

# Current Biology

## Confirmation Bias through Selective Overweighting of Choice-Consistent Evidence

### Highlights

- People's interpretation of new evidence is often biased by their previous choices
- Talluri, Urai et al. developed a new task for probing the underlying mechanisms
- Evidence consistent with an observer's initial choice is processed more efficiently
- This "choice-induced gain change" affects both perceptual and numerical decisions

### Authors

Bharath Chandra Talluri, Anne E. Urai, Konstantinos Tsetsos, Marius Usher, Tobias H. Donner

### Correspondence

bharathchandra.talluri@gmail.com (B.C.T.),  
t.donner@uke.de (T.H.D.)

### In Brief

Committing to a categorical choice biases subsequent decision-making, a phenomenon called confirmation bias. Talluri, Urai et al. developed a new behavioral task to probe into the underlying mechanism. They find that choices selectively increased the weighting of choice-consistent evidence on subsequent decisions, in a way resembling attention.



# Confirmation Bias through Selective Overweighting of Choice-Consistent Evidence

Bharath Chandra Talluri,<sup>1,5,\*</sup> Anne E. Urai,<sup>1,2,5</sup> Konstantinos Tsetsos,<sup>1,6</sup> Marius Usher,<sup>3,4,6</sup> and Tobias H. Donner<sup>1,2,6,7,\*</sup>

<sup>1</sup>Department of Neurophysiology & Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany

<sup>2</sup>Department of Psychology, University of Amsterdam, 1018 WS Amsterdam, the Netherlands

<sup>3</sup>School of Psychology, Tel-Aviv University, 69989 Ramat-Aviv, Tel-Aviv, Israel

<sup>4</sup>Sagol School of Neuroscience, Tel-Aviv University, 69989 Ramat-Aviv, Tel-Aviv, Israel

<sup>5</sup>These authors contributed equally

<sup>6</sup>Senior author

<sup>7</sup>Lead Contact

\*Correspondence: [bharathchandra.talluri@gmail.com](mailto:bharathchandra.talluri@gmail.com) (B.C.T.), [t.donner@uke.de](mailto:t.donner@uke.de) (T.H.D.)

<https://doi.org/10.1016/j.cub.2018.07.052>

## SUMMARY

People's assessments of the state of the world often deviate systematically from the information available to them [1]. Such biases can originate from people's own decisions: committing to a categorical proposition, or a course of action, biases subsequent judgment and decision-making. This phenomenon, called confirmation bias [2], has been explained as suppression of post-decisional dissonance [3, 4]. Here, we provide insights into the underlying mechanism. It is commonly held that decisions result from the accumulation of samples of evidence informing about the state of the world [5–8]. We hypothesized that choices bias the accumulation process by selectively altering the weighting (gain) of subsequent evidence, akin to selective attention. We developed a novel psychophysical task to test this idea. Participants viewed two successive random dot motion stimuli and made two motion-direction judgments: a categorical discrimination after the first stimulus and a continuous estimation of the overall direction across both stimuli after the second stimulus. Participants' sensitivity for the second stimulus was selectively enhanced when that stimulus was consistent with the initial choice (compared to both, first stimuli and choice-inconsistent second stimuli). A model entailing choice-dependent selective gain modulation explained this effect better than several alternative mechanisms. Choice-dependent gain modulation was also established in another task entailing averaging of numerical values instead of motion directions. We conclude that intermittent choices direct selective attention during the evaluation of subsequent evidence, possibly due to decision-related feedback in the brain [9]. Our results point to a recurrent interplay between decision-making and selective attention.

