bootstrapped parameter estimates with 500 bootstraps per subject. Dashed line is the identity line. p-values are computed using permutation test. N = 14 observers.

Choice trials showed reduced sensitivity to evidence in X_2 compared to No-Choice trials (Figure 6A). We found that this Choice effect was present in Consistent trials and Inconsistent trials separately.

The averaged model weights from the evidence-based selective gain model did not show this choice effect (Figure 6B). This could be because the evidence-based selective gain model is primarily defined to capture the variance in the data due to changes in consistency and is less sensitive to the Choice effect found in ROC indices. Hence, we fit a separate Global gain model and redid the model-free ROC analysis, separating Choice trials and No-Choice trials irrespective of their consistency. Model-free ROC indices

(Figure 6C) and model weights from the Global gain model (Figure 6D) both show that overt-choices result in reduced sensitivity to evidence in X_2, replicating the findings of Bronfman et al. 2015 in a low-level perceptual task. This suggests that committing to a choice mid-way through a trial reduces overall sensitivity to later evidence.

Comparing the BIC values across different models show that the Global gain model provides a better fit to the data than Unbiased model (Group ΔBIC value on Choice trials where $X_2 \neq 0°$: Global gain model vs. Unbiased model = -138). Moreover, the Choice-based selective gain model provides a better fit to the data than the Global gain model (Group ΔBIC value on Choice trials where $X_2 \neq 0°$: Choice-based selective gain model vs Global gain model = -484). These comparisons suggest that among all models, the Choice-based selective gain model provides a better fit to the data while the Unbiased model provides the worst fit to the data.

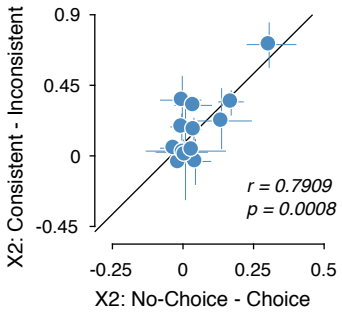

Figure 7. Correlation between Global and Selective reductions in perceptual weighting. For each observer, we quantified the degree to which they show reduced sensitivity to subsequent evidence following an overt choice (x-axis), and the degree to which they reduce their sensitivity to evidence that is inconsistent with their previously made binary choice (y-axis). Across observers, these two choice-induced changes in sensitivity were positively correlated (Pearson's r = 0.7909, p = 0.0008). Error bars show 95% confidence intervals from the model bootstrap.

We observed two types of choice-induced biases in the down-weighting of subsequent evidence: a 'global' effect, which lead observers to reduce their sensitivity to any incoming evidence after an overt choice, and a 'selective' effect, which lead observers to reduct their sensitivity to incoming evidence that was inconsistent with their binary choice. Intriguingly, we observed that these two biases were positively correlated across observers Figure 7. Although this post-hoc analysis has limited power with n = 14, and thus warrants further replication, it is interesting to speculate that both choice-induced biases share a common substrate, with some observers being overall more influenced by their overt binary choice both in comparison to No-Choice trials and in selectively down-weighting evidence inconsistent with that choice.

Discussion

Perceptual inference combines bottom-up sensory evidence with top-down factors. Earlier work on choice-induced biases in perceptual estimation explained these biases as by-products of mechanisms of sensory decoding (Jazayeri and Movshon 2007) or due to the tendency of observers to be self-consistent with their

choices in the absence of external feedback (Stocker and Simoncelli 2008; Luu and Stocker 2016). However, these studies did not explicitly test the effect of choice-induced biases on new sensory information. Here, we quantified these choice-induced biases using an experimental paradigm where new information was presented after the choice and observers were required to use this new information in their estimation judgements. The continuous estimation judgment observers provided is a direct readout of the perceptual inference process where they infer the posterior distribution over the motion directions given the random dot motion input.

Using statistical modeling, we quantified distinct components of the perceptual inference process, and investigated the degree to which they are affected by overt choices. We found that new information that was inconsistent with the observer's intermittent choice in a trial was down-weighted. This selective bias against inconsistent evidence may be a consequence of observers' self-consistency to their choices. As proposed in an earlier work, in the absence of external feedback observers assume their choices to be correct and base subsequent inference on this assumption (Stocker and Simoncelli 2008). This tendency for self-consistency could then give rise to a bias that up-weighs subsequent incoming information that supports this assumption and down-weigh subsequent evidence that contradicts this assumption. Further support for this is our observation that the choice-based selective gain model (defining consistency on the basis of reported intermittent choice) provides a better explanation to the data than the evidence-based selective gain model (defining consistency on the basis of sensory evidence).

Modeling studies have highlighted the role of interaction between top-down and bottom-up signals in perceptual decision making (Glaze et al. 2015; Wimmer et al. 2015). This interaction has also been corroborated by physiological studies in non-human primates (Gold and Shadlen 2007; Nienborg and Cumming 2009; Siegel et al. 2015; Goris et al. 2017). Our findings fit well into this framework where decisions not only influence subsequent decisions (Fründ et al. 2014; St. John-Saaltink et al. 2016; Braun et al. 2017; Fritsche et al. 2017) but also other forms of perceptual inference tasks.

Additionally, we find that overt choices led to an overall reduction in sensitivity to subsequent evidence. This is in line with the finding that overt choices reduce observers' sensitivity to subsequently presented information in both a numerical integration and luminance estimation task (Bronfman et al. 2015). The correspondence between these and our current findings show that choice-induced decreases in perceptual sensitivity occur in decision-making regardless of the precise experimental setup, or the modality of the stimulus to be integrated. This speaks to the generality of our findings, and further suggests that perceptual and value-based decision making can be understood within one common framework (Summerfield and Tsetsos 2012). Together, our findings demonstrate that rather than just being the end-point of a perceptual decision, overt choices themselves are an important context factor biasing subsequent perceptual inference.

Taking this generality one step further, we may speculate that choice-induced biases in sensitivity to sensory evidence may be a low-level form of confirmation bias. Such biases form a broad class (Nickerson 1998), which may or may not arise from a single computational mechanism. These are usually studied in the context of consumer choice, but such studies can be confounded by asking observers for the value of

each option as well as their preference (Chen and Risen 2010). Using low-level sensory stimuli to study such biases offers the advantage of being able to precisely manipulate the identity and reliability of the relevant signal, allowing for a more strictly controlled estimate of biases over and above the known sensory evidence. Moreover, extensive knowledge about the neural substrates of visual perceptual decisions (Nichols and Newsome 2002; Shadlen and Kiani 2013) allows for a natural link to biologically realistic models of perceptual inference in the visual domain (Liu and Wang 2008) which may show signs of choice-induced sensitivity reduction.

Acknowledgements

We thank Ana Vojvodic for help with data collection. This research was supported by the German Academic Exchange Service (DAAD) (to AEU), a Marie Curie Individual Fellowship (to KT) and the German Research Foundation (DFG), SFB 936/A7, SFB 936/Z1, DO 1240/2-1 and DO 1240/3-1 (to THD). We acknowledge computing resources for model fitting provided by NWO Physical Sciences.

Supplementary Figures

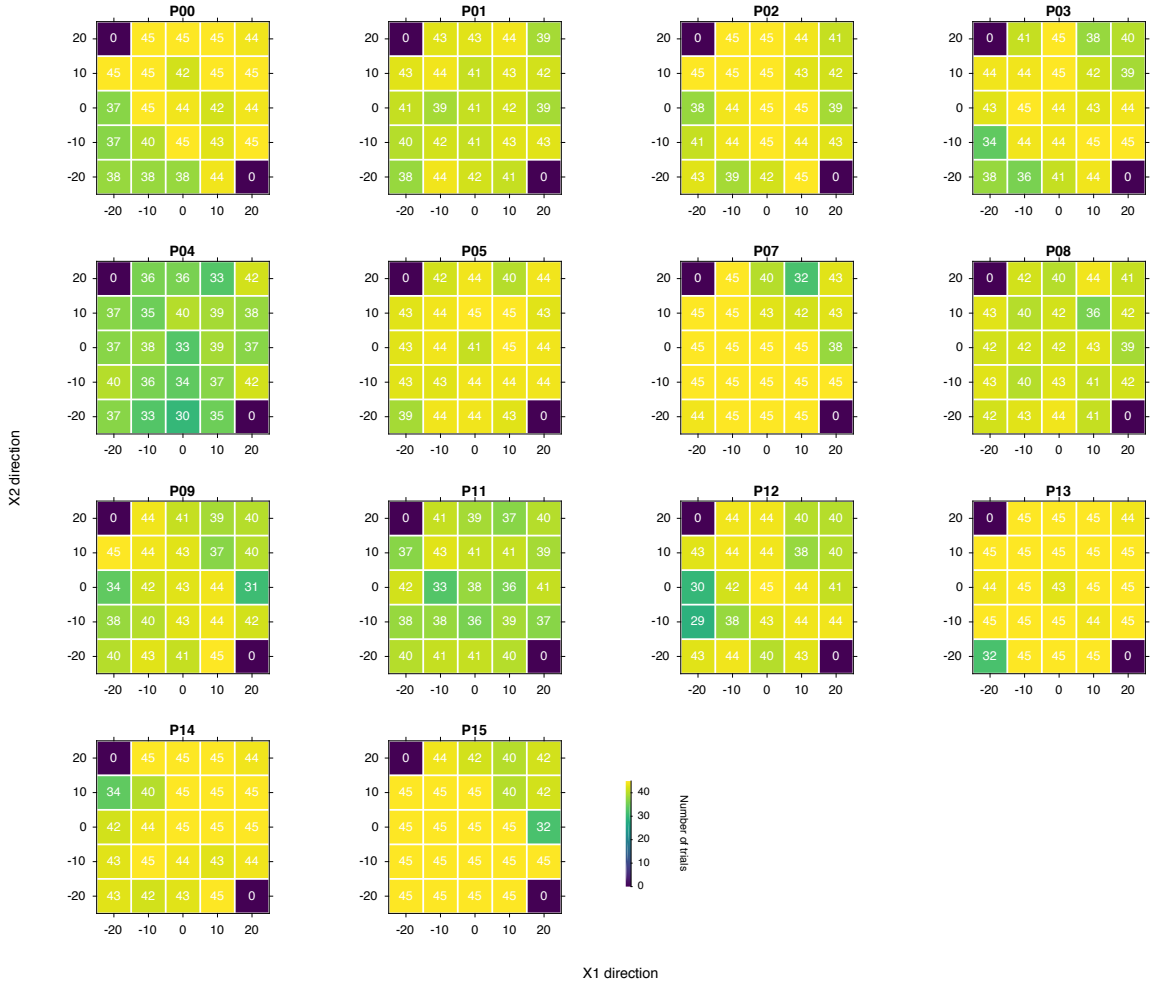

Figure S1. Trial distributions for Choice trials. The distribution matrices of X_1 and X_2 directions in Choice trials for each observer. Each cell in the matrix represents the number of trials presented to the observer for the corresponding values of X_1 and X_2 directions where the observer made a valid response (see Task in Methods). The cells are color coded for better visualization.

