CHOICE HISTORY BIASES SUBSEQUENT EVIDENCE ACCUMULATION

Anne E. Urai[1,2,*], Jan Willem de Gee[1,2], Tobias H. Donner[1,2,3,*]

1 Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

2 Department of Psychology, University of Amsterdam, Amsterdam, The Netherlands

3 Amsterdam Brain & Cognition, University of Amsterdam, Amsterdam, The Netherlands

* Correspondence: anne.urai@gmail.com, t.donner@uke.de

ABSTRACT

Perceptual choices depend not only on the current sensory input, but also on the behavioral context. An important contextual factor is the history of one's own choices. Choice history often strongly biases perceptual decisions, and leaves traces in the activity of brain regions involved in decision processing. Yet, it remains unknown how such history signals shape the dynamics of later decision formation. Models of perceptual choice construe decision formation as the accumulation of sensory evidence towards decision bounds. In this framework, it is commonly assumed that choice history signals shift the starting point of accumulation towards the bound reflecting the previous choice. We here present results that challenge this idea. We fit a bounded accumulation ('drift diffusion') decision model to behavioral data from multiple perceptual choice tasks and sensory modalities, and estimated bias terms that dependent on observers' previous choices. Individual history biases in behavior were consistently explained by a history-dependent change in the evidence accumulation, rather than in its starting point. Choice history signals thus seem to affect the interpretation of current sensory input, akin to shifting endogenous attention towards (or away from) the previously selected interpretation.

INTRODUCTION

Decisions are not isolated events, but are embedded in a sequence of choices. Preceding choices can exert a large influence even on low-level perceptual judgments (Fernberger, 1920). Previous work on the mechanisms of perceptual decision-making has largely focused on processing of the current sensory evidence (Gold and Shadlen, 2007). Yet, humans (Fründ et al., 2014; Urai et al., 2017), monkeys (Gold et al., 2008) and rodents (Busse et al., 2011) base their choices not only on sensory input, but also on choice history, in a way that can be adjusted to serial correlations in the sensory environment (Abrahamyan et al., 2016; Kim et al., 2017; Braun et al., 2018). History biases are also prevalent in environments lacking such structure, and vary substantially across individuals (Fründ et al., 2014; Abrahamyan et al., 2016; Urai et al., 2017). Choice history biases are ubiquitous also in other domains of decision-making (Leopold et al., 2002; Allefeld et al., 2013; Padoa-Schioppa, 2013). They likely reflect a general principle of the mechanisms governing decision-making.

Computational theory (Gao et al., 2009; Glaze et al., 2015) and psychophysical data (Kim et al., 2017; Braun et al., 2018) indicate that choice history biases result from the accumulation of internal decision variables across trials, with a timescale governed by the decision-makers' internal model of the correlation structure of their environment. Neural signals reflecting previous choices have been found across the sensorimotor pathways of the cerebral cortex, from sensory to associative and motor regions (de Lange et al., 2013; Akaishi et al., 2014; Pape and Siegel, 2016; Purcell and Kiani, 2016; St. John-Saaltink et al.,

2016; Thura et al., 2016; Hwang et al., 2017; Scott et al., 2017). But how such choice history signals act on the formation of new decisions, has so far remained elusive.

Here, we provide insight into this mechanism. Current models of perceptual decision-making posit the accumulation of noise-corrupted sensory evidence over time, resulting in an internal decision variable that grows with time (Bogacz et al., 2006; Gold and Shadlen, 2007; Ratcliff and McKoon, 2008; Brody and Hanks, 2016). When this decision variable reaches one of two decision bounds, a choice is made and the corresponding motor response is initiated. In this framework, a bias can be brought about in two ways: (i) by shifting the starting point of accumulation towards one of the two bounds, or (ii) by selectively changing the rate at which evidence for one versus the other choice alternative is accumulated (Figure 1). The former can be conceptualized as adding an offset to that 'perceptual interpretation signal' during the generation of the response, whereas the latter as altering the perceptual interpretation of the current sensory input.

Existing computational models of choice history biases postulate a shift in the starting point of the decision variables towards the bound of the previous choice (Yu and Cohen, 2008; Gao et al., 2009). This prediction is natural, assuming that history biases are due to the slow (passive) decay of the decision variable into the next trial (Cho et al., 2002; Gao et al., 2009; Bonaiuto et al., 2016). However, the cerebral cortex is equipped with a hierarchy of timescales (Murray et al., 2014; Wei and Wang, 2016), and choice history biases might originate from the accumulation of decision variables across trials, which might interact with several different stages of the decision-making process. Recent experimental work quantifying history biases in perceptual choice did not analyze the within-trial dynamics of decision formation (Busse et al., 2011; de Lange et al., 2013; Akaishi et al., 2014; Fründ et al., 2014; Urai et al., 2017; Braun et al., 2018), only allowed starting point to vary with choice history (Cho et al., 2002; Gold et al., 2008; Yu and Cohen, 2008; Wilder et al., 2009; Bode et al., 2012; Zhang et al., 2014) or analyzed only mean reaction times, rather than their full distributions (Gao et al., 2009; Jones et al., 2013). Consequently, it is unknown whether choice history acts by shifting the starting point or altering the evidence accumulation process.

We addressed this issue by fitting a computational decision model to human behavioral data from five studies of perceptual decision-making, covering a variety of task protocols and sensory modalities. Surprisingly, we found that across all five data sets, history biases evident in observers' overt choice behavior were explained by a history-dependent change in the accumulation bias rather than the starting point. This indicates that the interaction between choice history signals and decision formation is more complex than previously thought. Choices may act like an endogenous cue for selective attention that biases evidence accumulation towards (or away from) the previous chosen perceptual interpretation of the sensory input.

RESULTS

We fit the drift diffusion model to behavioral data (choices and reaction times, RT) from a total of 162 human participants (Materials and Methods). The drift diffusion model (DDM) is a popular variant of sequential sampling models of decision-making (Bogacz et al., 2006; Ratcliff and McKoon, 2008). It provides good fits to RT and choice patterns from a vast array

of two-choice task. It quantifies latent computational variables (model parameters) that naturally map onto elemental mechanisms situated at different stages of the sensory-motor pathways of the brain (Wong and Wang, 2006; Gold and Shadlen, 2007). The shape of RT distributions and choice fractions jointly constrain the model parameters. We here estimated the following parameters: non-decision time (the time needed for sensory encoding and response execution), starting point of the decision variable, separation of the decision bounds, mean drift rate, and a stimulus-independent constant added to the mean drift. We refer to the latter parameter - termed drift criterion by Ratcliff and McKoon (2008) - as 'drift bias'.

Within the DDM, choice history could bias the decision process by two mechanisms, illustrated in Figure 1. Either the starting point or the drift bias could shift towards the bound corresponding to the previous choice. A shift in starting point would be most prevalent early on in the decision process: it would affect the leading edge of the RT distribution, shifting its mode. It predicts that the majority of history-dependent choice biases occur on trials with fast reaction times. A shift in the drift bias instead grows with time. Therefore, it would affect the trailing edge of the distribution with little effect on the mode. In contrast to starting point bias, drift bias alters choice fractions across the whole range of reaction times, well into the tail of the RT distribution. Both mechanisms generate predictions that are indistinguishable in terms of the proportion of choices alone; the shape of RT distributions is the diagnostic feature that distinguishes between them.

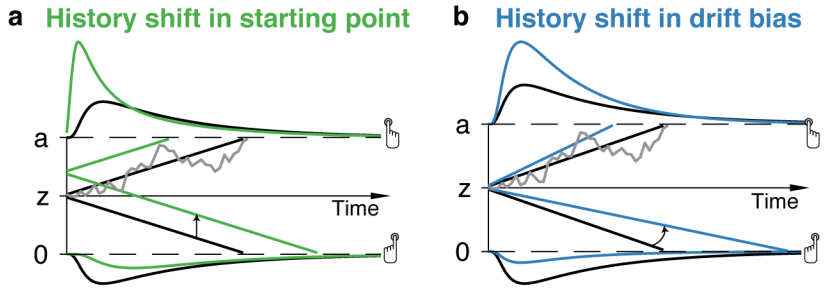

Figure 1. Two biasing mechanisms within the drift diffusion model. *(**a**) Choice history-dependent shift in starting point. The model postulates that noisy sensory evidence is accumulated over time, until the resulting decision variable y reaches one of two bounds (dashed black lines at y=0 and y=a) for the two choice options. Repeating this process over many trials yields RT distributions for both choices (plotted above and below the bounds). Gray line: example trajectory of decision variable from single trial. Black lines: mean drift and resulting RT distributions under unbiased conditions. Green lines: mean drift and RT distributions under biased starting point. (**b**) As (a), but for choice history-dependent shift in drift bias. Blue lines: mean drift and RT distributions under biased drift. Both mechanisms differentially affect the shape of RT distributions. See main text for details. RT distributions were simulated using equations from Ratcliff and Tuerlinckx (2002) and code from Murphy et al. (2016).*

We fit different variants of the DDM to data from five different experiments. These covered a range of task protocols and sensory modalities commonly used in the study perceptual decision-making (see Figure 2a-d and Materials and Methods for details): two alternative forced-choice, two interval forced-choice, as well as yes-no (simple forced choice) tasks; RT as well as so-called fixed duration tasks; visual motion direction and coherence discrimination; and visual contrast and auditory detection. As found in previous work

(Abrahamyan et al., 2016; Urai et al., 2017), observers exhibited a wide range of idiosyncratic choice history biases across all experiments, quantified here in terms of the probability to repeat the previous choice (Figure 2e).