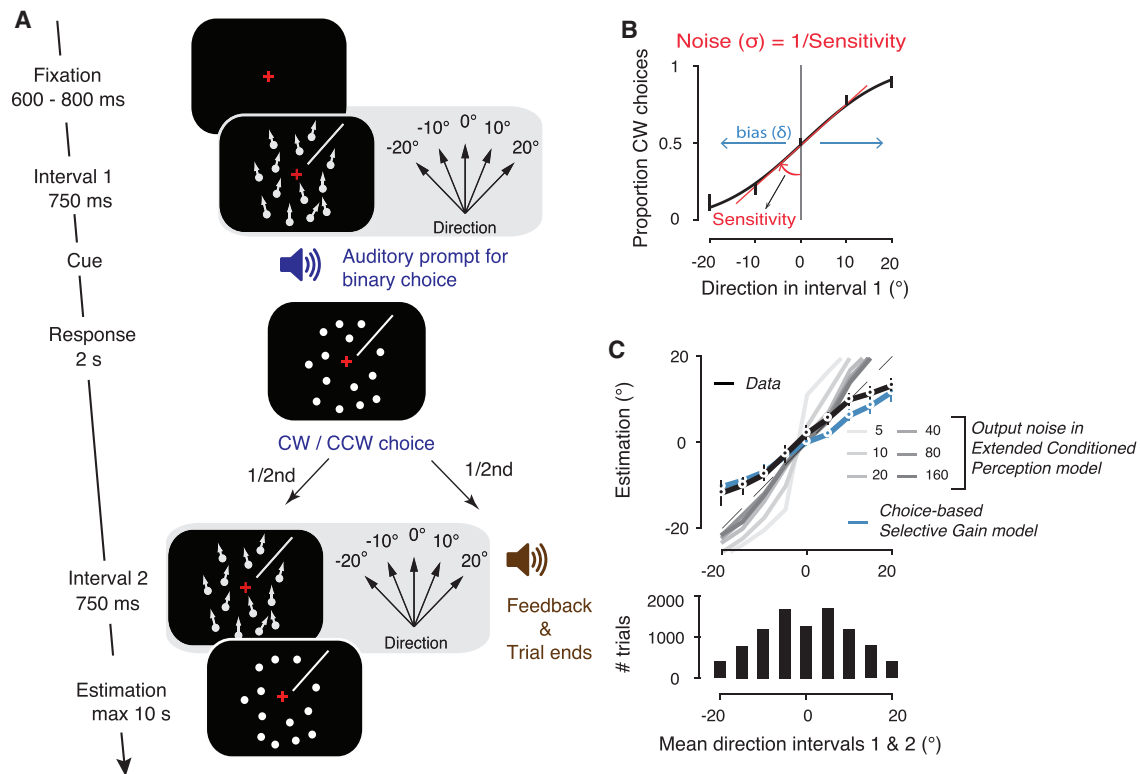
## RESULTS

Brain regions implicated in evidence accumulation, decision-making, and attentional control maintain their activity states over long timescales and send feedback to regions encoding the incoming evidence [9–11]. We thus reasoned that the consistency of new evidence with a previous choice might affect the decision-maker's sensitivity to the new evidence. Specifically, we hypothesized that a categorical choice induces a multiplicative gain modulation of new evidence, selectively boosting the sensitivity to consistent evidence. Such a selective gain modulation is commonly observed when explicit cues direct feature-based attention [12–14].

Previous studies have identified gain modulations in evidence accumulation by presenting multiple samples of evidence in succession and asking participants to report a binary choice based on the mean evidence at the end of the sequence [15–18]. Those studies did not assess the effect of intermittent choices in biasing the accumulation process. Other work has probed the interaction between categorical choices and continuous estimations by combining discrimination and estimation judgments based on the same evidence presented before [19–22]. Here, choice-related estimation biases may be a by-product of the bottom-up sensory decoding (i.e., weighting of sensory neurons) being tailored to the discrimination judgment [19] (but see [20, 22]). Whether a categorical choice occurring during a protracted stream of decision-relevant evidence selectively modulates the gain of evidence subsequent to that choice has remained unknown. We addressed this question by combining the above two approaches.

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





**Figure 1. Perceptual Task with Discrimination and Estimation Judgments**

(A) Schematic sequence of events within a trial. A first dot motion stimulus was shown on all trials for 750 ms and then paused. On two-thirds of trials, an auditory prompt instructed a direction discrimination judgment (CW or CCW with respect to reference line, at  $45^{\circ}$  in this example trial) as shown here. A third of trials, not analyzed here, did not require a choice. After half of the discrimination judgments, feedback was given, and the trial terminated. After the other half, a second motion stimulus was presented (equal coherence as first but independent direction), and participants were asked to estimate the mean direction of both stimuli.

(B) Proportion of CW choices as a function of stimulus direction, along with psychometric function fit.

(C) Top: Continuous estimations as function of mean direction across both stimuli. Bottom: Distribution of mean directions across trials. Black, data; blue, predictions generated from best-fitting parameters of Choice-based Selective Gain model; data points, group mean; error bars, SEM; gray, predictions by Extended Conditioned Perception model under several levels of output noise for average subject. Stimulus directions and estimations were always expressed as the angular distance from the reference, the position of which varied from trial to trial ( $0^{\circ}$  equals reference).