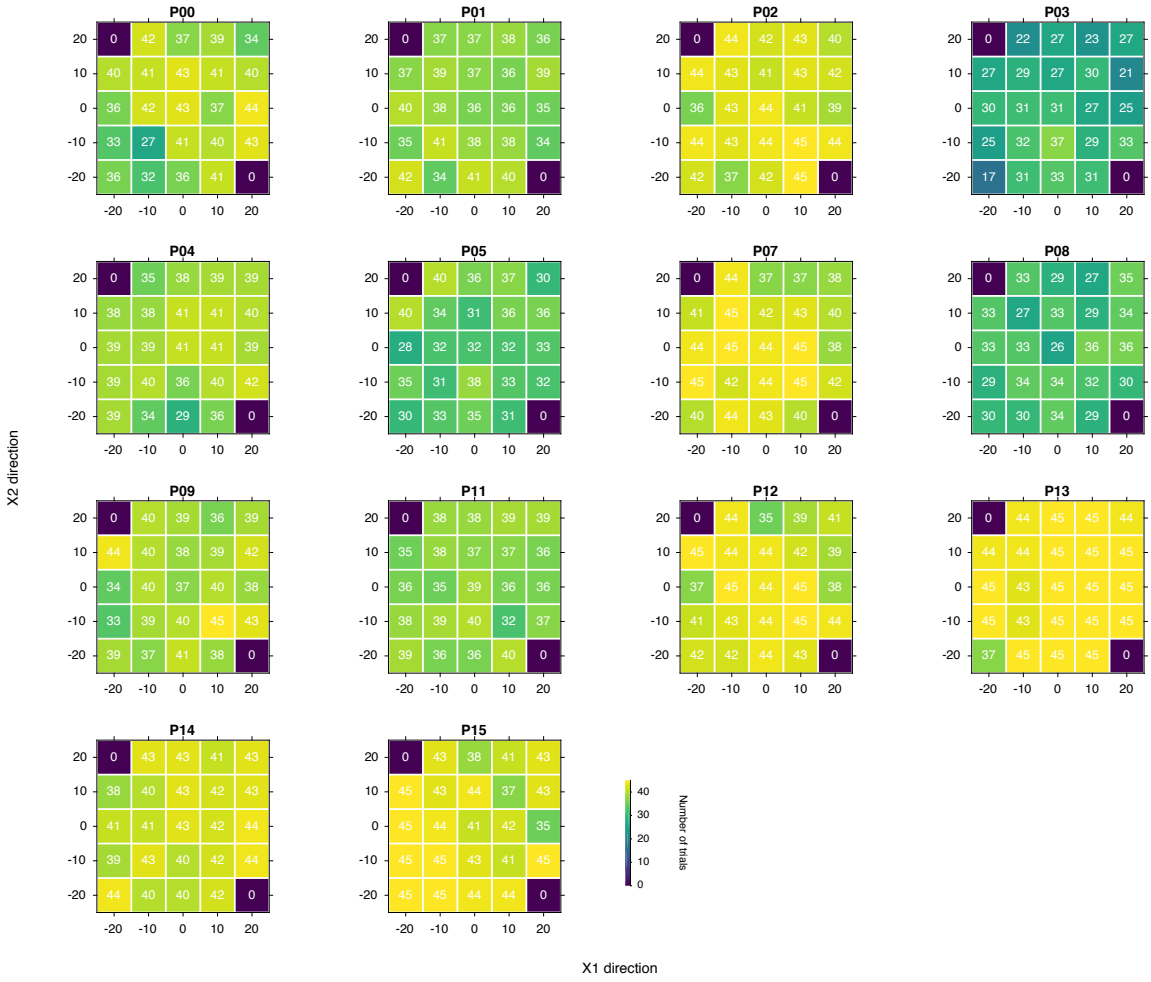

Figure S2. Trial distributions for No-Choice trials. The distribution matrices of X_1 and X_2 directions in Choice trials for each observer. Each cell in the matrix represents the number of trials presented to the observer for the corresponding values of X_1 and X_2 directions where the observer made a valid response (see Task in Methods). The cells are color coded for better visualization.

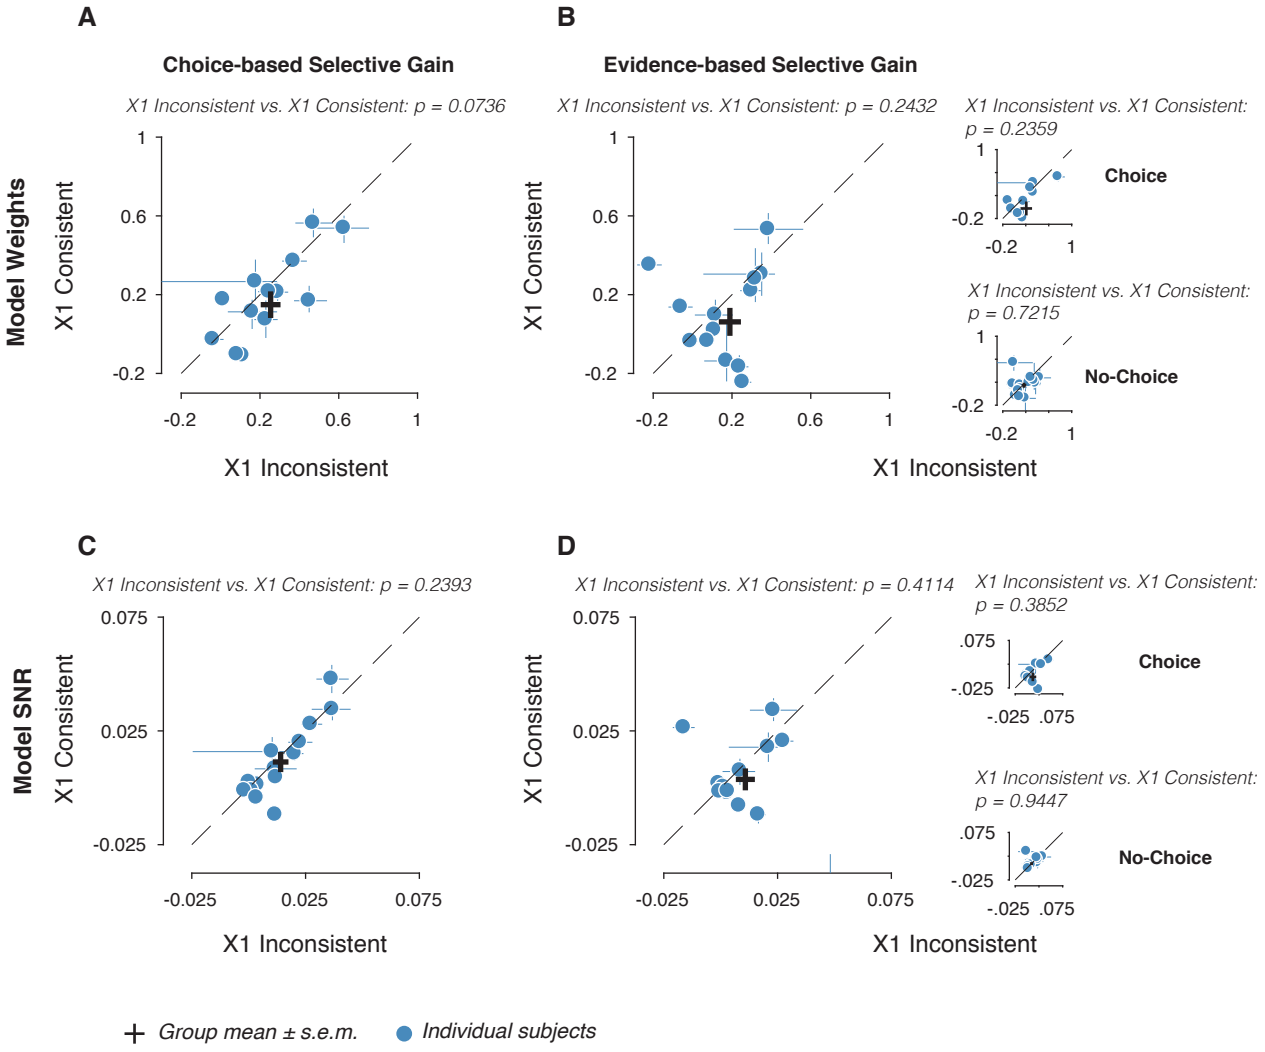

Figure S3. Selective gain model parameters for X_1. Model weights (top row) and Model SNR (middle row) quantifying the effect of choice-based (**A, C, E**) and evidence-based (**B, D, F**) consistency on evidence presented in X_1. Error bars show 95% confidence intervals from bootstrapped estimates with 500 iterations in each subject. Dashed line is the identity line. p values were computed using permutation test; N = 14 observers.