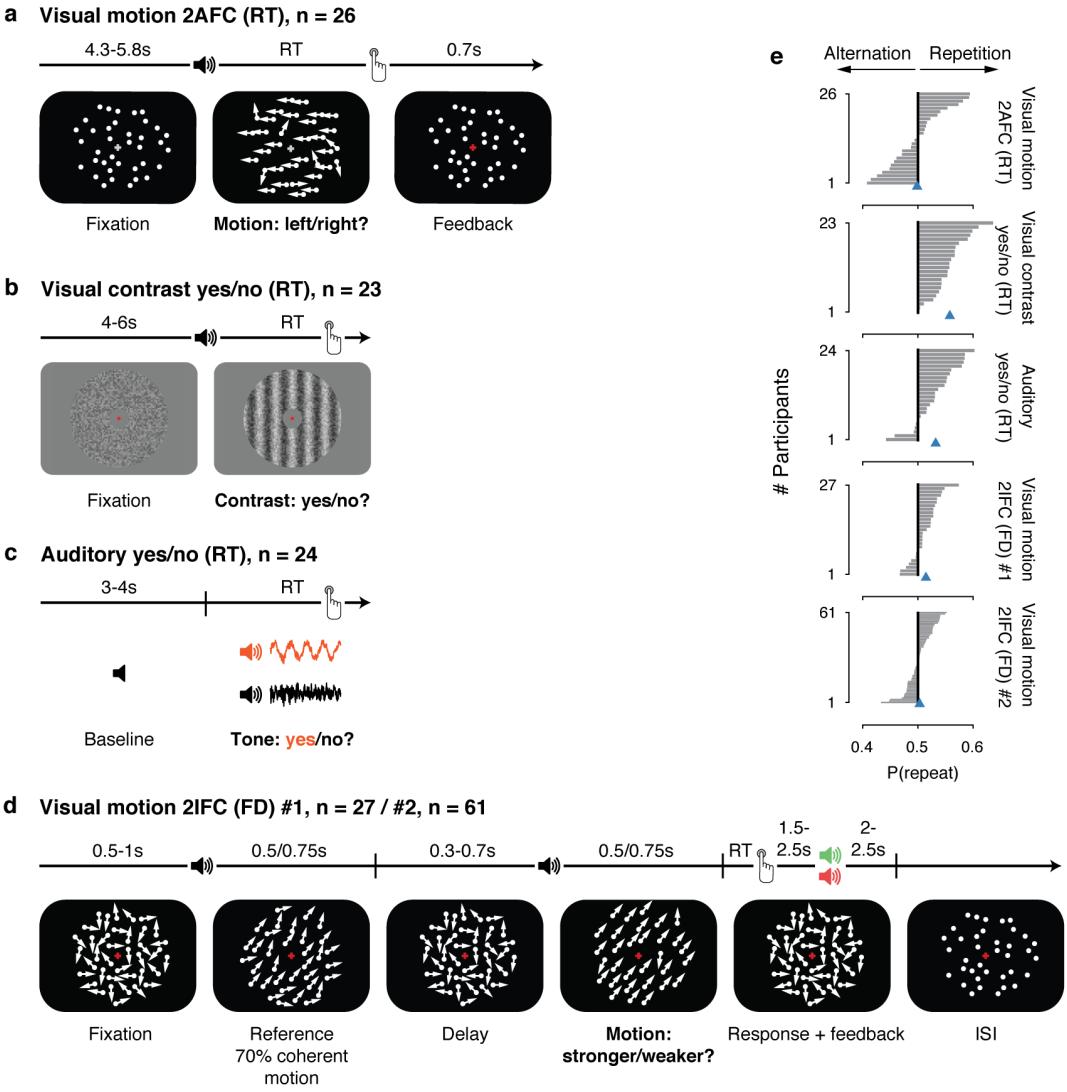

Figure 2. Behavioral tasks. *(**a**) Visual motion 2AFC (RT) task. (**b**) Visual contrast yes/no (RT) task. (**c**) Auditory yes/no (RT) task. (**d**) Visual motion 2IFC (FD) task, which was used in two separate data sets. See Methods for details. (**e**) Distributions of individual choice history biases for each data set. Choice history bias was quantified as the probability of choice repetition, using 0.5 as reference. Blue triangles, mean over participants.*

The DDM fit the data from all experiments well, and the parameter estimates behaved as expected (Figure 3). We first fit a basic version of the DDM that contained the above-described parameters, but without allowing bias parameters to vary with choice history. In all data sets, individual drift rate estimates correlated with individual perceptual sensitivity (d', Figure 3a). In the one data set, in which the strength of sensory evidence was manipulated systematically (Visual motion 2IFC (FD) #1), the mean drift rate increased monotonically as a function of evidence strength (Figure 3a). The RT distributions, shown separately for each combination of stimuli and choices in Figure 3b, closely matched the model predictions

(darker colors indicate predicted RTs obtained through model simulations). Note that this basic DDM model included an overall drift bias, which captured an often observed conservative tendency towards choosing 'no' in the yes/no tasks (de Gee et al., 2014, 2017).

In fixed duration tasks, the decision-maker does not need to set a bound for terminating the decision (Bogacz et al., 2006), so the bounded diffusion process described by the DDM might seem inappropriate. Yet, the success of the DDM in fitting these data was consistent with previous work (e.g. Ratcliff, 2006; Bode et al., 2012; Jahfari et al., 2012) and might have reflected the fact that observers set implicit decision bounds also when they do not control the stimulus duration (Kiani et al., 2008; but see Tsetsos et al., 2015).

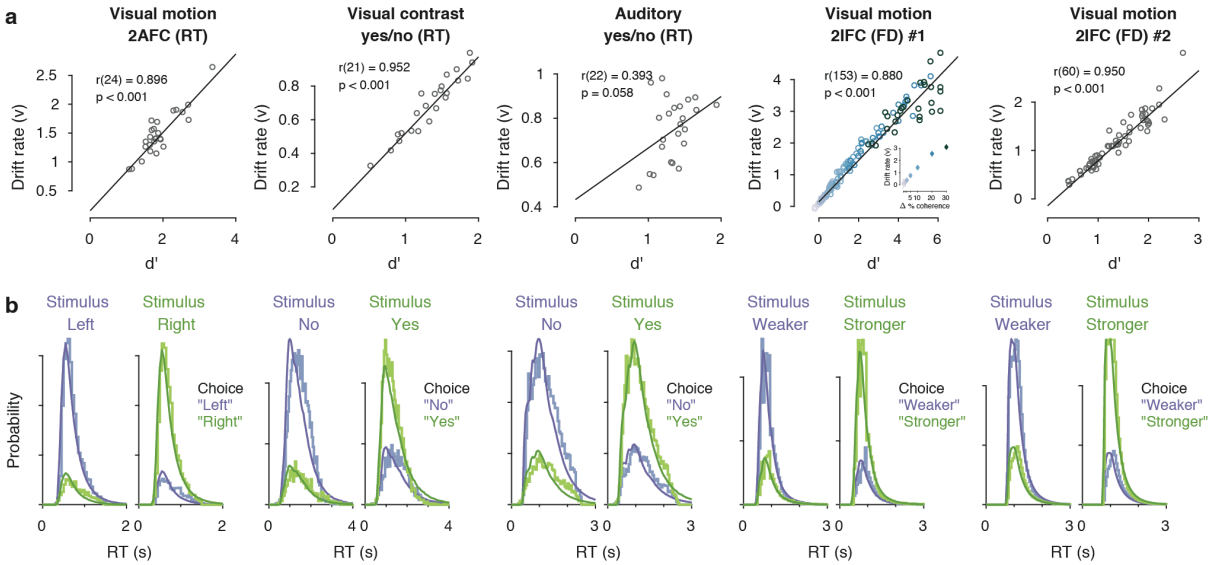

Figure 3. Drift diffusion model fits all data sets well. *(**a**) Correlation between drift rate (v) and d'. In the 2IFC #1 data set, drift rates were estimated separately for each level of stimulus difficulty and the correlation coefficient displayed is a partial correlation between v and d', while accounting for stimulus difficulty (indicated in the inset). (**b**) Measured and predicted RT distributions, across all trials and observers within each dataset. Observed (light) and predicted (dark) RT distributions are separately shown for each combination of stimulus and choice (green/purple), with the low-probability distributions indicating error trials.*

HISTORY-DEPENDENT SHIFTS IN DRIFT BIAS, NOT STARTING POINT, EXPLAIN INDIVIDUAL CHOICE HISTORY BIASES

We then fit the DDM while also allowing starting point, drift bias, or both to vary as a function of the observer's choice on the previous trial (Supplementary Figure 1). Models with history-dependent biases better explained the data than the baseline model without such history dependence (Figure 4a), corroborating the observation that observers' behavior showed considerable dependence on previous choices (Figure 2). The model with both history-dependent starting point and drift bias provided the best fit to all five data sets (Figure 4a), based on the Deviance Information Criterion (DIC, see Materials and Methods).

The above formal model comparison pointed to the importance of including a history-dependency into the model. Yet, the DIC is prone to selecting more complex models (Wiecki et al., 2013). We further examined the ability of each model to explain the diagnostic features in the data (Palminteri et al., 2017) that distinguished starting point from drift bias. As shown

above (Figure 1), a history-dependent shift in the starting point leads to biased choices primarily when responses are fast (early RT quantiles), whereas history-dependent shift in drift leads to biased choices across all trials, including those with slow responses (Figure 1). We simulated choices and RTs from the four different model variants and computed so-called 'conditional bias functions' (see White and Poldrack (2014) and Materials and Methods): the fraction of choices in line with each observer's choice repetition tendency (i.e., repetition probability) within each quantile of their RT distribution. For observers whose choice repetition probability was > 0.5, this was the fraction of repetitions; for the remaining observers, this was the fraction of alternations. Consistent with a shift in drift, observers exhibited history-dependent choice biases across the entire range of RTs across data sets (Figure 4b). In particular, the biased choices on slow RTs could only be captured by models that included a history-dependent shift in drift (Figure 4c, blue bars).