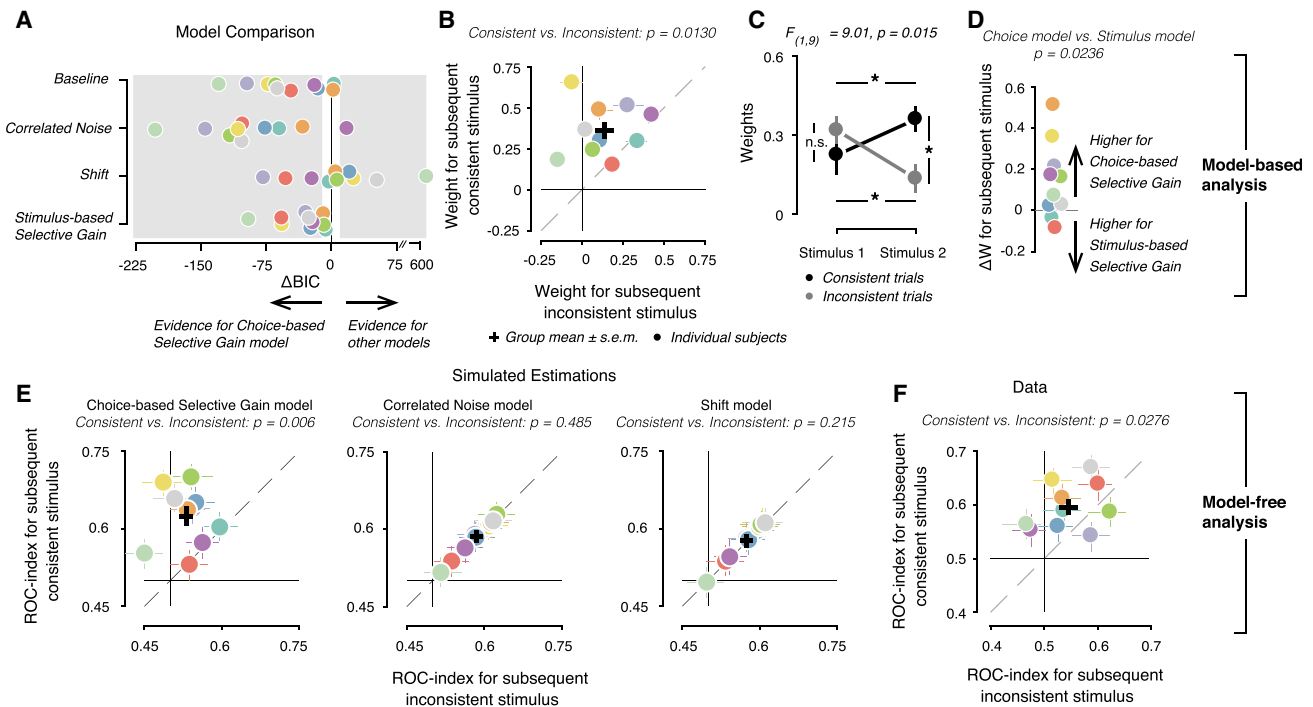
See also STAR Methods, Figure S1, and Video S1.

Participants made use of the stimulus information for both judgments: the fraction of CW choices increased as a function of the direction of the first stimulus from the reference (Figures 1B and S1A), and continuous estimations scaled with the mean stimulus direction across both intervals (Figures 1C, top, S1A, and S1B). The estimations were generally attracted toward the reference (Figure 1C, top, compare black and gray dashed line), in line with the non-uniform distribution of the mean stimulus directions (Figure 1C, bottom; STAR Methods).

Post-decisional selective gain modulation predicts that evidence subsequent to a choice produces larger (smaller) deviations in the overall estimations when these new directions are consistent (inconsistent) with that choice. We used two complementary approaches to test this prediction. The first approach modeled the overall estimations as a noisy weighted average of the directional evidence in both stimulus intervals (see STAR Methods). The weight for each stimulus quantified its gain in the estimation process. Trials with second stimulus directions consistent or inconsistent with the choice were modeled separately. This model, referred to as the Choice-based Selective

Gain model in the following (STAR Methods), provided a good account of observers' estimation reports (Figures 1C and 2). Smaller values of Bayes information criterion (BIC) within the majority of individual participants indicated that Choice-based Selective Gain explained the data better than a Baseline model without choice-dependent change in evidence weighting (Figure 2A; STAR Methods). Further, choice-consistent second stimuli received larger weight than choice-inconsistent second stimuli (Figure 2B; see Figure S2A for noise estimates). This weight difference was not evident for the first stimuli (Figure 2C). Indeed, weights were increased compared to the first stimulus for choice-consistent second stimuli and reduced for choice-inconsistent stimuli (Figure 2C). In sum, observers prioritized choice-consistent evidence after the categorical choice, in a way resembling feature-based attention.

The second, complementary approach corroborated this conclusion (Figures 2E and 2F). We developed a model-free measure based on the receiver-operating characteristic (ROC) that quantified the sensitivity to the second stimulus. ROC indices measured the extent to which single-trial estimations



**Figure 2. Sensitivity to Second Stimulus Dependent on Consistency with Initial Choice**

(A) Comparison between Choice-based Selective Gain and alternative models. Negative values are evidence for Choice-based Selective Gain. Gray,  $|\Delta BIC| > 10$ , indicating very strong evidence for model with smaller BIC (STAR Methods).