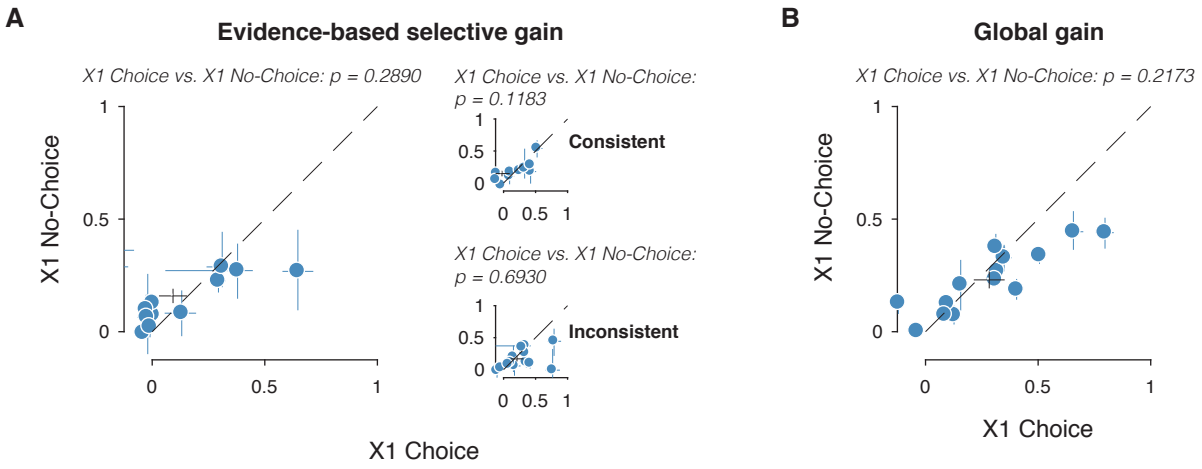

Figure S4. Global gain model parameters for X_1. Model weights showing the effect of overt choice on the estimation of evidence in X_1. (**A**) Values calculated by pooling together the values from evidence-based selective gain model (**B**) Values from the global gain model. Error bars show 95% confidence intervals from bootstrapped estimates with 500 iterations in each observer. Dashed line is the identity line. p values were computed using permutation test. N = 14 observers.

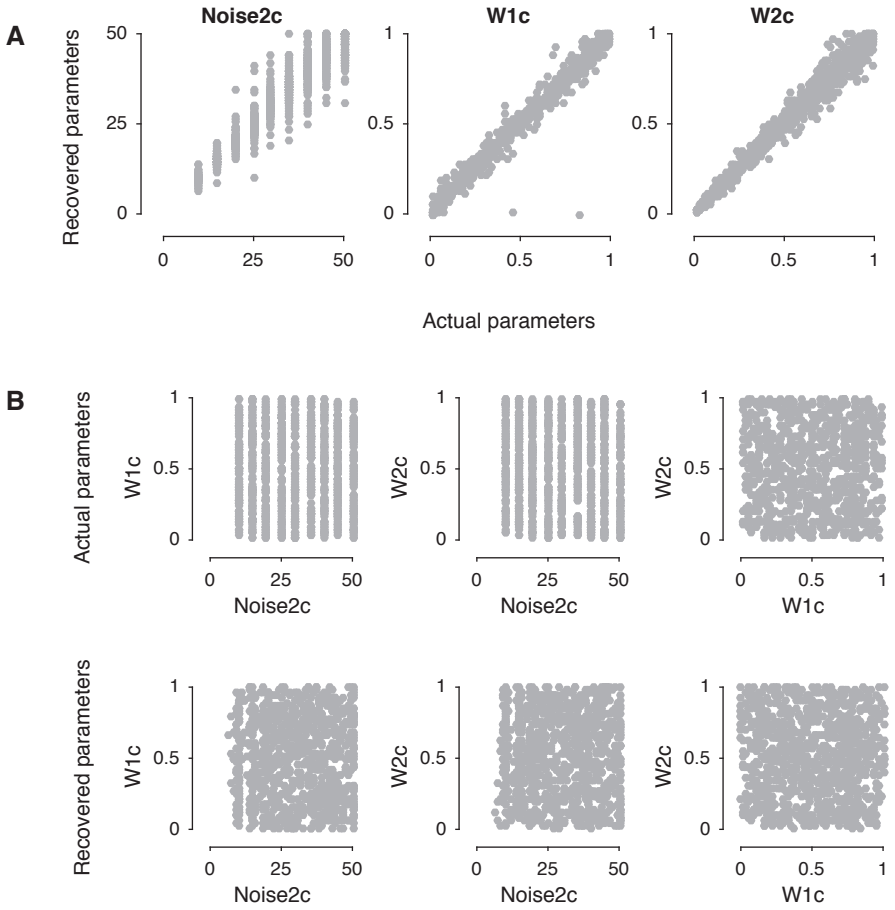

Figure S5. Parameter recovery. (**A**) Actual parameters vs. Recovered parameters in Consistent trials and Inconsistent trials using Selective gain model. The plots show that the model does a good job at recovering the parameters from simulated data. (**B**) Actual parameters were drawn independent of each other and hence show no correlation between each other (top panel). The recovered parameters (bottom panel) do not show any correlations similar to the actual parameters suggesting that the model does not introduce any spurious coupling between the parameters.

3 | Discussion

Was there ever in anyone's life span a point free in time, devoid of memory, a night when choice was any more than the sum of all the choices gone before?

— *Joan Didion*, Run, River

Decisions about the sensory environment do not only depend on the momentary sensory input, but also on the history of preceding choices and stimuli, which robustly biases even highly trained decision-makers. Although such choice biases were first identified in two-alternative forced choice tasks about a century ago, our understanding of the underlying mechanisms remains incomplete. In this thesis, I investigate the computational and mechanistic basis of choice-induced biases in perceptual decision-making.

Chapter 2.1 identifies a link between decision uncertainty, the brain's arousal state, and serial biases in decision-making. When making difficult perceptual decisions, observers' pupils dilated as predicted by a statistical model of decision uncertainty, with the largest pupil dilation after maximally uncertain decisions. Pupil-linked brainstem arousal centers have traditionally been viewed as low-level systems that govern slow and automatic fluctuations of central arousal state; in contrast, our findings add to a growing realization that these systems can rapidly change their activity in accord with high-level computational variables, such as decision uncertainty. Moreover, pupil responses predicted the degree to which observers subsequently change their behavioral bias, reducing an overall tendency to repeat previous choices. These findings pinpoint a novel functional role of decision uncertainty in behavior: modulating serial dependencies between choices.

Chapter 2.2 first investigates whether serial choice patterns are mostly driven by perceptual decisions or the motor responses used to report those decisions. By instructing participants about the stimulus-response mapping on each trial, we decoupled perceptual decisions from the associated motor responses. We found that decisions, but not motor responses, had a strong effect on subsequent choice bias. Second, we report data from an experiment where observers performed a perceptual choice task in 'biased environments', where subsequent stimuli predominantly repeated or alternated. In the absence of external feedback about choice outcome, observers adapted their choice patterns to these environmental statistics. This adaptation was beneficial for performance: observers who adapted most strongly to the stimulus sequence showed the

highest choice accuracy. Extending our findings of uncertainty-modulated serial biases from chapter 2.1, we found that observers' adaptation to biased environmental statistics was reduced by decision uncertainty (as measured by choice correctness and reaction times). These observations indicate that serial choice biases may result from the across-trial accumulation of the internal, graded decision variables on which binary choices are based.

Chapter 2.3 extends the characterization of serial choice bias into the temporal dimension, and differentiates between two possible algorithmic implementations of choice bias. We fit a hierarchical drift diffusion model (DDM) to six psychophysical datasets, allowing both starting point (the offset of evidence accumulation) and drift bias (an evidence-independent constant added to the decision variable at every time step) to vary as a function of previous choices. While these two types of bias can result in the same pattern of choices, they have different effects on response time distributions. Model comparison showed that DDMs with a history-shift (a dependence on the previous choice) in starting point and a history-shift in drift bias were best able to explain most datasets. We then correlated parameter estimates with observers' individual choice patterns, measured as their model-free probability of repeating subsequent choices. Across datasets, repetition probability was robustly correlated with the history-shift in drift bias, but not the history-shift in starting point. This suggests that previous choices primarily bias the rate of evidence accumulation towards a particular choice option. Intriguingly, only for those datasets with biased stimulus sequences did we observe an additional correlation between individual serial biases and their history-shift in starting point. These finding set the stage for further detailed analyses of choice patterns and reaction times in biased environments.

Chapter 2.4 describes data from a large group of observers who performed a perceptual decision-making task while we measured their brain activity using magnetoencephalography. We first replicated the findings described in chapter 2.1, showing that both post-decision pupil responses and reaction times reflect decision uncertainty and reduce choice repetition. We then extracted neural signatures of visual sensory encoding, evidence accumulation and response preparation, and tested whether these neural markers of decision-making were affected by previous choices. At the beginning of the trial, beta-band motor lateralization reflected the previous trial's button press. Visual gamma-band responses, reflecting incoming sensory evidence, also scaled with the previous choice. Both these single-trial neural signatures could be used to predict observers' upcoming choices over and above the sensory evidence and the previous choice. Moreover, the build-up slope of beta-band lateralization during decision formation, a measure of evidence integration, was biased by previous choices. This suggests that previous choices can affect neural activity as early as sensory encoding, possibly through feedback connections from higher decision-making circuits, which may then give rise to serial choice patterns in behaviour.

Chapter 2.5 investigates choice-induced biases in perceptual inference, using a task where observers made continuous perceptual estimation judgments (reporting the average direction of motion across two intervals of moving dots). We independently varied the sensory evidence (i.e. the direction of motion) in those two intervals, and asked observers to make a binary choice about the perceived motion direction (left

or right of a reference mark) after the first interval. Using a computational model of the final perceptual estimation judgment, we quantified observers' sensitivity to information presented in the second interval, in trials where that information was either consistent or inconsistent with the binary choice made before. Observers showed reduced sensitivity to incoming evidence that was inconsistent with their previous choice. Additionally, we quantified sensitivity to incoming evidence on a subset of trials were observers were cued to make a choice-neutral motor act, rather than a binary choice. Observers reduced their overall sensitivity to subsequent information after a choice, as compared to the choice-neutral button press. We thus observed choice-induced biases in perceptual inference, where choices reduce the influence to inconsistent incoming evidence. This may reflect a low-level version of confirmation bias that has been described in value-based choice.

Uncertainty-weighted across-trial accumulation

In neutral (unpredictable) environments and in environments where stimuli consistently tended to repeat or alternate, observers showed uncertainty-dependent serial choice biases: more confident choices had a larger impact on later behaviour (Braun et al. 2017; Urai et al. 2017). This may point to an underlying process of evidence accumulation not just within, but across trials, that gives rise to serial choice patterns and to their dependence on decision uncertainty. Specifically, the integrated decision variable on which binary choices are based can also reflect the degree of confidence in that decision (Kiani and Shadlen 2009; Hebart et al. 2016). Observers may combine these trial-by-trial graded decision variables through a (leaky) across-trial integration process. If a confident trial has a high decision variable, it will strongly influence the state of the accumulator, which then exerts greater bias on subsequent choices. The modulation of serial choice patterns by confidence may thus naturally arise from the graded nature of the decision variables that are accumulated.