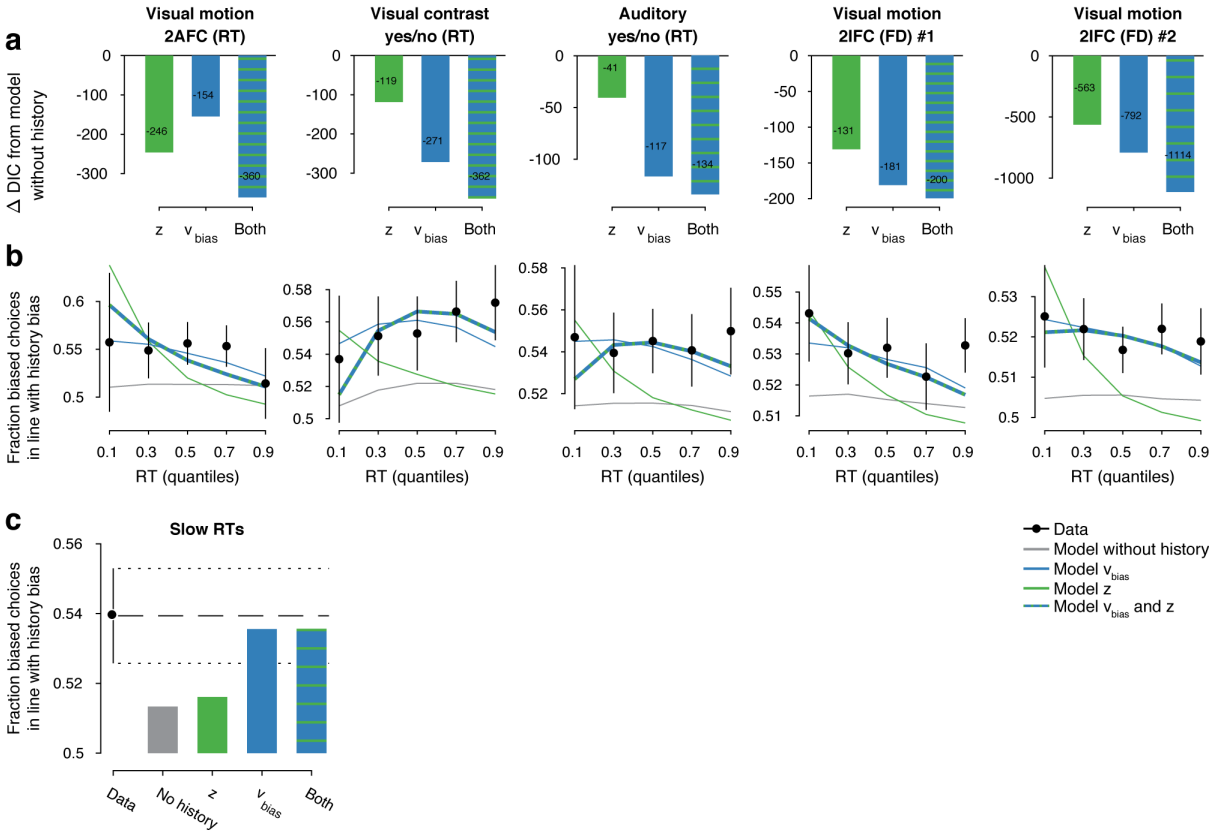

Figure 4. Model comparison and simulations. *(**a**) For each dataset, we compared the DIC (Materials and Methods) between models where drift bias, starting point bias or both were allowed to vary as a function of previous choice. The DIC for a model without history dependence was used as a baseline for each data set. Lower DIC values indicate a model that is better able to explain the data, after taking into account the model complexity; a DIC of 10 is generally taken as a threshold for considering one model a sufficiently better fit. (**b**) Conditional bias functions. For each of four simulated models, as well as the observed data, we divided all trials into quantiles of the RT distribution. For each quantile, the fraction of choices biased towards each individual's history bias (repetition or alternation) indicates the degree to which behavior is biased, within that range of RTs. For the 2IFC data set, conditional bias functions were computed separately for each difficulty level and subsequently averaged to visualize the effect of history bias over and above the effect of stimulus difficulty on both RTs and choices. The best-fitting model is shown in thicker lines. (**c**) Choice bias on slow response trials (last three quantiles*

of the RT distribution) can be captured only by models that include history-dependent drift bias. Black error bars indicate mean ± 95% confidence interval across all data sets, bars indicate the predicted fraction of choices in late RT quantiles.

We used the parameter estimates obtained from the full model (with both history-dependent starting point and drift bias) to investigate how the choice history-dependent variations in starting point and drift bias related to each individual's tendency to repeat their previous choices. We call each bias parameter's dependence on the previous choice its 'history shift'. For instance, in the left vs. right motion discrimination task, the history shift in starting point was computed as the difference between the starting point estimate for previous 'left' and previous 'right' choices. Across all five data sets, the history shift in drift bias, but not the history shift in starting point, was robustly correlated to the individual probability of choice repetition (Figure 5, significant correlations indicated with solid regression lines). In four out of five data sets, the correlation with the history shift in drift bias was significantly stronger than the correlation with the history shift in starting point (Figure 5, Δr values).

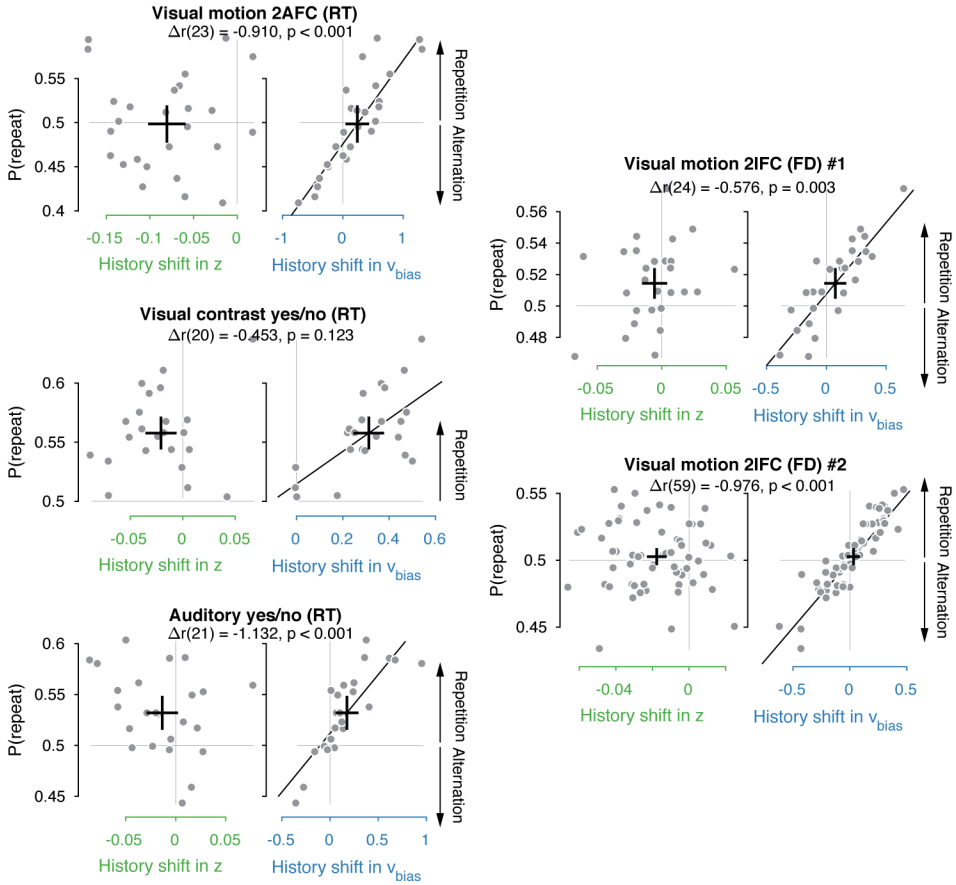

Figure 5. Individual choice history biases are explained by history-dependent changes in drift bias, not starting point*. Correlations between individual choice repetition probabilities, P(repeat), and history shift in starting point (left column, green) and drift (right column, blue). Parameter estimates were obtained from a model in which both bias terms were allowed to vary with previous choice. Horizontal and vertical lines, unbiased references. Thick black crosses, group mean ± s.e.m. in both directions. Black lines best fit of a linear regression (only plotted for significant correlations). Δr quantifies the extent to which the two DDM parameters are differentially able to predict individual choice repetition (p-values from Steiger's test).*

We quantified the total evidence by computing a Bayes factor for each correlation (Wetzels and Wagenmakers, 2012), and multiplying these across data sets (Scheibehenne et al., 2016). This further confirmed that individual choice history biases were not captured by history shifts in starting point, but consistently captured by history shifts in drift (Figure 6). Specifically, the Bayes factor for the history shift in starting point approached zero, indicating strong evidence for the null hypothesis of no correlation. The Bayes factor for the history shift in drift indicated strong evidence for a correlation (Kass and Raftery, 1995).

The same qualitative pattern of results was obtained with an alternative fitting procedure (non-hierarchical G^2 optimization procedure (Ratcliff and Tuerlinckx, 2002), Figure S3a), as well as a model that allowed for additional across-trial variability in starting point (Figure S3b). These findings are thus robust to specifics of the model and fitting method. The fifth data set (Visual motion 2IFC #2) also entailed two distinct pharmacological interventions in two sub-groups of participants (Materials and Methods). The same effects as in Figures 5 and 6 were found for both sub-groups as well as the placebo group (Figure S5).