(B) Model weights for second stimulus in Consistent and Inconsistent conditions. Error bars, 66% bootstrap confidence intervals; black cross, mean and SEM; dashed line, Consistent = Inconsistent; data points above dashed line, Consistent > Inconsistent.

(C) Mean model weights for both stimulus intervals in Consistent and Inconsistent. Error bars, SEM; F-statistic, interaction between interval and condition (2-way ANOVA).

(D) Difference between effect strength (difference: Consistent-Inconsistent) for second stimulus, in weights obtained from Choice-based Selective Gain and Stimulus-based Selective Gain models.

(E) ROC indices for second Consistent and Inconsistent stimulus, predicted by simulations of alternative models as indicated above (individual trial distributions and best fitting model parameters).

(F) As (D) but for measured data. Data points in all but (C) are participants, with identical color scheme. p values, permutation tests (100,000 permutations) comparing weights or ROC indices between Consistent and Inconsistent across participants ( $n = 10$ ).

See also Figure S2.

separated between second stimuli of nearby directions (i.e.,  $10^\circ$  versus  $20^\circ$ , or  $-10^\circ$  versus  $-20^\circ$ ; see STAR Methods for details). Simulations confirmed that the difference between these ROC indices, computed separately for choice-consistent and choice-inconsistent stimuli, captured the choice-dependent gain modulation described by the Choice-based Selective Gain model (Figures 2E, left, and S2B). Critically, for the actual data, ROC indices were larger for the Consistent than Inconsistent condition (Figure 2F). In sum, the model-free analysis also revealed a selective modulation of sensitivity to additional evidence, in line with feature-based attention.

This consistency-dependent change in sensitivity for subsequent evidence, as quantified by the ROC indices, could not be explained by other mechanisms lacking multiplicative gain modulation. In a first alternative model, biases shared among choice and subsequent estimations resulted from slow fluctuations in noise corrupting both judgments, without any genuine effect of the choice. This so-called Correlated Noise model (STAR Methods) provided a worse account of estimation reports (in 9 out of 10) than Choice-based Selective Gain (Figure 2A) and

could not produce the consistency-dependent ROC effect neither for the individually fitted parameters (Figure 2E, middle) nor for any combination of parameters that we simulated (Figure S2B).

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

data (Figure S2B). It is possible that an additive shift and multiplicative gain modulation jointly governed choice-induced biases in the overall estimation behavior (see Discussion).

Taken together, the analyses presented so far indicate that consistency-dependent gain modulation was necessary to account for certain features of participants' behavior. Further analyses indicated that this gain modulation was, in fact, induced by the intermittent choice (i.e., participants' categorization of the first stimulus) rather than by the first stimulus itself (Figures S2C and S2D) or by the disparity between first and second stimulus (Figure S2E). We fitted a variant of the Selective Gain model, in which the consistency of the second stimulus was defined based on the first physical stimulus direction, rather the participants' choice (STAR Methods). This so-called Stimulus-based Selective Gain model provided a worse account of the data than the Choice-based Selective Gain (Figure 2A). Critically, the selective gain effect was larger for the parameters estimated by Choice-based Selective Gain model (Figure 2D). In sum, the selective modulation in sensitivity was linked to the participants' categorical choice.

A recent Bayesian account of post-decision biases has proposed that perceptual inference is "conditioned" on choice in order to ensure consistency between binary discrimination and continuous estimation judgments of the same stimulus [20, 22]. This account is framed at a different level of description (Bayesian inference), but the notion of a choice-dependent prior for estimation is similar to our idea of a choice-induced top-down modulation. Could choice-based conditioning of internal representations explain the present results? Our task and analyses isolated the impact of binary choice on the processing of subsequent evidence for continuous estimation, requiring additional assumptions about the conditioning operation. If only the representation of the first stimulus was conditioned, this would yield an offset of the representation of the second stimulus—equivalent to the Shift model considered above, which did not account consistency-effect on ROC indices observed in the data (Figures 2E, right, and S2B). If also the representation of the second stimulus was conditioned on the choice (referred to as Extended Conditioned Perception, see STAR Methods), this reproduced the ROC-effect (Figure S2B, right). However, the later model did not account well for the relationship between overall estimations and mean stimulus direction (gray lines in Figure 1C; for further comparison between Extended Conditioned Perception and Choice-based Selective Gain, see Figures S2G and S2H). Future work should develop biologically plausible and dynamic approximations of choice-based conditioning operation in order to unravel possible links to choice-dependent gain modulation.