Confidence-weighted serial dependencies also arise in a Bayesian model, where each trial's prior depends to a certain degree on the previous trial's posterior (van Bergen 2017, chapter 5). However, whether or not the neural representations underlying uncertainty-dependent perceptual decision-making are based on on fully probabilistic computations remains to be experimentally tested (Rahnev 2017).

Idiosyncratic serial choice patterns

Why are decisions biased by previous choices, even when observers know that stimulus sequences are random? One appealing idea is that observers learn the temporal statistics of real-world visual input, and bring this expectation with them to randomized laboratory experiments. Specifically, observers' choice strategies may have developed to maximize rewards in naturalistic environments, which are rarely composed of discrete, independent consecutive trials. The pervasiveness of local stationarities in real-world sensory data leads observers to accumulate these local patterns, and use them to inform their subsequent decisions.

When put into an artificial laboratory situation with independent trials, this strategy produces maladaptive sequential effects (Yu and Cohen 2008).

A recent model formalized the process of across-trial accumulation, showing that optimal behaviour in dynamic environments requires the integration of decision-relevant evidence with dynamics that depend on the environment's expected instability (Glaze et al. 2015). The timescale of accumulation depends on a parameter called the *hazard rate*, which quantifies an individual's beliefs about the stability of the environment. The hazard rate varies between zero (stability) and one (continuous alternation), with a value of 0.5 reflecting an unpredictable input sequence. The model assumes a fixed hazard rate, which is specified as a free parameter; with the correct hazard rate (i.e. reflecting the generative stability of the input) this model can be shown to produce statistically optimal behaviour. Our adaptation results from chapter 2.2 suggest that observers can (at least partially) adjust their hazard rate to the statistics of the environment. The prevalence of intrinsic serial choice biases in randomized psychophysical experiments can be interpreted within the same framework; observers may use a non-zero subjective hazard rate, that does not exactly match the generative statistics of the task. Specifically, repeaters estimate the environment to be stable, whereas alternators expect the generative input to consistently change.

While this overarching framework for serial choice biases is attractive, it does not specify *how* observers determine their hazard rate and adapt it to the dominant statistics of biased environments. An interesting possibility is that an observers' belief about the stability of the environment is continually estimated from the data itself. In such *dynamic belief models*, individuals assume that the stability of the environment may itself change unannounced, and combine previous observations into a time-varying hazard rate estimate (Yu and Cohen 2008; Zhang et al. 2014). This hazard rate then governs across-trial accumulation of decision-relevant evidence, giving rise to intrinsic serial choice patterns even in uncorrelated experimental contexts. An exciting idea arising from this framework, which remains to be experimentally addressed, is that each observer's serial choice bias might (at least partially) be attributed to the precise sequence of stimuli they encountered before.

Starting point and rate of evidence accumulation

While models of across-trial evidence accumulation may explain how people combine information from previous choices, they leave open the question of how exactly this accumulated information affects subsequent decision dynamics. Most previous studies on serial choice bias have conceptualized and implemented serial biases as an offset (i.e. starting point) of the evidence accumulation process (Cho et al. 2002; Gold et al. 2008; Bode et al. 2012; Goldfarb et al. 2012; Zhang et al. 2014). However, in the framework of sequential sampling, biased choice patterns can also arise from a bias in the rate of evidence accumulation. When simultaneously allowing both starting point as well as a bias in the rate of evidence accumulation to vary with previous choices, we found that individual repetition patterns were more closely related to a bias in the rate of evidence accumulation (chapter 2.3).

One recent study compared the effect of previous choices on both the starting point and rate of evidence accumulation (Kim et al. 2017). They fit the LATER model (which represents the rise of a single accumulator to a bound, Carpenter and Williams 1995) to data from a saccadic reaction time task. Kim et al. found that the starting point of saccade initiation - but not the rate at which the sensory cue was processed to determine the saccade direction - reflected the accumulation of previous decision-relevant information. These findings stand in contrast to the data presented in chapter 2.3, where we observed the most consistent effect of previous choices on the bias in drift, not starting point.

Future studies could directly compare different response modalities (saccades vs button presses) and different sequential sampling models (LATER vs DDM), to see in what cases this leads to different conclusions about the way in which previous choices affect within-trial decision dynamics. It is also interesting to speculate that these two components governing decision dynamics (drift bias vs starting point) may arise from automatic and strategic processes, respectively. Specifically, trial-by-trial updates in starting point agree with predictions from a normative model (Glaze et al. 2015; Kim et al. 2017), and are only correlated to individual choice biases in our datasets with biased transition probabilities (chapter 2.3). History-dependent drift biases, on the other hand, reflect individual choice biases in neutral environments, possibly through history-dependent allocation of selective attention (see below).

An important next step will be to determine the degree to which serial choice biases arise from similar neural and computational processes in neutral and biased environments. Lastly, an important future extension of these ideas will be to investigate serial choice patterns in the context of more flexible and even biologically realistic models of evidence accumulation (Jones et al. 2002; Gao et al. 2009; Bonaiuto et al. 2016).

Choice-induced biases in decision-making

What can these results tell us about decision-making more broadly? It is tempting to speculate that the biases in perceptual decision-making described in this thesis are reflective of cognitive biases more generally. Specifically, *confirmation bias* reflects the tendency to seek out or assign value to new evidence that is compatible with previously held beliefs or choices (Nickerson 1998), and may arise due to *cognitive dissonance* experienced when incoming information conflicts with people's existing beliefs or values (Festinger 1957).

Confirmation bias is most often studied in the context of economic choice; in the 'free-choice paradigm', people assign higher value to an item they have chosen from several equally desirable items. However, since the value of each item can only determined from people's ratings, it can be hard or even impossible to disentangle existing preference differences from bias caused by the choice itself (Chen and Risen 2010; Izuma and Murayama 2013). Studying choice-induced bias in perceptual tasks, where the decision-relevant evidence is exactly known and can be tightly controlled, allowed us to precisely identify the effects of choices on later evidence integration. Specifically, we found that observers' sensitivity to incoming visual evidence depended on the consistency of that evidence with the observer's previous choice. This finding suggests

that some form of confirmation bias may be present in low-level visual perceptual decision-making. This sets the stage for attempts to further bridge theoretical accounts of perceptual and economical decision processes (Summerfield and Tsetsos 2012).

Internal brain states driving serial choice biases

What may be the neural bases underlying serial choice patterns in behaviour? Our finding, that individual serial choice patterns are mainly reflected in a history-dependent shift in the rate of evidence accumulation, is in line with two possible scenarios at the neural level. Previous choices could either affect neural activity at the level of visual processing, or the efficiency with which sensory information is integrated into a choice.

Cortical feedback and belief states

The idea that serial choice biases are reflected at the level of visual cortex is borne out by the findings presented in chapter 2.4, where we describe an effect of previous choices on visual cortical gamma-band responses to decision-relevant motion stimuli. These effects could arise from a local memory in visual circuits, perhaps akin to sensory adaptation (Kohn 2007). It seems unlikely, however, that purely local effects in visual cortex can account for behavioural choice patterns: Akaishi et al. (2014) found that serial biases in behaviour remained intact even when observers viewed strong, irrelevant sensory stimuli between trials, which would presumably interfere with ongoing activity at the level of visual cortex. Moreover, 'choice probability' in visual cortical neurons cannot fully be explained through bottom-up effects originating in visual cortex (Nienborg and Cumming 2009; Wimmer et al. 2015).

A more probable account is that choice-selective signals in visual cortex arise through feedback connections from higher decision-related areas. Such top-down signals may reflect the allocation of selective attention, communicate behavioural relevance, or reflect a cortical 'belief state' about the most likely visual input (reviewed in Nienborg and Roelfsema 2015; Roelfsema and de Lange 2016). Some part of the networks involved in decision-making may keep a memory of the previous choice and feed this information back to lower-level sensory areas when a new stimulus is presented, which then biases the input to the decision process towards the previous choice. This account would predict that the same trial-to-trial variability in visual cortex that reflects previous choices also feeds back into the decision process, producing serial biases in behaviour.

Biased sampling and long-range cortical connections

Serial biases may also arise through the biased sampling of sensory information by higher decision-related brain regions: frontal and parietal cortical circuits may selectively 'read-out' those sensory neural populations encoding the previously chosen stimulus. Long-range cortical connections can gate information flow in cortical networks, flexibly enhancing specific channels conveying sensory evidence (Desimone and Duncan 1995; Siegel et al. 2008; Gilbert and Li 2013; Siegel et al. 2015). Such selective 'read-out' may increase the

rate at which sensory evidence for one over the other choice alternative is sampled and integrated into the decision variable, giving rise to behavioural serial choices patterns.

This idea of history-dependent 'biased sampling', where the brain uses previous choices to selectively allocate attention, is consistent with our findings of a history-dependent change in the rate of evidence accumulation (chapter 2.3) as well as the decrease in perceptual sensitivity to evidence inconsistent with previous choices (chapter 2.5). Even if the input to the evidence accumulation process remains unaffected by choice history, the efficiency with which this evidence is integrated by higher decision circuits may be asymmetrically biased as a function of previous choices.

Central arousal systems

The release of modulatory neurotransmitters from brainstem arousal systems may be an important factor that alters the neural processes underlying serial choice bias. In chapters 2.1 and 2.4, we found that post-decisional pupil responses, a proxy for phasic neuromodulatory activity, reduced trial-by-trial choice repetition. Similar findings were obtained in a dynamic learning task, where pupil responses caused by a sudden change in the environment increased the probability of discarding previously learned information (Nassar et al. 2012; Krishnamurthy et al. 2017).

These findings are line with the idea that the locus coeruleus-noradrenaline system, one of the brainstem centers linked to pupil dilation (Joshi et al. 2016; de Gee et al. 2017), can cause a 'network reset' in cortical circuits, enabling rapid behavioural adaptation (Bouret and Sara 2005; Karlsson et al. 2012). Further work - for example using pharmacological manipulation (chapter 4) or single-trial measurements of brainstem arousal centers (de Gee et al. 2017) - is needed to pinpoint how different neuromodulators interact with information processing in cortical circuits to bring about the neural state changes that produce serial choice bias.