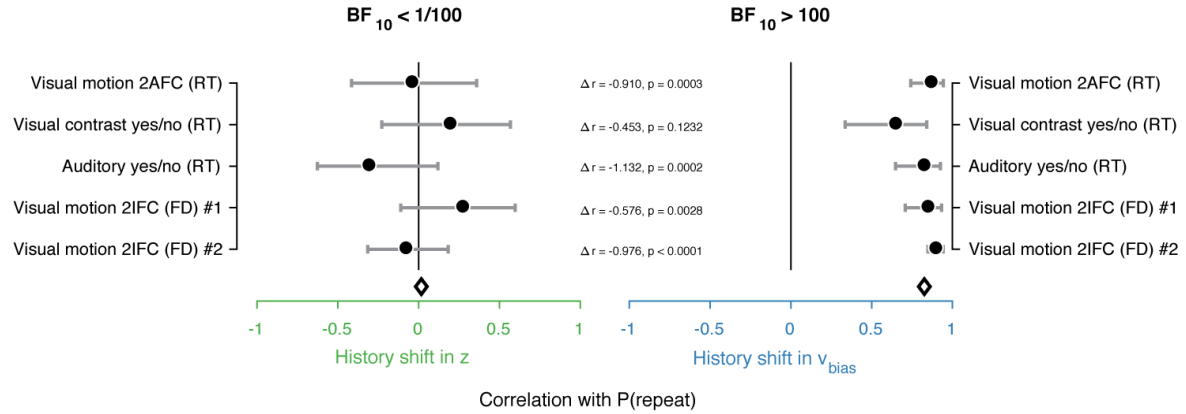

Figure 6. Consistent evidence across data sets. *Summary of the correlations between individual choice repetition probability and the history shifts in starting point (green; left) and drift bias (blue; right). Error bars indicate the 95% confidence interval of the correlation coefficient. Δr quantifies the extent to which the two DDM parameters are differentially able to predict individual choice repetition probability, p-values from Steiger's test. The black diamond indicates the mean correlation coefficient across data sets. The Bayes factor (BF_{10}) quantifies the relative evidence for the alternative over the null hypothesis, with values < 1 indicating evidence for the null hypothesis of no correlation, and > 1 indicating evidence for a correlation.*

DISCUSSION

Observers' perceptual choices often depend on choices made before, even when subsequent stimuli in the environment are uncorrelated. Such choice biases are ubiquitous across domains of decision-making. So far, it has remained unknown how these history biases shape the dynamics of the subsequent decision process. We here fit bounded accumulation (drift diffusion) decision models to behavioral data from a range of psychophysical experiments. This allowed us to tease apart two alternative effects of choice history on the current decision process: a shift in the starting point of the decision variable, or a change in the rate at which evidence for one versus the other option is accumulated. These two scenarios can lead to identical (history-dependent) choice patterns, but they can

be distinguished based on their effects on RT distributions. Surprisingly, we obtained strong and consistent evidence for the second scenario, a history-dependent shift of the drift. This result calls for a revision of current models of history biases (Yu and Cohen, 2008; Zhang et al., 2014).

A number of previous studies have used the drift diffusion model to tease apart shifts in starting point and drift bias as a consequence 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; but see Hanks et al., 2011). Similarly, a higher reward associated with one response option moves the starting point towards that option (Rorie et al., 2010; Gao et al., 2011; Leite and Ratcliff, 2011; Mulder et al., 2012; White and Poldrack, 2014). In both these cases, shifts in starting point correspond to optimal behavior; 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). By contrast, drift bias 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).

This previous work used bias manipulations that were under experimental control: either constant within experimental blocks, or cued at the single-trial level. By contrast, the choice history biases studied here emerge spontaneously, in an idiosyncratic fashion (with different signs and largely varying magnitudes between individuals, Figure 2e). While these biases are systematic properties of each individual’s behavior, they vary dynamically from trial to trial. For example, in a ‘repeating’ observer performing the visual motion 2AFC task, a leftward choice will be followed by tendency to accumulate towards the bound for leftward choices, and a rightward choice will be followed by a tendency to accumulate towards the bound for rightward choices. Traditional analyses of perceptual decision-making mistakenly attribute such dynamic biases to ‘noise’ in the decision process by ignoring choice history (Shadlen et al., 1996; Renart and Machens, 2014; Wyart and Koechlin, 2016), thus underestimating the true perceptual sensitivity of the observer (Fründ et al., 2014; Abrahamyan et al., 2016). Specifically, our present results imply that choice history systematically contributes to the trial-to-trial variability in drift. This is analogous to trial-to-trial variations of drift bias due to phasic boosts of arousal as indexed by pupil dilation (de Gee et al., 2017). Tracking ‘contextual factors’ such as choice history allows for the partitioning of drift variability into systematic components, and residual variability that might reflect noise at the level of the underlying neural computations (Glimcher, 2005; Renart and Machens, 2014).

The lack of a correlation between history-dependent starting point shifts and individual choice history bias is surprising in light of previous accounts (Yu and Cohen, 2008; Gao et al., 2009; Glaze et al., 2015). History shift in starting point were mostly negative (i.e., tendency towards choice alternation) across participants, regardless of their individual tendency towards choice repetition or alternation (Figure S4, significant in 2 out of five data sets). This small but consistent effect likely explains why our formal model comparison favored a model with history-dependent shifts in both, drift and starting point over one with

shifts in drift only. Critically, however, in all five data sets only the history-dependent shift in drift accounted for the individual differences in participants' overt choice history biases. Further, the model with only history-dependent shift in drift accounted for the diagnostic features of RT distributions just as well as the model with shifts in both starting point and drift (Figure 4d, compare two most rightward bars).

The DDM is a valid reduction of biophysically plausible (high-dimensional) decision-making models for certain parameter regimes (Bogacz et al., 2006; Wong and Wang, 2006). But in more general settings it may fail to capture subtle patterns in the behavioral data. The lower dimensionality of the DDM enables fits to behavioral data, a prerequisite for the insights gained here. It remains to be seen whether alternative reductions of biophysical decision-making models (Roxin and Ledberg, 2008) might identify similar history-dependent changes in evidence accumulation, possibly without shifts in starting point. Further, it is possible that on a subset of trials, observers' behavior results from processes not incorporated in the bounded accumulation framework underlying the DDM – for example, fast guesses (Noorbaloochi et al., 2015), automatic decision processing (Servant et al., 2014; Ulrich et al., 2015), or post-accumulation bias (Erlich et al., 2015).

In bounded accumulation models like the DDM, dynamic changes in the decision variable can be mimicked by dynamic changes in the decision bounds during the course of the trial (Supplementary Figure 6). Specifically, simulations show that the combination of a linearly collapsing bound (for the favored choice) and a linearly expanding bound (for the non-favored choice) has the same effect on choice fractions and RTs as the drift criterion effect which we focused on here (Supplementary Figure 6, compare b and c). We favor history-dependent shifts in the drift of the decision variable as a mechanistic explanation for our observations, for several reasons. First, we are not aware of any evidence, neural or behavioral, for asymmetric changes in decision bounds that would be necessary to implement choice bias. Data from other top-down modulations of the decision process show that premotor build-up signals in the brain reach a threshold level of firing just before motor response, that remains fixed across manipulations of prior probability (Hanks et al., 2011), and even under manipulations of speed-accuracy trade-off (Hanks et al., 2014; Murphy et al., 2016). Instead, the dynamics of these neural build-up signals are altered, in line with dynamic changes in drift. Second, a large body of evidence on attention establishes that the brain is equipped with a powerful mechanism for biasing the evidence accumulation process, across multiple stages of the cortical hierarchy: selectively modulating the gain of neuronal responses to specific features of the sensory input (Maunsell and Treue, 2006). In light of this previous evidence as well as our current model fits, we propose the history-dependent shift in the drift as the most parsimonious explanation of choice history biases.

Our interpretation has direct implications for the candidate neural mechanisms underlying choice history bias, which can help guide future neurophysiological work. Shifts in starting point reflect additive effects at the stage of (Platt and Glimcher, 1999) or downstream from (Basso and Wurtz, 1997) the evidence accumulator, likely comprising posterior parietal (Gold and Shadlen, 2007; but see Brody and Hanks, 2016) or motor cortex (Pape and Siegel, 2016). By contrast, changes in drift bias can come about either through a selective modulation of the sensory representation that feeds into the accumulator (i.e., feature-based attention, see above), or through a change in the read-out of this sensory

representation by the accumulator (i.e., the effective connectivity between sensory cortices and downstream brain regions). The robust effect on drift bias we observed is clearly in line with the latter scenario. It indicates that choice history signals could bias sensory representations, possibly through feedback connections from decision-related areas to sensory cortex (Nienborg and Cumming, 2009; Wimmer et al., 2015; St. John-Saaltink et al., 2016). This is in line with the idea that choice history signals bias perceptual decision-making in a manner akin to top-down (feature-based) attention: previous choices may act as endogenous 'cues' that direct attention towards or away from previously chosen options.

Materials and Methods

Data sets: Behavioral tasks and participants

We analyzed five different data sets, four of which were previously published. These spanned different modalities (visual or auditory), decision-relevant features (motion direction, contrast, tone presence, motion coherence), and tasks (detection or discrimination). Those tasks where the decision-relevant sensory evidence was presented until the observer generated a response were called reaction time (RT) tasks; those tasks where the sensory evidence is presented for a fixed duration, and its offset cues the observer's response, were called fixed duration (FD) tasks in line with the terminology from Mazurek et al. (2003). These two protocols have also been termed 'free response protocol' and 'interrogation protocol', respectively (Bogacz et al., 2006). In all data sets, stimulus strength (i.e., decision difficulty) was kept constant, or varied systematically across five levels, within all main experimental sessions that were used for fitting the DDM.