The post-decisional biasing effect in the visual perceptual task resembled well-documented effects in reasoning [2] and preference reports [4, 23]. It is unknown, however, whether the latter high-level post-decision biases are mediated by selective gain modulations akin to attention. To test for this, we re-analyzed and modeled previously published [24] data from a numerical averaging task that also required the combination of evidence presented before and after a choice into an overall estimation (Figure 3A; see STAR Methods for task and analysis details). Again, the weights were larger on Consistent than Inconsistent conditions, specifically for evidence after choice (Figure 3B) again with an interaction between interval and consistency

(Figure S3D). Likewise, the ROC indices were also larger for Consistent than Inconsistent conditions (Figure 3C). In sum, the choice-induced biasing mechanism we uncovered for perceptual decision-making, including the selective gain modulation, also accounts for post-decision biases in higher-level decisions based on numerical evidence.

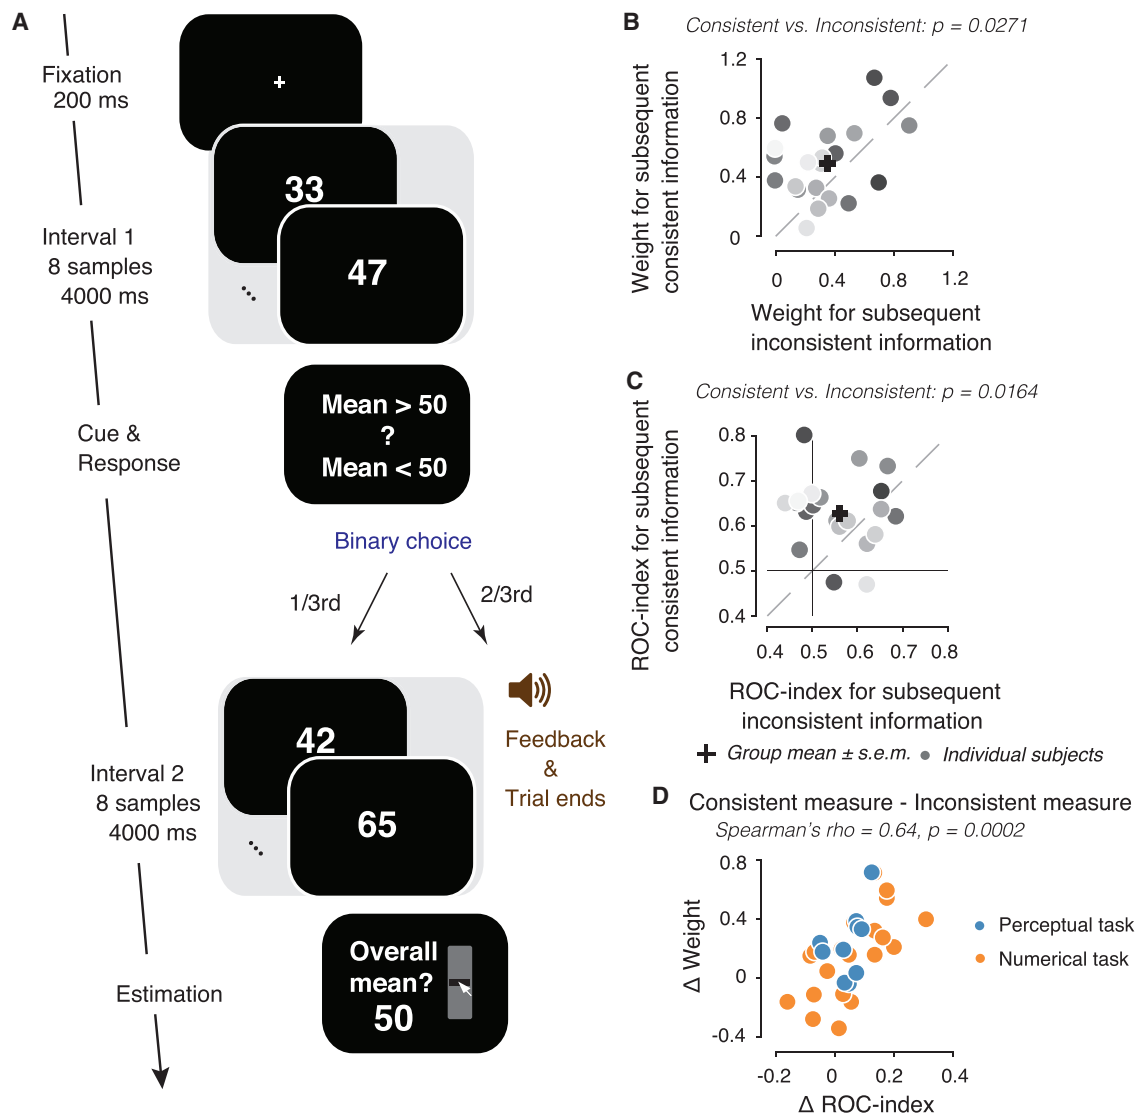
## DISCUSSION

Decision-makers are often systematically influenced by their own choices: committing to a categorical hypothesis or choosing a course of action biases the subsequent evaluation of the decision-relevant evidence [2, 4]. The mechanisms underlying such post-decisional confirmation biases have so far remained unknown. Here, we have shown that choices selectively increased the gain of subsequent evidence that was consistent with that choice, for perceptual as well as numerical decisions. A selective modulation of the gain of sensory responses is commonly observed during attention to certain stimulus features [12–14]. In sum, our results illuminate the linkage between decision-making and attention—two capacities commonly studied in isolation but interacting in real-life behavior. Our findings indicate that an agent's decision acts like a cue for selective attention, biasing subsequent decision processing.

Evidence inconsistent with an initial choice may induce post-decisional dissonance, possibly related to conflict between competing cognitive states or motor responses [3, 25]. Previous work has shown that such conflict boosts top-down control, increasing task performance and response caution on subsequent trials [26, 27]. But this line of work has not associated conflict with subsequent decision biases. In particular, it has not shown that conflict induces selective modulations of new information that is consistent with respect to a previous choice.

We have recently established that sensitivity for new information is generally reduced after an overt choice, compared to no overt choice [24]. To this end, we assessed a non-selective reduction in sensitivity for *any* post-decision evidence. Our current work goes beyond this by uncovering a *selective* mechanism of confirmation bias: preferentially sampling the evidence that confirms one's prior belief. This effect indicates a more refined mechanism than the non-selective reduction in overall sensitivity due to an overt choice. Identifying this effect was afforded by an improved modeling approach (see STAR Methods) combined with a model-free behavioral readout of selective gain modulation (ROC analysis), both yielding consistent results (Figure 3D). It is conceivable that a non-selective gain reduction due to overt choice (possibly reflecting reduced arousal and/or cortical attractor dynamics [24]) and selective attention toward choice-consistent evidence conspire to shape overall estimation behavior.