Concluding remarks

In this thesis, I present several new insights into the computational and mechanistic basis of serial choice biases. These biases are pervasive in perceptual decision-making, can be flexibly adapted to environmental demands, and depend not only on previous choices but also the associated uncertainty. Moreover, serial choice biases may arise from biased brain states underlying selective sensory encoding or read-out, which bias the rate of evidence accumulation in line with previous choices.

The precise nature of the neural network dynamics underlying serial choice biases, and their interaction with the brain's arousal systems, remain to be explored. Another exciting question for future research is to assess the nature of individual differences in serial biases, and their possible underlying causes. All in all, the findings presented here start to elucidate how previous choices affect our decisions, in ways that may reflect fundamental principles of decision-making in an uncertain world.

5 | Bibliography

Abrahamyan A et al. (2016). Adaptable History Biases in Human Perceptual Decisions. *Proceedings of the National Academy of Sciences* 113(25):E3548–E3557.

Adelson EH, Bergen JR (1985). Spatiotemporal Energy Models for the Perception of Motion. *Journal of the Optical Society of America A* 2(2):284–299.

Akaishi R, Umeda K, Nagase A, Sakai K (2014). Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia. *Neuron* 81(1):195–206.

Allefeld C et al. (2013). Sequential Dependencies between Trials in Free Choice Tasks. *arXiv* 1311:0753.

Ashby FG (1983). A Biased Random Walk Model for Two Choice Reaction Times. *Journal of Mathematical Psychology* 27(3):277–297.

Aston-Jones G, Cohen JD (2005). An Integrative Theory of Locus Coeruleus-Norepinephrine Function: Adaptive Gain and Optimal Performance. *Annual Review of Neuroscience* 28:403–450.

Baron-Cohen S et al. (2001). The Autism-Spectrum Quotient (AQ): Evidence from Asperger Syndrome/High-Functioning Autism, Males and Females, Scientists and Mathematicians. *Journal of Autism and Developmental Disorders* 31(1):5–17.

Benjamini Y, Hochberg Y (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)*:289–300.

Bennur S, Gold JI (2011). Distinct Representations of a Perceptual Decision and the Associated Oculomotor Plan in the Monkey Lateral Intraparietal Area. *Journal of Neuroscience* 31(3):913–921.

Berens P (2009). CircStat: A MATLAB Toolbox for Circular Statistics. *Journal of Statistical Software* 31(10):1–21.

Bitzer S, Bruineberg J, Kiebel SJ (2015). A Bayesian Attractor Model for Perceptual Decision Making. *PLoS Computational Biology* 11(8):e1004442.

Blair RC, Karniski W (1993). An Alternative Method for Significance Testing of Waveform Difference Potentials. *Psychophysiology* 30(5):518–524.

Bode S et al. (2012). Predicting Perceptual Decision Biases from Early Brain Activity. *Journal of Neuroscience* 32(36):12488–12498.

Bogacz R, Cohen JD (2004). Parameterization of Connectionist Models. *Behavior Research Methods, Instruments, & Computers* 36(4):732–741.

Bogacz R, Wagenmakers EJ, Forstmann BU, Nieuwenhuis S (2010). The Neural Basis of the Speed–accuracy Tradeoff. *Trends in Neurosciences* 33(1):10–16.

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(4):700–765.

Bonaiuto JJ, Berker A de, Bestmann S (2016). Response Repetition Biases in Human Perceptual Decisions Are Explained by Activity Decay in Competitive Attractor Models. *eLife* 5:e20047.

Botvinick MM et al. (2001). Conflict Monitoring and Cognitive Control. *Psychological Review* 108(3):624.

Bouret S, Sara SJ (2005). Network Reset: A Simplified Overarching Theory of Locus Coeruleus Noradrenaline Function. *Trends in Neurosciences* 28(11):574–582.

Braun A, Urai AE, Donner TH (2017). Confidence-Dependent Accumulation of Past Decision Variables Biases Perceptual Choice. *bioRxiv*:172049.

Britten KH et al. (1996). A Relationship between Behavioral Choice and the Visual Responses of

Neurons in Macaque MT. *Visual Neuroscience* 13(01):87–100.

Brody CD, Hanks TD (2016). Neural Underpinnings of the Evidence Accumulator. *Current Opinion in Neurobiology* 37:149–157.

Bronfman ZZ et al. (2015). Decisions Reduce Sensitivity to Subsequent Information. *Proceedings of the Royal Society B: Biological Sciences* 282(1810).

Brooks JL (2012). Counterbalancing for Serial Order Carryover Effects in Experimental Condition Orders. *Psychological Methods* 17(4):600–614.

Buffalo EA et al. (2010). A Backward Progression of Attentional Effects in the Ventral Stream. *Proceedings of the National Academy of Sciences* 107(1):361–365.

Busse L et al. (2011). The Detection of Visual Contrast in the Behaving Mouse. *Journal of Neuroscience* 31(31):11351–11361.

Carpenter RHS, Williams MLL (1995). Neural Computation of Log Likelihood in Control of Saccadic Eye Movements. *Nature* 377(6544):59–62.

Carver CS, White TL (1994). Behavioral Inhibition, Behavioral Activation, and Affective Responses to Impending Reward and Punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology* 67(2):319.

Cavanagh JF, Wiecki TV, Kochar A, Frank MJ (2014). Eye Tracking and Pupillometry Are Indicators of Dissociable Latent Decision Processes. *Journal of Experimental Psychology: General* 143(4):1476–1488.

Chamberlain SR et al. (2009). Atomoxetine Modulates Right Inferior Frontal Activation During Inhibitory Control: A Pharmacological Functional Magnetic Resonance Imaging Study. *Biological Psychiatry* 65(7):550–555.

Chen MK, Risen JL (2010). How Choice Affects and Reflects Preferences: Revisiting the Free-Choice Paradigm. *Journal of Personality and Social Psychology* 99(4):573–594.

Cho RY et al. (2002). Mechanisms Underlying Dependencies of Performance on Stimulus History in a Two-Alternative Forced-Choice Task. *Cognitive, Affective, & Behavioral Neuroscience* 2(4):283–299.

Crapse TB, Basso MA (2015). Insights into Decision Making Using Choice Probability. *Journal of Neurophysiology* 114(6):3039–3049.

Cross DV (1973). Sequential Dependencies and Regression in Psychophysical Judgments. *Perception & Psychophysics* 14(3):547–552.

Dayan P, Kakade S, Montague PR (2000). Learning and Selective Attention. *Nature Neuroscience* 3:1218–1223.

Dayan P, Yu AJ (2006). Phasic Norepinephrine: A Neural Interrupt Signal for Unexpected Events. *Network: Computation in Neural Systems* 17(4):335–350.

de Berker AO et al. (2016). Computations of Uncertainty Mediate Acute Stress Responses in Humans. *Nature Communications* 7:10996.

de Gee JW, Knapen T, Donner TH (2014). Decision-Related Pupil Dilation Reflects Upcoming Choice and Individual Bias. *Proceedings of the National Academy of Sciences* 111(5):E618–E625.

de Gee JW et al. (2017). Dynamic Modulation of Decision Biases by Brainstem Arousal Systems. *eLife* 6:e23232.

de Lafuente V, Jazayeri M, Shadlen MN (2015). Representation of Accumulating Evidence for a Decision in Two Parietal Areas. *Journal of Neuroscience* 35(10):4306–4318.

de Lange FP, Rahnev DA, Donner TH, Lau H (2013). Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *Journal of Neuroscience* 33(4):1400–1410.

Desimone R, Duncan J (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience* 18(1):193–222.

Donders FC (1868). Over de snelheid van psychische processen. *Onderzoekingen gedaan in het Physiologisch Laboratorium der Utrechtsche Hoogeschool (1968–1869)* 2:92–120.

Donner TH, Nieuwenhuis S (2013). Brain-Wide Gain Modulation: The Rich Get Richer. *Nature Neuroscience* 16(8):989–990.

Donner TH, Siegel M, Fries P, Engel AK (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology* 19(18):1581–1585.

Dutilh G et al. (2012). How to Measure Post-Error Slowing: A Confound and a Simple Solution. *Journal of Mathematical Psychology* 56(3):208–216.

Ebitz RB, Platt ML (2015). Neuronal Activity in Primate Dorsal Anterior Cingulate Cortex Signals Task Conflict

and Predicts Adjustments in Pupil-Linked Arousal. *Neuron* 85(3):628–640.

Edwards W (1965). Optimal Strategies for Seeking Information: Models for Statistics, Choice Reaction Times, and Human Information Processing. *Journal of Mathematical Psychology* 2(2):312–329.

Efron B, Tibshirani R (1986). Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. *Statistical Science* 1(1):54–75.

Einhäuser W, Stout J, Koch C, Carter O (2008). Pupil Dilation Reflects Perceptual Selection and Predicts Subsequent Stability in Perceptual Rivalry. *Proceedings of the National Academy of Sciences* 105(5):1704–1709.

Eldar E, Cohen JD, Niv Y (2013). The Effects of Neural Gain on Attention and Learning. *Nature Neuroscience* 16(8):1146–1153.

Fechner GT (1860). *Elemente der Psychophysik*. Leipzig: Breitkopf und Härtel.

Fernberger SW (1920). Interdependence of Judgments within the Series for the Method of Constant Stimuli. *Journal of Experimental Psychology* 3(2):126.

Festinger L (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.

Filimon F et al. (2013). How Embodied Is Perceptual Decision Making? Evidence for Separate Processing of Perceptual and Motor Decisions. *Journal of Neuroscience* 33(5):2121–2136.

Fischer J, Whitney D (2014). Serial Dependence in Visual Perception. *Nature Neuroscience* 17(5):738–743.

Freitag CM et al. (2007). Evaluation Der Deutschen Version Des Autismus-Spektrum-Quotienten (AQ) - Die Kurzversion AQ-K. *Zeitschrift für Klinische Psychologie und Psychotherapie* 36(4):280–289.

Fritsche M, Mostert P, de Lange FP (2017). Opposite Effects of Recent History on Perception and Decision. *Current Biology* 27(4):590–595.

Fründ I, Wichmann FA, Macke JH (2014). Quantifying the Effect of Intertrial Dependence on Perceptual Decisions. *Journal of Vision* 14(7):9–9.