2AFC visual motion direction discrimination task (RT)

These data were previously published (Murphy et al., 2014), and are available at https://doi.org/10.5061/dryad.tb542. Twenty-six observers (22 women and 4 men, aged 18-29) performed a motion direction (left vs. right) discrimination task. Stationary white dots were presented on a black screen for an interval of 4.3-5.8 s. After this fixation interval, the decision-relevant sensory evidence was presented: some percentage of dots (the 'motion coherence' level) moved to the left or the right. The coherence was individually titrated to yield an accuracy level of 85% correct, estimated from a psychometric function fit, before the start of the main experiment, and kept constant afterwards. The moving dots were presented until observers indicated their choice with a button press. After the response, the fixation cross changed color for 700 ms to indicate single-trial feedback. Each observer performed 500 trials of the task (one session). We refer to this task as 'Visual motion 2AFC (RT)'.

Visual contrast yes/no detection task (RT)

These data were previously published (de Gee et al., 2014), and are available at https://doi.org/10.6084/m9.figshare.4806559. Twenty-nine observers (14 women and 15 men, aged 18–38) performed a yes/no contrast detection task. During a fixation interval of 4-6 seconds, observers viewed dynamic noise (a binary noise pattern that was refreshed each frame, at 100 Hz). A beep indicated the start of the decision-relevant sensory evidence: on half the trials, a vertical grating was superimposed onto the dynamic noise; on the other half

of trials, only the dynamic noise was shown. The sensory evidence (signal+noise or noise-only) was presented until the observers reported their choice (yes, grating was present; or no, grating was absent), or after a maximum of 2.5s. The signal contrast was individually titrated to yield an accuracy level of 75% correct using a method of constant stimuli before the main experiment, and kept constant throughout the main experiment. Observers performed between 480–800 trials over 6-10 sessions. Six observers in the original paper (de Gee et al., 2014) performed a longer version of the task in which they also reported their confidence levels and received feedback; these were left out of the current analysis, leaving twenty-three subjects to be included. We refer to this task as 'Visual contrast yes/no (RT)'.

AUDITORY TONE YES/NO DETECTION TASK (RT)

These data were previously published (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. After an inter-trial interval of 3-4 seconds, decision-relevant sensory evidence was presented: on half the trials, a sine wave (2 KHz) superimposed onto dynamic noise (so-called TORCs, (McGinley et al., 2015) was presented); on the other half of trials only the dynamic noise was presented. The sensory evidence was presented until the participant reported their choice button press or after a maximum of 2.5s. No feedback was provided. Each individual's signal volume was titrated to an accuracy level of 75% correct using an adaptive staircase procedure before the start of the main experiment, and kept constant throughout the main experiment. Participants performed between 1320 and 1560 trials each, divided over two sessions. We refer to this task as 'Auditory yes/no (RT)'.

VISUAL MOTION COHERENCE DISCRIMINATION 2IFC TASK (FD): DATA SET 1

These data were previously published in (Urai et al., 2017), and are available at http://dx.doi.org/10.6084/m9.figshare.4300043. Twenty-seven observers (17 women and 10 men, aged 18-43) performed a two-interval motion coherence discrimination task. They viewed two consecutive intervals of random dot motion, containing coherent motion signals in a constant direction towards one of the four diagonals (counterbalanced across participants) and judged whether the second interval (variable coherence) contained stronger or weaker motion than the first (constant coherence) interval. After a fixation interval of 0.5-1s, they viewed two consecutive intervals of 500 ms each, separated by a delay of 300-700 ms. The decision-relevant sensory evidence (i.e., the difference in motion coherence between intervals), was chosen pseudo-randomly for each trial from the set (0.625, 1.25, 2.5, 5, 10, 20, 30%). Observers received auditory feedback on their choice after a delay of 1.5-2.5s. After continuing to view noise dots for 2-2.5 s, stationary dots indicated an inter-trial interval. Observers self-initiated the start of the next trial (range of median inter-trial intervals across observers: 0.68–2.05 s). Each observer performed 2500 trials of the task, divided over five sessions. We refer to this task as 'Visual motion 2IFC (FD) #1'.

2IFC VISUAL MOTION COHERENCE DISCRIMINATION TASK (FD): DATA SET 2

These data were not previously published, so we here provide a more detailed description of the participants and informed consent, in addition to the task protocol.

Participants and informed consent. Sixty-one 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.

Task protocol. 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 behavioral lab (1500 trials divided over 15 blocks per session). The task was as described above for 'Visual motion 2IFC (FD) #1', with the following exceptions. The strength of the decision-relevant sensory evidence was individually titrated to an accuracy level of 70% correct, estimated from a psychometric function fit, before the start of the main experiment and kept constant for each individual throughout the main experiment. In the MEG sessions, auditory feedback was presented 1.5-3 s after response, and an inter-trial interval with stationary dots started 2-3 s after feedback. Participants initiated the next trial with a button press (across-subject range of median inter-trial interval duration: 0.64 to 2.52 s, group average: 1.18 s). In the training sessions, auditory feedback was presented immediately after the response. This was followed by an inter-trial interval of 1 s, after which the next trial started. We refer to this task as 'Visual motion 2IFC (FD) #2'. In this experiment, three sub-groups of observers received different pharmacological treatments prior to each session, receiving placebo, atomoxetine (a noradrenaline reuptake inhibitor), or donepezil (an acetylcholinesterase inhibitor). See Figure S5a for a description of the treatment. These groups did not differ in their choice history bias and were pooled for the purpose of the present study (Figure S5b).

MODEL-FREE ANALYSIS OF SENSITIVITY AND CHOICE HISTORY BIAS

We quantified perceptual sensitivity in terms of signal detection-theoretic d' (Green and Swets, 1966):

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

where Φ was the normal cumulative distribution function, H was the fraction of hits and FA the fraction of false alarms. In the 2AFC and 2IFC data sets, one of the two stimulus categories was arbitrarily treated as signal absent. 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 estimated d' separately for each individual and, for the visual motion 2IFC (FD) #1 data, for each level of sensory evidence.

We quantified individual choice history bias in terms of the probability of repeating a choice, termed P(repeat), regardless of the category of the (previous or current) stimulus. This yielded a measure of bias that ranged between 0 (maximum alternation bias) and 1 (maximum repetition bias), whereby 0.5 indicated no bias.

DRIFT DIFFUSION MODEL (DDM) FITS

DDM: GENERAL

This section describes the general DDM, with a focus on the biasing mechanisms described in Results and illustrated in Figure 1 (Ratcliff and McKoon, 2008). In the DDM, the accumulation of noisy sensory evidence is given by a drift diffusion process:

$$dy = s \cdot v \cdot dt + cdW \quad (2)$$

where y is the decision variable (gray example traces in Figure 1), s is the stimulus category (coded as [-1,1]), *v* is the drift rate, and cdW is Gaussian distributed white noise with mean 0 and variance c^2dt (Bogacz et al., 2006). In an unbiased case, the starting point of the decision variably $y(0) = z$, is situated midway between the two decision bounds 0 and a:

$$y(0) = z = \frac{a}{2} \quad (3)$$

where a is the separation between the two decision bounds. A bias in the starting point is implemented by an additive offset z_{bias} from the midpoint between the two bounds (Figure 1a):

$$y(0) = z = \frac{a}{2} + z_{bias} \quad (4)$$

A drift bias can be implemented by adding a stimulus-independent constant v_{bias}, also referred to as drift bias (Ratcliff and McKoon, 2008), to the (stimulus-dependent) mean drift (Figure 1b):

$$dy = (s \cdot v + v_{bias})dt + cdW \quad (5)$$

This adds a bias to the drift that linearly grows with time. Previous work has used non-linear changes in the decision variable to account for choice biases in case of unequal prior probability of the two choice options (Hanks et al., 2011). In the absence of specific experimental manipulations or data that might help constrain the form of nonlinearities in the bias, we here opted for fitting a linear drift bias for the sake of simplicity. While we acknowledge the possibility that the bias dynamics might be non-linear also in the case of history biases, our choice was supported by the good fits of the DDM to all data sets.

These two biasing mechanisms result in the same (asymmetric) fraction of choices, but they differ in terms of the resulting shapes of RT distributions (Figure 1). We allowed both bias parameters to vary as a function of observers' previous choice, to test their relative contributions to the individual differences in overt choice history biases.

ESTIMATING DDM BIAS PARAMETERS

We used hierarchical drift diffusion modeling as implemented in the HDDM toolbox (Wiecki et al., 2013) to fit the model and estimate its parameters. RTs faster than 250 ms were discarded from the model fits. As recommended by the HDDM toolbox, 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 (Ratcliff and Tuerlinckx, 2002). We fit the DDM to RT distributions for the two choice categories, conditioned on the stimulus category for each trial (s in eq. 2) - a procedure referred to as 'stimulus coding'. This fitting method deviates from a widely used expression of the model, where RT distributions for correct and incorrect choices are fit (also called 'accuracy coding'). Only the former can fit decision biases towards one choice over the other.

First, we estimated a model without history-dependent bias parameters. Overall drift rate, boundary separation, non-decision time, starting point, and drift bias were estimated for each individual (Figure S1). Across-trial variability in drift rate was estimated at the group-level only (Ratcliff and Childers, 2015). For the data set including variations of sensory evidence strength (Visual motion 2IFC (FD) #1, see above), we allowed drift rate to vary with

evidence strength. This model was used to confirm that the DDM was able to fit all data sets well, and to serve as a baseline model when comparing diagnostic features.

Second, we estimated three different models of history bias, allowing (i) starting point, (ii) drift or (iii) both to vary as a function of the observer's immediately preceding choice (thus capturing only so-called first-order sequential effects; cf Gao et al., 2009; Wilder et al., 2009). The effect of the preceding choice on each bias parameter was then termed its 'history shift'. For example, for the visual motion direction discrimination task we separately estimated the starting point parameter for trials following 'left' and 'right' choices. The difference between these two parameters then reflects individual observers' history shift in starting point, computed such that a positive value reflected a tendency towards repetition and a negative value reflected a tendency towards alternation. The history shift in drift bias was computed in the same way.