Our analysis of the perceptual task also revealed, in some of the participants, an additive shift in the direction of the chosen category, on top of the gain modulation. This additive shift may reflect previously identified choice-induced biases [19]. This additive shift could not, however, account for the consistency-dependent change in sensitivity (Figure 2E), which we found in the data (Figure 2F). The co-existence of additive and multiplicative effects may relate to the observation that common



**Figure 3. Numerical Task and Sensitivity for Choice-Consistent versus Choice-Inconsistent Information**

(A) Schematic sequence of events within a trial entailing intermittent binary choice. After the first sequence of eight numbers, participants discriminated the mean as larger or smaller than 50 (a quarter of trials, not analyzed here, did not require a choice; see STAR Methods and [24]). Following the discrimination report, the trial terminated with feedback (two-thirds of trials), or a second sequence of eight numbers was presented (mean independent from that of first interval). Participants were then asked to report the mean of the whole number sequence.

(B and C) Model-based (B) and model-free (C) measures of consistency-dependent sensitivity modulation (as Figures 2B and 2F).

(D) Correlation between consistency effect in model-based and model-free analyses across participants from both tasks. Effect strength is Consistent-Inconsistent difference in model weights or ROC. Data points, participants.  $p$  values in (B) and (C) from permutation tests across participants (100,000 permutations;  $n = 21$ ). See also Figure S3.

manipulations of selective attention produce effects on both sensitivity and decision criteria, which are dissociable at behavioral and neural levels [28, 29]. Our present experimental manipulation does not allow for distinguishing between an additive baseline shift in the sensory response and a shift of the starting point of the decision variable accumulating the sensory response. Future experiments could manipulate the duration of the second evidence to dissociate these two scenarios.

Our work contributes to recent progress in the understanding of history-dependent biases in perceptual choice [30–40]. One

class of mechanism contributing to such biases is stimulus-selective adaptation, which can cause repulsion [36] or attraction [41–44], possibly owing to adaptation dynamics at different processing levels. Low-level adaptation after prolonged stimulus exposure as in our task (Figure 1A) commonly produces repulsive effects [36], due to suppressing sensory cortical responses [45, 46]. This is inconsistent with our results because it predicts stronger sensitivity loss for congruent than incongruent stimuli. Higher-level adaptation mechanisms can cause attraction, especially in the face of ambiguous evidence [41–43], and has



been linked to gain modulation induced by the stimulus sequence [44]. We found that the consistency-dependent gain modulation was more strongly tied to observers' choices than the physical stimuli, implying a higher-level source. This, combined with the multiplicative nature of the effect, naturally links it to feature-based attention. Whether a local adaptation mechanism with such functional characteristics exists remains to be tested.