Gallitto E, Leth-Steensen C (2015). Autistic Traits and Adult Attachment Styles. *Personality and Individual Differences* 79:63–67.

Gao J, Tortell R, McClelland JL (2011). Dynamic Integration of Reward and Stimulus Information in Perceptual Decision-Making. *PLoS ONE* 6(3):e16749.

Gao J et al. (2009). Sequential Effects in Two-Choice Reaction Time Tasks: Decomposition and Synthesis of Mechanisms. *Neural Computation* 21(9):2407–2436.

Gilbert CD, Li W (2013). Top-down Influences on Visual Processing. *Nature Reviews Neuroscience* 14(5):350–363.

Glaze CM, Kable JW, Gold JI (2015). Normative Evidence Accumulation in Unpredictable Environments. *eLife* 4:e08825.

Glimcher PW (2005). Indeterminacy in Brain and Behavior. *Annual Review of Psychology* 56(1):25–56.

Glimcher PW (2011). Understanding Dopamine and Reinforcement Learning: The Dopamine Reward Prediction Error Hypothesis. *Proceedings of the National Academy of Sciences* 108:15647–15654.

Gold JI, Law CT, Connolly P, Bennur S (2008). The Relative Influences of Priors and Sensory Evidence on an Oculomotor Decision Variable During Perceptual Learning. *Journal of Neurophysiology* 100(5):2653–2668.

Gold JI, Shadlen MN (2002). Banburismus and the Brain. *Neuron* 36(2):299–308.

Gold JI, Shadlen MN (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience* 30(1):535–574.

Gold JI, Stocker AA (2017). Visual Decision-Making in an Uncertain and Dynamic World. *Annual Review of Vision Science* 3:227–250.

Goldfarb S et al. (2012). Can Post-Error Dynamics Explain Sequential Reaction Time Patterns? *Cognitive Science* 3:213.

Good IJ (1979). Studies in the History of Probability and Statistics. XXXVII A. M. Turing's Statistical Work in World War II. *Biometrika* 66(2):393–396.

Goris RLT et al. (2017). Dissociation of Choice Formation and Choice-Correlated Activity in Macaque Visual Cortex. *Journal of Neuroscience* 37(20):5195–5203.

Gottlieb J (2012). Attention, Learning, and the Value of Information. *Neuron* 76(2):281–295.

Green DM, Swets JA (1966). *Signal Detection Theory and Psychophysics*. John Wiley and Sons.

Gross J et al. (2001). Dynamic Imaging of Coherent Sources: Studying Neural Interactions in the Human

Brain. *Proceedings of the National Academy of Sciences* 98(2):694–699.

Harris KD, Thiele A (2011). Cortical State and Attention. *Nature Reviews Neuroscience* 12:509–523.

Hebart MN, Hesselmann G (2012). What Visual Information Is Processed in the Human Dorsal Stream? *Journal of Neuroscience* 32(24):8107–8109.

Hebart MN, Schriever Y, Donner TH, Haynes JD (2016). The Relationship between Perceptual Decision Variables and Confidence in the Human Brain. *Cerebral Cortex* 26(1):118–130.

Hebb DO (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley.

Hipp JF, Siegel M (2013). Dissociating Neuronal Gamma-Band Activity from Cranial and Ocular Muscle Activity in EEG. *Frontiers in Human Neuroscience* 7.

Hoeks B, Ellenbroek BA (1993). A Neural Basis for a Quantitative Pupillary Model. *Journal of Psychophysiology* 7:315–324.

Holmes CJ et al. (1998). Enhancement of MR Images Using Registration for Signal Averaging. *Journal of Computer Assisted Tomography* 22(2):324–333.

Hunt LT, Hayden BY (2017). A Distributed, Hierarchical and Recurrent Framework for Reward-Based Choice. *Nature Reviews Neuroscience* 18(3):172–182.

Insabato A, Pannunzi M, Rolls ET, Deco G (2010). Confidence-Related Decision Making. *Journal of Neurophysiology* 104(1):539–547.

Izuma K, Murayama K (2013). Choice-Induced Preference Change in the Free-Choice Paradigm: A Critical Methodological Review. *Frontiers in Psychology* 4:41.

Jazayeri M, Movshon JA (2007). Integration of Sensory Evidence in Motion Discrimination. *Journal of Vision* 7(12):7–7.

Jazayeri M, Movshon JA (2006). Optimal Representation of Sensory Information by Neural Populations. *Nature Neuroscience* 9(5):690–696.

Jones AD et al. (2002). A Computational Model of Anterior Cingulate Function in Speeded Response Tasks: Effects of Frequency, Sequence, and Conflict. *Cognitive, Affective & Behavioral Neuroscience* 2(4):300–317.

Joshi S, Li Y, Kalwani RM, Gold JI (2016). Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron* 89(1):221–234.

Karlsson MP, Tervo DGR, Karpova AY (2012). Network Resets in Medial Prefrontal Cortex Mark the Onset of Behavioral Uncertainty. *Science* 338(6103):135–139.

Kass RE, Raftery AE (1995). Bayes Factors. *Journal of the American Statistical Association* 90(430):773–795.

Katahira K (2016). How Hierarchical Models Improve Point Estimates of Model Parameters at the Individual Level. *Journal of Mathematical Psychology* 73:37–58.

Katz LN, Yates JL, Pillow JW, Huk AC (2016). Dissociated Functional Significance of Decision-Related Activity in the Primate Dorsal Stream. *Nature* 535(7611):285–288.

Kelly SP, O'Connell RG (2013). Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. *Journal of Neuroscience* 33(50):19434–19441.

Kepecs A, Mainen ZF (2012). A Computational Framework for the Study of Confidence in Humans and Animals. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367(1594):1322–1337.

Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008). How Uncertainty Boosts Learning: Dynamic Updating of Decision Strategies. *Society for Neurocience*. Washington DC.

Kiani R, Hanks TD, Shadlen MN (2008). Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment. *Journal of Neuroscience* 28(12):3017–3029.

Kiani R, Shadlen MN (2009). Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* 324(5928):759–764.

Kim TD, Kabir M, Gold JI (2017). Coupled Decision Processes Update and Maintain Saccadic Priors in a Dynamic Environment. *Journal of Neuroscience* 37(13):3632–3645.

Kira S, Yang T, Shadlen MN (2015). A Neural Implementation of Wald's Sequential Probability Ratio Test. *Neuron* 85(4):861–873.

Kiyonaga A, Scimeca JM, Bliss DP, Whitney D (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences* 21(7):493–497.

Klein SA (2001). Measuring, Estimating, and Understanding the Psychometric Function: A Commentary. *Perception & Psychophysics* 63(8):1421–1455.

Kleiner M et al. (2007). What's New in Psychtoolbox-3. *Perception* 36(14):1–1.

Kloosterman NA et al. (2015). Pupil Size Tracks Perceptual Content and Surprise. *European Journal of Neuroscience* 41(8):1068–1078.

Knapen T, de Gee JW (2016). FIRDeconvolution - Python Class That Performs Finite Impulse Response Fitting on Time Series Data. *Zenodo*:dx.doi.org/10.5281/zenodo.46216.

Kohn A (2007). Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology* 97(5):3155–3164.

Kok P, Brouwer GJ, van Gerven MAJ, de Lange FP (2013). Prior Expectations Bias Sensory Representations in Visual Cortex. *Journal of Neuroscience* 33(41):16275–16284.

Kok P, Mostert P, de Lange FP (2017). Prior Expectations Induce Prestimulus Sensory Templates. *Proceedings of the National Academy of Sciences* 114(39):10473–10478.

Komura Y et al. (2013). Responses of Pulvinar Neurons Reflect a Subject's Confidence in Visual Categorization. *Nature Neuroscience* 16(6):749–755.

Krishnamurthy K, Nassar MR, Sarode S, Gold JI (2017). Arousal-Related Adjustments of Perceptual Biases Optimize Perception in Dynamic Environments. *Nature Human Behaviour* 1:0107.

Kruschke JK (2013). Bayesian Estimation Supersedes the t Test. *Journal of Experimental Psychology: General* 142(2):573.

Lak A et al. (2014). Orbitofrontal Cortex Is Required for Optimal Waiting Based on Decision Confidence. *Neuron* 84(1):190–201.

Laming D (1968). *Information Theory of Choice-Reaction Times*. London: Academic Press Inc.

Latimer KW et al. (2015). Single-Trial Spike Trains in Parietal Cortex Reveal Discrete Steps during Decision-Making. *Science* 349(6244):184–187.

Lee SH, Dan Y (2012). Neuromodulation of Brain States. *Neuron* 76(1):209–222.

Leite FP, Ratcliff R (2011). What Cognitive Processes Drive Response Biases? A Diffusion Model Analysis. *Judgment & Decision Making* 6(7):651–687.

Lempert KM, Chen YL, Fleming SM (2015). Relating Pupil Dilation and Metacognitive Confidence during Auditory Decision-Making. *PLoS ONE* 10(5):e0126588.

Leopold DA, Wilke M, Maier A, Logothetis NK (2002). Stable Perception of Visually Ambiguous Patterns. *Nature Neuroscience* 5(6):605.

Licata AM et al. (2017). Posterior Parietal Cortex Guides Visual Decisions in Rats. *Journal of Neuroscience* 37(19):4954–4966.

Link SW, Heath RA (1975). A Sequential Theory of Psychological Discrimination. *Psychometrika* 40(1):77–105.

Liu F, Wang XJ (2008). A Common Cortical Circuit Mechanism for Perceptual Categorical Discrimination and Veridical Judgment. *PLoS Computational Biology* 4(12):e1000253.

Logothetis NK, Schall JD (1989). Neuronal Correlates of Subjective Visual Perception. *Science* 245(4919):761–763.

Luce RD (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press.

Luu L, Stocker AA (2016). Choice-Induced Biases in Perception. *bioRxiv* 043224.

Ma WJ, Jazayeri M (2014). Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience* 37(1):205–220.

Macmillan NA, Creelman CD (2004). *Detection Theory: A User's Guide*. Mahwah, NJ: Lawrence Erlbaum Associates.

Marder E (2012). Neuromodulation of Neuronal Circuits: Back to the Future. *Neuron* 76(1):1–11.

Maris E, Oostenveld R (2007). Nonparametric Statistical Testing of EEG- and MEG-Data. *Journal of Neuroscience Methods* 164(1):177–190.