MODEL FITTING PROCEDURES

The HDDM (Wiecki et al., 2013) uses Markov-chain Monte Carlo sampling for generating posterior distributions over model parameters (Andrieu et al., 2003). Two features of this method deviate from more standard model optimization. First, the Bayesian MCMC generates full posterior distributions over parameter estimates, quantifying not only the most likely parameter value but also the uncertainty associated with that estimate (see e.g. Figure S4). Second, the hierarchical nature of the model assumes that all observers in a dataset are drawn from a group, with specific group-level prior distributions that are informed by the literature (Figure S1 and Wiecki et al., 2013). In practice, this results in more stable parameter estimates for individual subjects, who are constrained by the group-level inference. Note that we also repeated our model fits with more traditional G^2 optimization (Ratcliff and Tuerlinckx, 2002) and obtained qualitatively identical results.

For each variant of the model, we ran 15 separate Markov chains with 10000 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 2500 samples per chain, which were concatenated across chains. Individual parameter estimates were then estimated from the posterior distributions across the resulting 37500 samples. All group-level chains were visually inspected to ensure convergence. Additionally, we computed the Gelman-Rubin $\hat{R}$ statistic (which compares within-chain and between-chain variance) and checked that all group-level parameters had an $\hat{R}$ between 0.98-1.02.

Formal comparison between the different model variants (see above) was performed using the Deviance Information Criterion (Spiegelhalter et al., 2002), a commonly used method for assessing the goodness of fit in hierarchical models, for which a unique 'likelihood' is not defined, and the effective number of degrees of freedom is often unclear. Lower DIC values indicate a better fit, while taking into account the complexity of each model. A difference in DIC values of more than 10 is considered evidence for the winning model to reflect the data significantly better.

CONDITIONAL BIAS FUNCTIONS

For each variant of the model and each data set, we simulated data using the best-fitting parameters. Specifically, we simulated 100 responses (choices and RTs) for each trial

performed by the observers. These predicted patterns for the 'baseline model' (without history-dependent bias parameters) were first used to compare the observed and predicted patterns of choices and RTs (Figure 3b).

We also used the simulated data, as well as the participants' choices and RTs, to visualize specific features in our data that distinguish the different biased models (Palminteri et al., 2017). Specifically, we computed conditional bias functions (White and Poldrack, 2014) that visualize choice history bias as a function of RTs. Each choice was recoded into a repetition (1) or alternation (0) of the previous choice. We then expressed each choice as being either in line with or against the observer's individual bias (classified into 'repeaters' and 'alternators' depending on choice repetition probability). This allowed us to visualize the effect of choice history bias as a function of time within each trial, together for all observers in each dataset.

To generate conditional bias functions, we divided each (simulated or real) observer's RT distribution into five quantiles (0.1, 0.3, 0.5, 0.7 and 0.9) and computed the fraction of choices within each quantile. The shape of the conditional bias functions for models with z_{bias} and v_{bias} confirm that z_{bias} predominantly produces biased choices with short RTs, whereas v_{bias} leads to biased choices across the entire range of RTs (Figure 4b).

STATISTICAL TESTS

We quantified across-subject correlations between P(repeat) and the individual history components in DDM bias parameter estimates using Pearson's correlation coefficient. 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, P(repeat) - without inflation of the false positive rate (Katahira, 2016). The difference between two correlation coefficients that shared a common variable, and its associated p-value, was computed using Steiger's test (Steiger, 1980).

We used Bayes factors to quantify the strength of evidence across our different data sets. We first computed the Bayes factor for each correlation (between P(repeat) and the history shift in starting point, and between P(repeat) and the history shift in drift bias) (Wetzels and Wagenmakers, 2012). We then multiplied these Bayes factors across data sets to quantify the total evidence in favor or against the null hypothesis of no correlation (Scheibehenne et al., 2016). BF_{10} quantifies the evidence in favor of the null or the alternative hypothesis, where $BF_{10} = 1$ indicates inconclusive evidence to draw conclusions from the data. $BF_{10} < 1/10$ or > 10 is taken to indicate substantial evidence for H_0 or H_1, respectively (Kass and Raftery, 1995).

CODE AVAILABILITY

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

AUTHOR CONTRIBUTIONS

AEU, Conceptualization, Investigation, Formal Analysis, Software, Visualization, Writing – Original Draft, Writing – Review and Editing; JWdG, Investigation, Formal Analysis, Writing –

Review and Editing; THD, Conceptualization, Resources, Supervision, Writing – Original Draft, Writing – Review and Editing.

ACKNOWLEDGEMENTS

We thank Gilles de Hollander and Peter Murphy for discussion. Christiane Reissmann, Karin Deazle, Samara Green and Lina Zakarauskaite helped with participant recruitment and data acquisition for the 2IFC #2 MEG study. This research was supported by the German Academic Exchange Service (DAAD) (to A.E.U.) and the German Research Foundation (DFG) grants DO 1240/2-1, DO 1240/3-1, SFB 936/A7, and SFB 936/Z1 (to T.H.D.).

BIBLIOGRAPHY

Abrahamyan A, Silva LL, Dakin SC, Carandini M, Gardner JL (2016) Adaptable history biases in human perceptual decisions. Proc Natl Acad Sci 113:E3548–E3557.

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

Allefeld C, Soon CS, Bogler C, Heinzle J, Haynes J-D (2013) Sequential dependencies between trials in free choice tasks. arxiv.org.

Andrieu C, Freitas N de, Doucet A, Jordan MI (2003) An Introduction to MCMC for Machine Learning. Mach Learn 50:5–43.

Basso MA, Wurtz RH (1997) Modulation of neuronal activity by target uncertainty. Nature 389:66–69.

Bode S, Sewell DK, Lilburn S, Forte JD, Smith PL, Stahl J (2012) Predicting Perceptual Decision Biases from Early Brain Activity. J Neurosci 32:12488–12498.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev 113:700–765.

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.

Braun A, Urai AE, Donner TH (2018) Adaptive history biases result from confidence-weighted accumulation of past choices. J Neurosci:2189–17.

Brody CD, Hanks TD (2016) Neural underpinnings of the evidence accumulator. Curr Opin Neurobiol 37:149–157.

Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Schölvinck ML, Zaharia AD, Carandini M (2011) The Detection of Visual Contrast in the Behaving Mouse. J Neurosci 31:11351–11361.

Chamberlain SR, Hampshire A, Müller U, Rubia K, del Campo N, Craig K, Regenthal R, Suckling J, Roiser JP, Grant JE, Bullmore ET, Robbins TW, Sahakian BJ (2009) Atomoxetine Modulates Right Inferior Frontal Activation During Inhibitory Control: A Pharmacological Functional Magnetic Resonance Imaging Study. Biol Psychiatry 65:550–555.

Cho RY, Nystrom LE, Brown ET, Jones AD, Braver TS, Holmes PJ, Cohen JD (2002) Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. Cogn Affect Behav Neurosci 2:283–299.

de Gee JW, Colizoli O, Kloosterman NA, Knapen T, Nieuwenhuis S, Donner TH (2017) Dynamic modulation of decision biases by brainstem arousal systems. eLife 6:e23232.

de Gee JW, Knapen T, Donner TH (2014) Decision-related pupil dilation reflects upcoming choice and individual bias. Proc Natl Acad Sci 111:E618–E625.

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

Edwards W (1965) Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. J Math Psychol 2:312–329.

Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD (2015) Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. eLife 4.

Fernberger SW (1920) Interdependence of judgments within the series for the method of constant stimuli. J Exp Psychol 3:126.

Fründ I, Wichmann FA, Macke JH (2014) Quantifying the effect of intertrial dependence on perceptual decisions. J Vis 14:9–9.

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

Gao J, Wong-Lin K, Holmes P, Simen P, Cohen JD (2009) Sequential Effects in Two-Choice Reaction Time Tasks: Decomposition and Synthesis of Mechanisms. Neural Comput 21:2407–2436.

Glaze CM, Kable JW, Gold JI (2015) Normative evidence accumulation in unpredictable environments. eLife 4:e08825.

Glimcher PW (2005) Indeterminacy in Brain and Behavior. Annu Rev Psychol 56:25–56.

Gold JI, Law C-T, Connolly P, Bennur S (2008) The Relative Influences of Priors and Sensory Evidence on an Oculomotor Decision Variable During Perceptual Learning. J Neurophysiol 100:2653–2668.

Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. Annu Rev Neurosci 30:535–574.

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

Hanks T, Kiani R, Shadlen MN (2014) A neural mechanism of speed-accuracy tradeoff in macaque area LIP. eLife 3:e02260.

Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J Neurosci 31:6339–6352.

Hwang EJ, Dahlen JE, Mukundan M, Komiyama T (2017) History-based action selection bias in posterior parietal cortex. Nat Commun 8:1242.

Jahfari S, Verbruggen F, Frank MJ, Waldorp LJ, Colzato L, Ridderinkhof KR, Forstmann BU (2012) How preparation changes the need for top-down control of the basal ganglia when inhibiting premature actions. J Neurosci Off J Soc Neurosci 32:10870–10878.

Jones M, Curran T, Mozer MC, Wilder MH (2013) Sequential effects in response time reveal learning mechanisms and event representations. Psychol Rev 120:628.

Kass RE, Raftery AE (1995) Bayes Factors. J Am Stat Assoc 90:773–795.

Katahira K (2016) How hierarchical models improve point estimates of model parameters at the individual level. J Math Psychol 73:37–58.

Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J Neurosci 28:3017–3029.

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

Kruschke JK (2013) Bayesian estimation supersedes the t test. J Exp Psychol Gen 142:573.

Leite FP, Ratcliff R (2011) What cognitive processes drive response biases? A diffusion model analysis. Judgm Decis Mak 6:651–687.