History biases in perceptual choice tasks requiring categorical judgments have specifically been linked to the history of previous choices [32, 36, 37, 39, 47] or choice outcomes [33, 48]. While these across-trial biases are idiosyncratic [33, 37], the predominant tendency is to repeat choices more often than expected by chance [32, 36, 37, 39], in line with the current choice-consistency bias established here within a single protracted decision process. Recent work on across-trial history biases in categorical choice points to a similar attentional mechanism giving rise to choice-repetition biases across trials [49].

The accumulation of fluctuating sensory evidence toward binary choices is well characterized at a neurophysiological level [6, 8, 50, 51]. Theoretical work points to an analogous mechanism underlying continuous decisions [7]. While less is known about continuous decisions based on two successive evidence streams, it is tempting to speculate that the selective re-weighting effect results from top-down feedback from cortical accumulator regions to regions that encode the evidence [9–11, 52]. Such feedback interactions might alter the decision-maker's interpretation of incoming information by the evolving belief state [53, 54].

Our results have broader implications. First, insight into the computational mechanisms producing confirmation biases has considerable ecological value because these biases are pervasive in daily life, shaping human judgment in cases of critical significance (e.g., scientific hypothesis testing) [2]. Second, our work sets the stage for probing into the neural mechanisms of confirmation biases in humans and animal models. Previous work into confirmation bias has focused on high-level judgment and reasoning [2], the neural bases of which remain elusive. By contrast, neuroscience has accumulated substantial knowledge about the neural signals that encode the sensory evidence and evolving decision about visual motion [8, 55]. The modulation of visual motion signals by attention is also well characterized [12]. Our current findings establish an analogous biasing mechanism in both domains—high-level judgment and perceptual decisions—along with an effective behavioral readout and computational signature that constrains for the candidate neural mechanisms.

## STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Perceptual task
  - Numerical task

- Modeling discrimination judgments
- Modeling estimation reports
- Model-free analysis of estimation reports
- Simulated estimations from the models
- QUANTIFICATION AND STATISTICAL ANALYSIS
- DATA AND SOFTWARE AVAILABILITY

## SUPPLEMENTAL INFORMATION

Supplemental Information includes three figures and one video and can be found with this article online at <https://doi.org/10.1016/j.cub.2018.07.052>.

## ACKNOWLEDGMENTS

We thank Ana Vojvodic for help with data collection and Zohar Bronfman and Noam Brezis for continued discussion and helpful feedback on the manuscript. This research was supported by grants from the German Research Foundation (DO 1240/2-1, DO 1240/3-1, and SFB 936/A7 to T.H.D.), the German Academic Exchange Service (to A.E.U.), and a Marie Curie Individual Fellowship (to K.T.). We acknowledge computing resources provided by NWO Physical Sciences.

## AUTHOR CONTRIBUTIONS

T.H.D., K.T., and M.U. formulated the idea for the study and designed the experiment. A.E.U., T.H.D., and M.U. specified the design for the perceptual tasks. A.E.U. programmed the task and collected data for the perceptual task. B.C.T., A.E.U., and K.T. analyzed data. T.H.D., M.U., and K.T. supervised the study. B.C.T., A.E.U., and T.H.D. wrote the paper. All authors revised and approved the manuscript.

## DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: February 9, 2018

Revised: June 6, 2018

Accepted: July 19, 2018

Published: September 13, 2018

## REFERENCES

1. Tversky, A., and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 1124–1131.
2. Nickerson, R.S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* 2, 175–220.
3. Festinger, L. (1957). *A theory of cognitive dissonance* (Stanford: Stanford University Press).
4. Brehm, J.W. (1956). Postdecision changes in the desirability of alternatives. *J. Abnorm. Psychol.* 52, 384–389.
5. Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.
6. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (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.
7. Liu, F., and Wang, X.-J. (2008). A common cortical circuit mechanism for perceptual categorical discrimination and veridical judgment. *PLoS Comput. Biol.* 4, e1000253.
8. Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron* 60, 215–234.
9. Wimmer, K., Compte, A., Roxin, A., Peixoto, D., Renart, A., and de la Rocha, J. (2015). Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat. Commun.* 6, 6177.