McDougal DH, Gamlin PDR (2008). Pupillary Control Pathways. *The Senses: A Comprehensive Reference*. Ed. by Albright TD et al. New York: Academic Press, pp. 521–536.

McGinley MJ, David SV, McCormick DA (2015a). Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. *Neuron* 87(1):179–192.

McGinley MJ et al. (2015b). Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* 87(6):1143–1161.

Meyer TJ, Miller ML, Metzger RL, Borkovec TD (1990). Development and Validation of the Penn State Worry

Questionnaire. *Behaviour Research and Therapy* 28(6):487–495.

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

Meyniel F, Sigman M, Mainen ZF (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron* 88(1):78–92.

Miller EK, Cohen JD (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience* 24(1):167–202.

Mitra PP, Pesaran B (1999). Analysis of Dynamic Brain Imaging Data. *Biophysical Journal* 76(2):691–708.

Morcos AS, Harvey CD (2016). History-Dependent Variability in Population Dynamics during Evidence Accumulation in Cortex. *Nature Neuroscience* 19(12):1672–1681.

Mulder MJ et al. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience* 32(7):2335–2343.

Muris P, Merckelbach H, Wessel I, van de Ven M (1999). Psychopathological Correlates of Self-Reported Behavioural Inhibition in Normal Children. *Behaviour Research and Therapy* 37(6):575–584.

Murphy PR, Boonstra E, Nieuwenhuis S (2016a). Global Gain Modulation Generates Time-Dependent Urgency during Perceptual Choice in Humans. *Nature Communications* 7:13526.

Murphy PR, van Moort ML, Nieuwenhuis S (2016b). The Pupillary Orienting Response Predicts Adaptive Behavioral Adjustment after Errors. *PLoS ONE* 11(3):e0151763.

Murphy PR, Vandekerckhove J, Nieuwenhuis S (2014). Pupil-Linked Arousal Determines Variability in Perceptual Decision Making. *PLoS Computational Biology* 10(9):e1003854.

Nassar MR et al. (2012). Rational Regulation of Learning Dynamics by Pupil-Linked Arousal Systems. *Nature Neuroscience* 15(7):1040–1046.

Newsome WT, Britten KH, Movshon JA (1989). Neuronal Correlates of a Perceptual Decision. *Nature* 341(6237):52–54.

Nichols MJ, Newsome WT (2002). Middle Temporal Visual Area Microstimulation Influences Veridical Judgments of Motion Direction. *Journal of Neuroscience* 22(21):9530–9540.

Nickerson RS (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2(2):175.

Nienborg H, Cumming BG (2006). Macaque V2 Neurons, But Not V1 Neurons, Show Choice-Related Activity. *Journal of Neuroscience* 26(37):9567–9578.

Nienborg H, Cumming BG (2009). Decision-Related Activity in Sensory Neurons Reflects More than a Neuron's Causal Effect. *Nature* 459(7243):89–92.

Nienborg H, Roelfsema PR (2015). Belief States as a Framework to Explain Extra-Retinal Influences in Visual Cortex. *Current Opinion in Neurobiology* 32:45–52.

Nolte G (2003). The Magnetic Lead Field Theorem in the Quasi-Static Approximation and Its Use for Magnetoencephalography Forward Calculation in Realistic Volume Conductors. *Physics in Medicine and Biology* 48(22):3637–3652.

O'Connell RG, Dockree PM, Kelly SP (2012). A Supramodal Accumulation-to-Bound Signal That Determines Perceptual Decisions in Humans. *Nature Neuroscience* 15(12):1729–1735.

Odoemene O, Nguyen H, Churchland AK (2017). Visual Evidence Accumulation Behavior in Unrestrained Mice. *bioRxiv* 195792.

Oostenveld R, Fries P, Maris E, Schoffelen JM (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience* (156869).

O'Reilly JX (2013). Making Predictions in a Changing World-Inference, Uncertainty, and Learning. *Frontiers in Neuroscience* 7:105.

Ossmy O et al. (2013). The Timescale of Perceptual Evidence Integration Can Be Adapted to the Environment. *Current Biology* 23(11):981–986.

Padoa-Schioppa C (2013). Neuronal Origins of Choice Variability in Economic Decisions. *Neuron* 80(5):1322–1336.

Pape AA, Siegel M (2016). Motor Cortex Activity Predicts Response Alternation during Sensorimotor Decisions. *Nature Communications* 7:13098.

Park-Wyllie LY et al. (2009). Cholinesterase Inhibitors and Hospitalization for Bradycardia: A Population-Based Study. *PLOS Medicine* 6(9):e1000157.

Pfeffer T et al. (2017). Catecholamines, Not Acetylcholine, Alter Cortical and Perceptual Dynamics in Line with Increased Excitation-Inhibition Ratio. *bioRxiv* 170613.

Polack PO, Friedman J, Golshani P (2013). Cellular Mechanisms of Brain State-Dependent Gain Modulation in Visual Cortex. *Nature Neuroscience* 16(9):1331–1339.

Pouget A, Drugowitsch J, Kepecs A (2016). Confidence and Certainty: Distinct Probabilistic Quantities for Different Goals. *Nature Neuroscience* 19(3):366–374.

Preuschoff K, 't Hart BM, Einhäuser W (2011). Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making. *Frontiers in Neuroscience* 5:115.

Rabbitt PMA (1968). Repetition Effects and Signal Classification Strategies in Serial Choice-Response Tasks. *Quarterly Journal of Experimental Psychology* 20(3):232–240.

Rahnev D (2017). The Case against Full Probability Distributions in Perceptual Decision Making. *bioRxiv* 108944.

Rahnev D, Denison R (2016). Suboptimality in Perception. *bioRxiv* 060194.

Ratcliff R (1978). A Theory of Memory Retrieval. *Psychological Review* 85(2):59.

Ratcliff R, Childers R (2015). Individual Differences and Fitting Methods for the Two-Choice Diffusion Model of Decision Making. *Decision* 2(4):237–279.

Ratcliff R, McKoon G (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation* 20(4):873–922.

Ratcliff R, Tuerlinckx F (2002). Estimating Parameters of the Diffusion Model: Approaches to Dealing with Contaminant Reaction Times and Parameter Variability. *Psychonomic Bulletin & Review* 9(3):438–481.

Read JCA (2015). The Place of Human Psychophysics in Modern Neuroscience. *Neuroscience* 296:116–129.

Reimer J et al. (2014). Pupil Fluctuations Track Fast Switching of Cortical States during Quiet Wakefulness. *Neuron* 84(2):355–362.

Ress D, Backus BT, Heeger DJ (2000). Activity in Primary Visual Cortex Predicts Performance in a Visual Detection Task. *Nature Neuroscience* 3(9):940–945.

Roelfsema PR, de Lange FP (2016). Early Visual Cortex as a Multiscale Cognitive Blackboard. *Annual Review of Vision Science* 2(1):131–151.

Rogers SL, Friedhoff LT (1998). Pharmacokinetic and Pharmacodynamic Profile of Donepezil HCl Following Single Oral Doses. *British Journal of Clinical Pharmacology* 46.

Roitman JD, Shadlen MN (2002). Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *Journal of Neuroscience* 22(21):9475–9489.

Rokem A, Silver MA (2010). Cholinergic Enhancement Augments Magnitude and Specificity of Visual Perceptual Learning in Healthy Humans. *Current Biology* 20(19):1723–1728.

Rouder JN, Morey RD, Speckman PL, Province JM (2012). Default Bayes Factors for ANOVA Designs. *Journal of Mathematical Psychology* 56(5):356–374.

Rowan TH (1990). Functional Stability Analysis Of Numerical Algorithms. PhD thesis. Austin: University of Texas.

Sanders JI, Hangya B, Kepecs A (2016). Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron* 90(3):499–506.

Sara SJ (2009). The Locus Coeruleus and Noradrenergic Modulation of Cognition. *Nature Reviews Neuroscience* 10(3):211–223.

Sauer JM, Ring BJ, Witcher JW (2005). Clinical Pharmacokinetics of Atomoxetine. *Clinical Pharmacokinetics* 44(6):571–590.

Scase MO, Braddick OJ, Raymond JE (1996). What Is Noise for the Motion System? *Vision Research* 36(16):2579–2586.

Schwarz G (1978). Estimating the Dimension of a Model. *The Annals of Statistics* 6(2):461–464.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996). A Computational Analysis of the Relationship between Neuronal and Behavioral Responses to Visual Motion. *Journal of Neuroscience* 16(4):1486–1510.

Shadlen MN, Kiani R (2013). Decision Making as a Window on Cognition. *Neuron* 80(3):791–806.

Shadlen MN, Shohamy D (2016). Decision Making and Sequential Sampling from Memory. *Neuron* 90(5):927–939.

Shaw AD et al. (2017). Neurophysiologically-Informed Markers of Individual Variability and Pharmacological

Manipulation of Human Cortical Gamma. *NeuroImage* 161:19–31.

Siegel M, Buschman TJ, Miller EK (2015). Cortical Information Flow during Flexible Sensorimotor Decisions. *Science* 348(6241):1352–1355.

Siegel M, Engel AK, Donner TH (2011). Cortical Network Dynamics of Perceptual Decision-Making in the Human Brain. *Frontiers in Human Neuroscience* 5.

Siegel M et al. (2008). Neuronal Synchronization along the Dorsal Visual Pathway Reflects the Focus of Spatial Attention. *Neuron* 60(4):709–719.

Siegel M et al. (2006). High-Frequency Activity in Human Visual Cortex Is Modulated by Visual Motion Strength. *Cerebral Cortex* 17(3):732–741.

Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A (2002). Bayesian Measures of Model Complexity and Fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(4):583–639.

Spielberg JM et al. (2011). Approach and Avoidance Profiles Distinguish Dimensions of Anxiety and Depression. *Cognitive Therapy and Research* 35(4):359–371.

St. John-Saaltink E, Kok P, Lau HC, de Lange FP (2016). Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex. *Journal of Neuroscience* 36(23):6186–6192.

Stanislaw H, Todorov N (1999). Calculation of Signal Detection Theory Measures. *Behavior Research Methods, Instruments, & Computers* 31(1):137–149.

Steiger JH (1980). Tests for Comparing Elements of a Correlation Matrix. *Psychological Bulletin* 87(2):245.