Leopold DA, Wilke M, Maier A, Logothetis NK (2002) Stable perception of visually ambiguous patterns. Nat Neurosci 5:605–609.

Maunsell JHR, Treue S (2006) Feature-based attention in visual cortex. Trends Neurosci 29:317–322.

Mazurek ME (2003) A Role for Neural Integrators in Perceptual Decision Making. Cereb Cortex 13:1257–1269.

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

Mulder MJ, Wagenmakers E-J, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. J Neurosci 32:2335–2343.

Murphy PR, Boonstra E, Nieuwenhuis S (2016) Global gain modulation generates time-dependent urgency during perceptual choice in humans. Nat Commun 7:13526.

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

Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang X-J (2014) A hierarchy of intrinsic timescales across primate cortex. Nat Neurosci 17:1661–1663.

Nienborg H, Cumming BG (2009) Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature 459:89–92.

Noorbaloochi S, Sharon D, McClelland JL (2015) Payoff Information Biases a Fast Guess Process in Perceptual Decision Making under Deadline Pressure: Evidence from Behavior, Evoked Potentials, and Quantitative Model Comparison. J Neurosci 35:10989–11011.

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

Palminteri S, Wyart V, Koechlin E (2017) The Importance of Falsification in Computational Cognitive Modeling. Trends Cogn Sci 21:425–433.

Pape A-A, Siegel M (2016) Motor cortex activity predicts response alternation during sensorimotor decisions. Nat Commun 7:13098.

Pfeffer T, Avramiea A-E, Nolte G, Engel AK, Linkenkaer-Hansen K, Donner TH (2017) Catecholamines, not acetylcholine, alter cortical and perceptual dynamics in line with increased excitation-inhibition ratio. bioRxiv 170613.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature 400:233–238.

Purcell BA, Kiani R (2016) Neural Mechanisms of Post-error Adjustments of Decision Policy in Parietal Cortex. Neuron 89:658–671.

Ratcliff R (2006) Modeling response signal and response time data. Cognit Psychol 53:195–237.

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

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput 20:873–922.

Ratcliff R, Tuerlinckx F (2002) Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. Psychon Bull Rev 9:438–481.

Renart A, Machens CK (2014) Variability in neural activity and behavior. Curr Opin Neurobiol 25:211–220.

Rogers SL, Friedhoff LT (1998) Pharmacokinetic and pharmacodynamic profile of donepezil HCl following single oral doses. Br J Clin Pharmacol 46.

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

Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of Sensory and Reward Information during Perceptual Decision-Making in Lateral Intraparietal Cortex (LIP) of the Macaque Monkey. PLOS ONE 5:e9308.

Roxin A, Ledberg A (2008) Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. PLoS Comput Biol 4:e1000046.

Sauer J-M, Ring BJ, Witcher JW (2005) Clinical pharmacokinetics of atomoxetine. Clin Pharmacokinet 44:571–590.

Scheibehenne B, Jamil T, Wagenmakers E-J (2016) Bayesian Evidence Synthesis Can Reconcile Seemingly Inconsistent Results: The Case of Hotel Towel Reuse. Psychol Sci 27:1043–1046.

Scott BB, Constantinople CM, Akrami A, Hanks TD, Brody CD, Tank DW (2017) Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. Neuron.

Servant M, Montagnini A, Burle B (2014) Conflict tasks and the diffusion framework: Insight in model constraints based on psychological laws. Cognit Psychol 72:162–195.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci 16:1486–1510.

Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A (2002) Bayesian measures of model complexity and fit. J R Stat Soc Ser B Stat Methodol 64:583–639.

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. J Neurosci 36:6186–6192.

Stanislaw H, Todorov N (1999) Calculation of signal detection theory measures. Behav Res Methods Instrum Comput 31:137–149.

Steiger JH (1980) Tests for comparing elements of a correlation matrix. Psychol Bull 87:245.

Thura D, Guberman G, Cisek P (2016) Trial-to-trial adjustments of speed-accuracy trade-offs in premotor and primary motor cortex. J Neurophysiol 117:665–683.

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

Ulrich R, Schröter H, Leuthold H, Birngruber T (2015) Automatic and controlled stimulus processing in conflict tasks: Superimposed diffusion processes and delta functions. Cognit Psychol 78:148–174.

Urai AE, Braun A, Donner TH (2017) Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. Nat Commun 8:14637.

van Ravenzwaaij D, Mulder MJ, Tuerlinckx F, Wagenmakers E-J (2012) Do the dynamics of prior information depend on task context? An analysis of optimal performance and an empirical test. Front Psychol 3:132.

Wei W, Wang X-J (2016) Inhibitory Control in the Cortico-Basal Ganglia-Thalamocortical Loop: Complex Regulation and Interplay with Memory and Decision Processes. Neuron 92:1093–1105.

Wetzels R, Wagenmakers E-J (2012) A default Bayesian hypothesis test for correlations and partial correlations. Psychon Bull Rev 19:1057–1064.

White CN, Poldrack RA (2014) Decomposing bias in different types of simple decisions. J Exp Psychol Learn Mem Cogn 40:385–398.

Wiecki TV, Sofer I, Frank MJ (2013) HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. Front Neuroinformatics 7.

Wilder M, Jones M, Mozer MC (2009) Sequential effects reflect parallel learning of multiple environmental regularities. In: Advances in Neural Information Processing Systems 22 (Bengio Y, Schuurmans D, Lafferty JD, Williams CKI, Culotta A, eds), pp 2053–2061. Curran Associates, Inc. Available at: http://papers.nips.cc/paper/3870-sequential-effects-reflect-parallel-learning-of-multiple-environmental-regularities.pdf [Accessed November 10, 2015].

Wimmer K, Compte A, Roxin A, Peixoto D, Renart A, de la Rocha J (2015) Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. Nat Commun 6:6177.

Wong K-F, Wang X-J (2006) A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. J Neurosci 26:1314–1328.

Wyart V, Koechlin E (2016) Choice variability and suboptimality in uncertain environments. Curr Opin Behav Sci 11:109–115.

Yu AJ, Cohen JD (2008) Sequential effects: Superstition or rational behavior? Adv Neural Inf Process Syst 21:1873–1880.

Zhang S, Huang CH, Yu AJ (2014) Sequential effects: a Bayesian analysis of prior bias on reaction time and behavioral choice. In: Proceedings of the 36th Annual Conference of the Cognitive Science Society Available at: https://mindmodeling.org/cogsci2014/papers/320/paper320.pdf [Accessed August 5, 2016].

SUPPLEMENTARY FIGURES

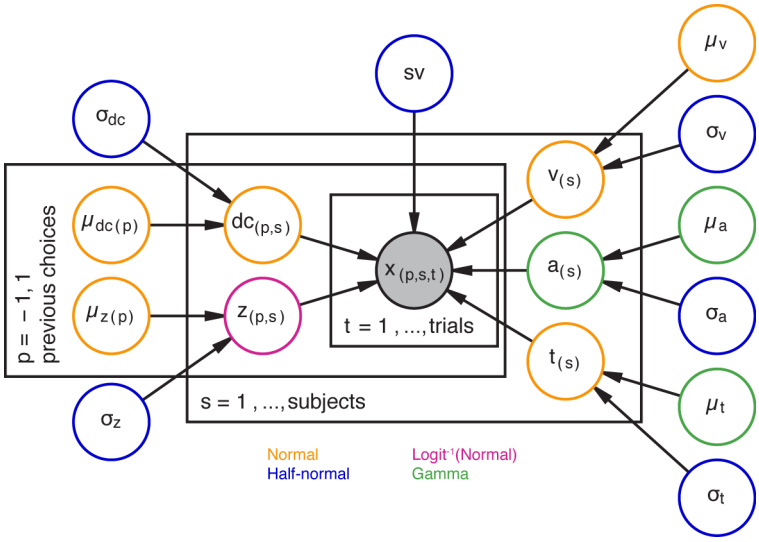

Figure S1. Graphical representation of the hierarchical model structure. The full model (with both history-dependent drift bias and starting point) is depicted. Round nodes represent random variables, and the shaded node x represented the observed data (choices and RTs for all observers within each task). Subject-specific parameter estimates were distributed according to the group-level posterior values, thereby 'shrinking' individual values towards the group average. Colors indicate the distributions used for each node; see (Wiecki et al., 2013) for the specific parameters defining each distribution. Note that for the 2IFC #1 dataset, we additionally estimated a separate drift rate (v) for each difficulty level (Figure 3a, inset).

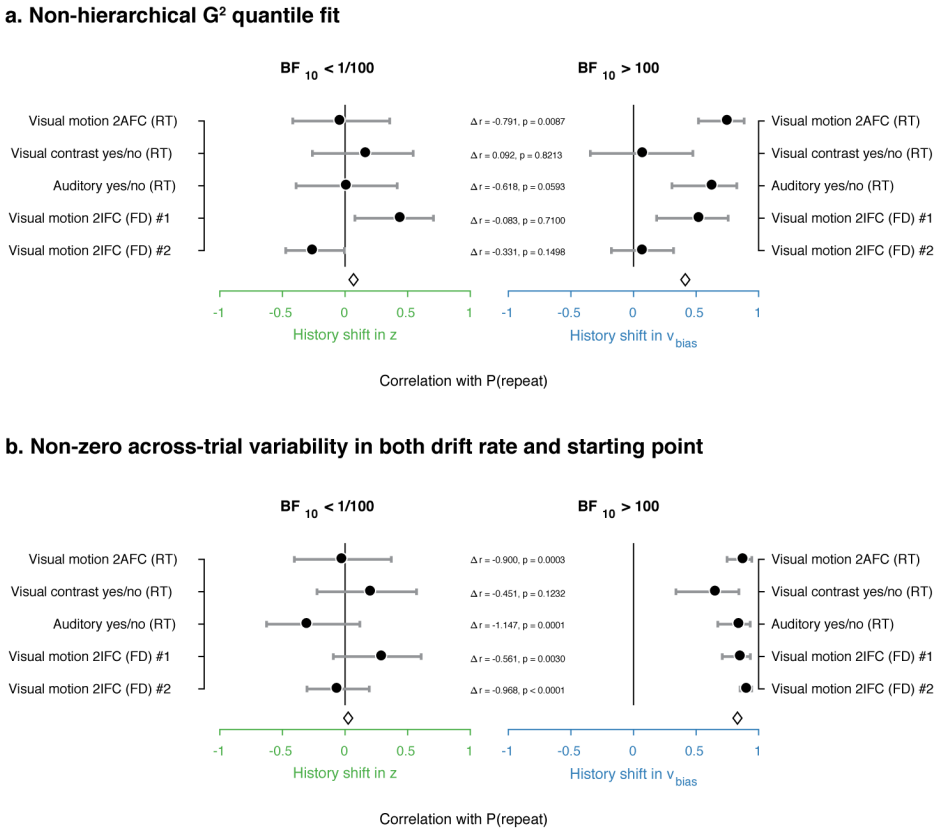

Figure S3. Control model fits. (**a**) Summary figure based on non-hierarchical G^2 fits (Ratcliff and Tuerlinckx, 2002). (**b**) Summary figure based on the full hierarchical model, where across-trial variability in starting point sz was added as a free parameter. Like the across-trial variability in drift rate sv, the sz parameter was only estimated at the group level (Ratcliff and Childers, 2015).

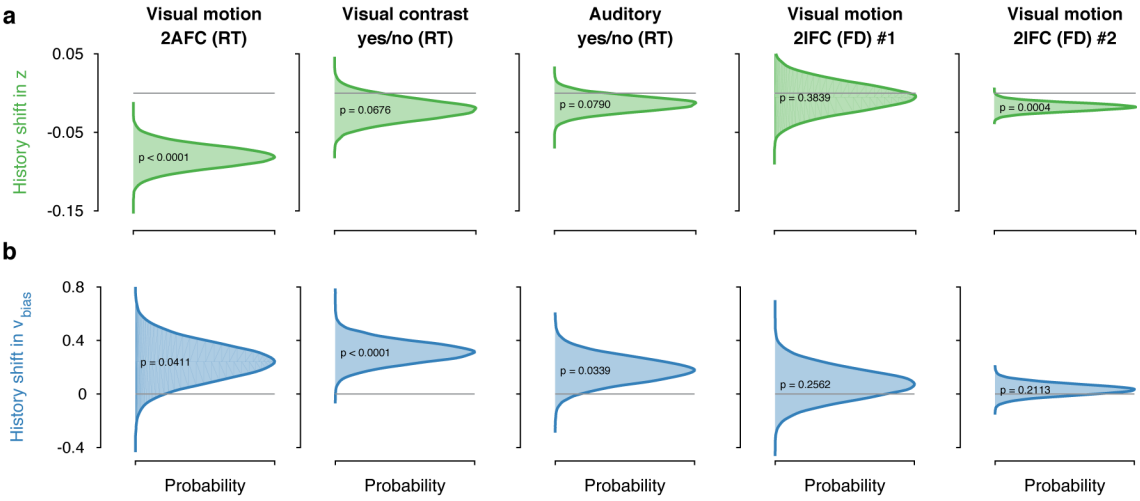

Figure S4. Group-level posterior distributions of history bias parameters. *History shift in (**a**) starting point and (**b**) drift bias, for each dataset. P-values were derived directly from the group posterior (Kruschke, 2013).*

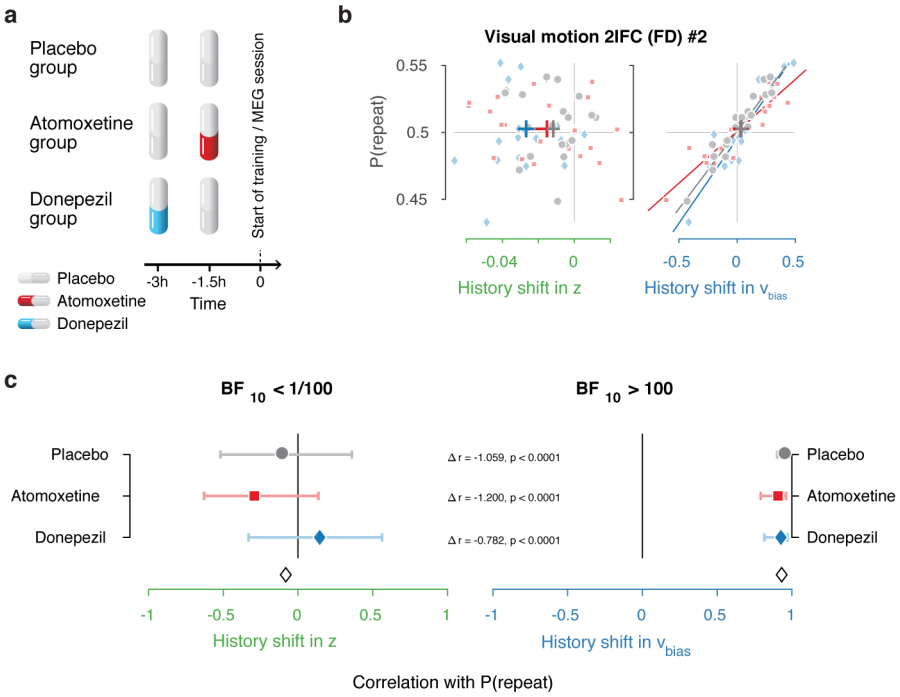

Figure S5. Same biasing mechanism under two pharmacological interventions. *(**a**) Participants in the MEG study were assigned to one of three pharmacological groups. At the start of each experimental session, they orally took 40 mg atomoxetine (Strattera®), 5 mg donepezil (Aricept®), or placebo. Since 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 placebo 1.5 hours before starting the experimental session. Participants in the atomoxetine group took 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. Schematic adapted from (Pfeffer et al., 2017) with permission. (**b, c**) Choice history biases separately for each pharmacological group. Since we did not observe differences in choice history bias between these groups, we pooled all observers for the main analyses.*

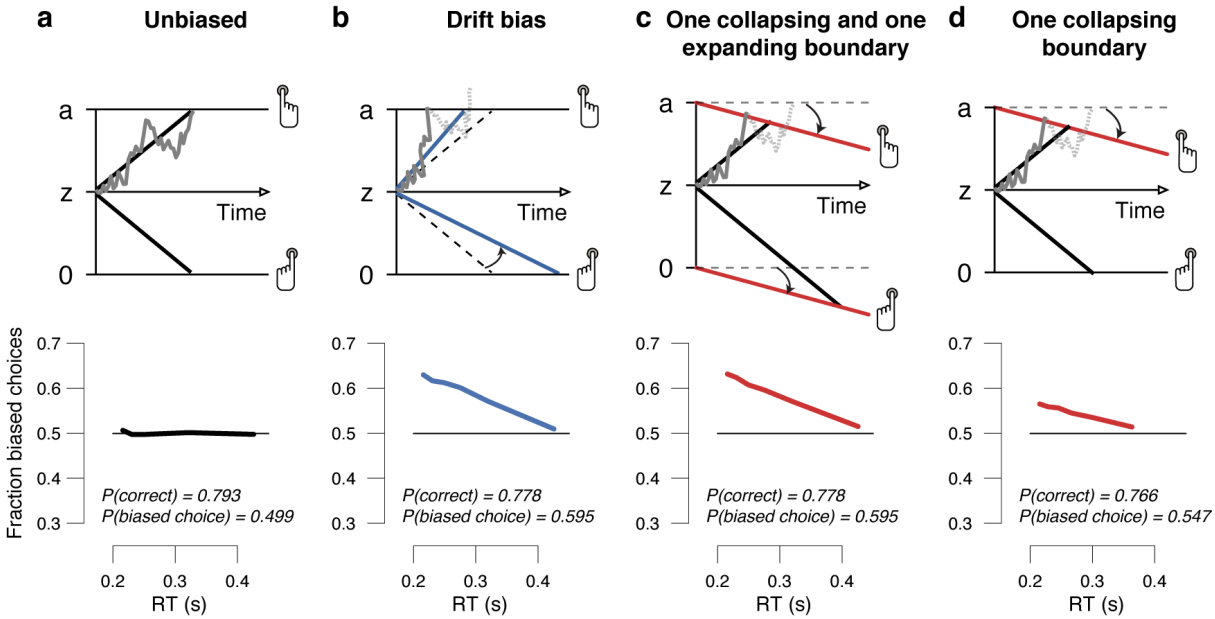

Figure S6. Drift bias and dynamic changes in decision bounds can produce indistinguishable behavioral results. *Time-dependent changes in the decision variable (because of drift bias) are equivalent to time-dependent changes in the decision bounds. Our simulations (200K trials per model) reveal that these different biasing mechanisms produce similar changes in the fraction of biased choices as a function of RT (conditional response functions, see Methods). (**a**) Top: Unbiased model. Black lines: mean drift under unbiased conditions. Bottom: Condition response function for unbiased model. (**b**) As panel (**a**), but for drift bias model (see also Figure 1). Dashed lines: mean drift rate. Blue lines: mean drift under bias drift (drift rate plus drift bias). (**c**) As panel (**a**), but for one collapsing and one expanding boundary model. Red lines: time-dependent changes in the boundary (upper, collapsing; bottom, expanding). (**d**) As in (**a**), but for one collapsing boundary model. Red line: time dependent change in the upper boundary (collapsing). All panels: Gray line: example trajectory of decision variable from single trial. The numbers printed in the condition response function plots are overall probabilities, regardless of RT.*