Steriade M (2000). Corticothalamic Resonance, States of Vigilance and Mentation. *Neuroscience* 101(2):243–276.

Stocker A, Simoncelli EP (2008). A Bayesian Model of Conditioned Perception. *Advances in Neural Information Processing Systems* 20:1409–1416.

Stolk A, Todorovic A, Schoffelen JM, Oostenveld R (2013). Online and Offline Tools for Head Movement Compensation in MEG. *NeuroImage* 68:39–48.

Sugrue LP, Corrado GS, Newsome WT (2004). Matching Behavior and the Representation of Value in the Parietal Cortex. *Science* 304(5678):1782–1787.

Summerfield C, de Lange FP (2014). Expectation in Perceptual Decision Making: Neural and Computational Mechanisms. *Nature Reviews Neuroscience* 15(11):745–756.

Summerfield C, Tsetsos K (2012). Building Bridges between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms. *Frontiers in Neuroscience* 6.

Sutton RS, Barto AG (1998). *Reinforcement Learning: An Introduction.* Cambridge, Massachsetts: A Bradford Book.

Teichert T, Ferrera VP, Grinband J (2014). Humans Optimize Decision-Making by Delaying Decision Onset. *PLoS ONE* 9(3):e89638.

Tervo DGR et al. (2014). Behavioral Variability through Stochastic Choice and Its Gating by Anterior Cingulate Cortex. *Cell* 159(1):21–32.

Thaler L, Schütz A, Goodale M, Gegenfurtner K (2013). What Is the Best Fixation Target? The Effect of Target Shape on Stability of Fixational Eye Movements. *Vision Research* 76:31–42.

Tsetsos K, Pfeffer T, Jentgens P, Donner TH (2015). Action Planning and the Timescale of Evidence Accumulation. *PLoS ONE* 10(6):e0129473.

Twomey DM, Kelly SP, O'Connell RG (2016). Abstract and Effector-Selective Decision Signals Exhibit Qualitatively Distinct Dynamics before Delayed Perceptual Reports. *Journal of Neuroscience* 36(28):7346–7352.

Twomey DM, Murphy PR, Kelly SP, O'Connell RG (2015). The Classic P300 Encodes a Build-to-Threshold Decision Variable. *European Journal of Neuroscience* 42(1):1636–1643.

Urai AE, Braun A, Donner TH (2017). Pupil-Linked Arousal Is Driven by Decision Uncertainty and Alters Serial Choice Bias. *Nature Communications* 8:14637.

Urai AE, Murphy PR (2016). Commentary: Sensory Integration Dynamics in a Hierarchical Network Explains Choice Probabilities in Cortical Area MT. *Frontiers in Systems Neuroscience* 10(37).

Urai AE, Pfeffer T (2014). An Action-Independent Signature of Perceptual Choice in the Human Brain. *Journal of Neuroscience* 34(15):5081–5082.

van Bergen RS (2017). Sensory Uncertainty and Response Variability in Human Visual Cortex. PhD thesis. Nijmegen: Radboud University.

van Ravenzwaaij D, Mulder MJ, Tuerlinckx F, Wagenmakers EJ (2012). Do the Dynamics of Prior

Information Depend on Task Context? An Analysis of Optimal Performance and an Empirical Test. *Frontiers in Psychology* 3:132.

Van Veen B, van Drongelen W, Yuchtman M, Suzuki A (1997). Localization of Brain Electrical Activity via Linearly Constrained Minimum Variance Spatial Filtering. *IEEE Transactions on Biomedical Engineering* 44(9):867–880.

Varazzani C, San-Galli A, Gilardeau S, Bouret S (2015). Noradrenaline and Dopamine Neurons in the Reward/Effort Trade-off: A Direct Electrophysiological Comparison in Behaving Monkeys. *Journal of Neuroscience* 35(20):7866–7877.

Vinck M, Batista-Brito R, Knoblich U, Cardin JA (2015). Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron* 86(3):740–754.

Wald A (1947). *Sequential Analysis*. Courier Corporation.

Wallis WA (1980). The Statistical Research Group, 1942–1945. *Journal of the American Statistical Association* 75(370):320–330.

Wandell BA, Dumoulin SO, Brewer AA (2007). Visual Field Maps in Human Cortex. *Neuron* 56(2):366–383.

Wang CA, Munoz DP (2015). A Circuit for Pupil Orienting Responses: Implications for Cognitive Modulation of Pupil Size. *Current Opinion in Neurobiology* 33:134–140.

Wang XJ (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron* 36(5):955–968.

Wang XJ (2008). Decision Making in Recurrent Neuronal Circuits. *Neuron* 60(2):215–234.

Weber EH (1846). Tastsinn Und Gemeingefühl. *Handwörterbuch Der Physiologie*. Leipzig: R. Wagner.

Wei XX, Stocker AA (2015). A Bayesian Observer Model Constrained by Efficient Coding Can Explain 'anti-Bayesian' Percepts. *Nature Neuroscience* 18(10):1509–1517.

Wernicke JF et al. (2003). Cardiovascular Effects of Atomoxetine in Children, Adolescents, and Adults. *Drug Safety* 26(10):729–740.

Wessel JR, Danielmeier C, Ullsperger M (2011). Error Awareness Revisited: Accumulation of Multimodal Evidence from Central and Autonomic Nervous Systems. *Journal of Cognitive Neuroscience* 23(10):3021–3036.

Wetzels R, Wagenmakers EJ (2012). A Default Bayesian Hypothesis Test for Correlations and Partial Correlations. *Psychonomic Bulletin & Review* 19(6):1057–1064.

White CN, Poldrack RA (2014). Decomposing Bias in Different Types of Simple Decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 40(2):385–398.

Wichmann FA, Hill NJ (2001a). The Psychometric Function: I. Fitting, Sampling, and Goodness of Fit. *Perception & Psychophysics* 63(8):1293–1313.

Wichmann FA, Hill NJ (2001b). The Psychometric Function: II. Bootstrap-Based Confidence Intervals and Sampling. *Perception & Psychophysics* 63(8):1314–1329.

Wiecki TV, Sofer I, Frank MJ (2013). HDDM: Hierarchical Bayesian Estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics* 7.

Wimmer K et al. (2015). Sensory Integration Dynamics in a Hierarchical Network Explains Choice Probabilities in Cortical Area MT. *Nature Communications* 6:6177.

Yeung N, Botvinick MM, Cohen JD (2004). The Neural Basis of Error Detection: Conflict Monitoring and the Error-Related Negativity. *Psychological Review* 111(4):931.

Yu AJ, Cohen JD (2008). Sequential Effects: Superstition or Rational Behavior? *Advances in Neural Information Processing Systems* 21:1873–1880.

Yu AJ (2005). Ach and NE: Bayes, Uncertainty, Attention and Learning. Ph.D. University College London (University of London).

Zhang M, Wang X, Goldberg ME (2014). A Spatially Nonselective Baseline Signal in Parietal Cortex Reflects the Probability of a Monkey's Success on the Current Trial. *Proceedings of the National Academy of Sciences* 111(24):8967–8972.

6 | Acknowledgements

If our ignorance is infinite, the only possible course of action is to muddle through as best we can.
— *Martin Schwarz*, The importance of stupidity in scientific research

I wouldn't have made it to the start, let alone finish this thesis without the support of a great number people.

Tobias, thank you for your supervision and mentorship. Under your guidance I have grown from a nervous MSc student, asking you if I could please do my PhD in your group, into the researcher I am now. Your curiosity, experimental rigour and intellectual breadth continue to inspire me, and I am grateful for all your encouragement, patience and seemingly endless enthusiasm for science.

I am grateful to Professor Andreas Engel for supporting me as a PhD student in the institute, and Professor Laura Busse and Professor Christoph Mulert for reviewing my thesis. Throughout the years I've had the opportunity to work with a number of brilliant students: thank you, I hope you learned as much from me as I did from you. I also want to thank all my participants for doing endless, boring random dot tasks and lending their brains to scientific discovery.

Many thanks to all members of the growing Thunderlab family. I immensely enjoyed our many discussions: from colormaps to Bayes factors, from the pros and cons of scientific toolboxes to the best way to parallelize thousands of model fits, and from Matlab vs. Python to the difficulties of interpreting four-dimensional cluster statistics. Thank you for all the pizza's at Martinos, movies in the guest room, rounds of pingpong and perfectly brewed espresso, mental support when it was most needed, complaining about reviewer 3, and agonizing over the problems that come with a career in science. Additional thanks goes out to those of you who proofread parts of this thesis, provided feedback and caught embarrassing errors.

I've been lucky to have the best office mates a scientist could wish for. Arriving in the morning and looking forward to our punctual coffee runs, sharing scientific and personal highs and lows or just semi-inspirational quotes of the month, has kept me sane. I will always remember our regular breaks to sing Christmas songs (appropriate any time of year), come up with baby names, or play newly invented throw-a-pingpong-ball-around-the-room games.

To all my colleagues in Amsterdam and Hamburg, thank you for the fun, wisdom, frustration and support we've shared, for our discussions on the philosophy of neuroscience, feminism, the trade-off between

biological plausibility and parameter identifiability, and the best way to grow cacti or teach the Fourier transform; and for all the journal clubs, lunches, beer tastings, toad fests, nerd nights and conference discos.

More broadly, I'm indebted to all my previous and current classmates, lab-buddies, summer school friends, writing partners, impromptu mentors and unknowing role models, Twitterati and conference companions. Sharing the challenges and joys of science makes it infinitely more fun and interesting. You know who you are; may our coffee be strong, our discussions be fruitful and our paths continue to cross.

I want to thank my friends for many conversations about life, the universe and everything, muddy festivals, spontaneous beers in the sun, long-planned weekends away, and the occasionally necessary reminder that the latest bugs in my code really aren't the most important thing in life.

Thanks to my parents, brother and sister, the people who know me longest and probably best, for always lovingly supporting my latest crazy trip or plan, and for always being a home to return to - whatever life throws at us. I'm very happy and thankful for our wonderful extended patchwork family.

Jeroen, thank you for your love, support and humour. I feel incredibly lucky to share my life with you.

I dedicate this thesis to Hilda Urai, who would have been immensely proud to see me become the third generation Dr. Urai.

And now, on to the next adventure!

7 | Curriculum Vitae

Lebenslauf aus datenschutzrechtlichen Gründen nicht enthalten.

8 | 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.

Anne Urai