10. Nienborg, H., and Cumming, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459, 89–92.
11. Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. *Science* 348, 1352–1355.
12. Maunsell, J.H.R., and Treue, S. (2006). Feature-based attention in visual cortex. *Trends Neurosci.* 29, 317–322.
13. Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. *Neuron* 61, 168–185.
14. Herrmann, K., Heeger, D.J., and Carrasco, M. (2012). Feature-based attention enhances performance by increasing response gain. *Vision Res.* 74, 10–20.
15. Tsetsos, K., Chater, N., and Usher, M. (2012). Saliency driven value integration explains decision biases and preference reversal. *Proc. Natl. Acad. Sci. USA* 109, 9659–9664.
16. Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., and Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proc. Natl. Acad. Sci. USA* 113, 3102–3107.
17. Drugowitsch, J., Wyart, V., Devauchelle, A.-D., and Kochlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron* 92, 1398–1411.
18. Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* 76, 847–858.
19. Jazayeri, M., and Movshon, J.A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* 446, 912–915.
20. Stocker, A.A., and Simoncelli, E.P. (2007). A Bayesian Model of Conditioned Perception. *Adv. Neural Inf. Process. Syst.* 2007, 1409–1416.
21. Zamboni, E., Ledgeway, T., McGraw, P.V., and Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. Biol. Sci.* 283, 20160263.
22. Luu, L., and Stocker, A.A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife* 7, e33334.
23. Chen, M.K., and Risen, J.L. (2010). How choice affects and reflects preferences: revisiting the free-choice paradigm. *J. Pers. Soc. Psychol.* 99, 573–594.
24. Bronfman, Z.Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., and Usher, M. (2015). Decisions reduce sensitivity to subsequent information. *Proc. Biol. Sci.* 282, 20150228.
25. van Veen, V., Schooler, J.W., and Carter, C.S. (2009). Neural activity predicts attitude change in cognitive dissonance. *Nat. Neurosci.* 12, 1469–1474.
26. Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., and Cohen, J.D. (2001). Conflict monitoring and cognitive control. *Psychol. Rev.* 108, 624–652.
27. Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
28. Luo, T.Z., and Maunsell, J.H.R. (2015). Neuronal Modulations in Visual Cortex Are Associated with Only One of Multiple Components of Attention. *Neuron* 86, 1182–1188.
29. Luo, T.Z., and Maunsell, J.H.R. (2018). Attentional Changes in Either Criterion or Sensitivity Are Associated with Robust Modulations in Lateral Prefrontal Cortex. *Neuron* 97, 1382–1393.e7.
30. Fernberger, S.W. (1920). Interdependence of judgments within the series for the method of constant stimuli. *J. Exp. Psychol.* 3, 126–150.
31. de Lange, F.P., Rahnev, D.A., Donner, T.H., and Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *J. Neurosci.* 33, 1400–1410.
32. Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron* 81, 195–206.
33. Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. USA* 113, E3548–E3557.
34. Pape, A.-A., and Siegel, M. (2016). Motor cortex activity predicts response alternation during sensorimotor decisions. *Nat. Commun.* 7, 13098.
35. Kim, T.D., Kabir, M., and Gold, J.I. (2017). Coupled Decision Processes Update and Maintain Saccadic Priors in a Dynamic Environment. *J. Neurosci.* 37, 3632–3645.
36. Fritsche, M., Mostert, P., and de Lange, F.P. (2017). Opposite Effects of Recent History on Perception and Decision. *Curr. Biol.* 27, 590–595.
37. Urai, A.E., Braun, A., and Donner, T.H. (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* 8, 14637.
38. Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *J. Vis.* 14, 9.
39. Braun, A., Urai, A.E., and Donner, T.H. (2018). Adaptive History Biases Result from Confidence-weighted Accumulation of Past Choices. *J. Neurosci.* 38, 2189–17.
40. Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* 554, 368–372.
41. Kanai, R., and Verstraten, F.A.J. (2005). Perceptual manifestations of fast neural plasticity: motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Res.* 45, 3109–3116.
42. Brascamp, J.W., Knapen, T.H.J., Kanai, R., Noest, A.J., van Ee, R., and van den Berg, A.V. (2008). Multi-timescale perceptual history resolves visual ambiguity. *PLoS ONE* 3, e1497.
43. Pearson, J., and Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends Cogn. Sci.* 12, 334–341.
44. Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Hecce Castañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron* 81, 1429–1441.
45. Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. *J. Neurophysiol.* 97, 3155–3164.
46. Krekelberg, B., Boynton, G.M., and van Wezel, R.J.A. (2006). Adaptation: from single cells to BOLD signals. *Trends Neurosci.* 29, 250–256.
47. Bonaiuto, J.J., Berker, A., and Bestmann, S. (2016). Response repetition biases in human perceptual decisions are explained by activity decay in competitive attractor models. *eLife* 5, e20047.
48. Busse, L., Ayaz, A., Dhruv, N.T., Katzner, S., Saleem, A.B., Schölvinc, M.L., Zaharia, A.D., and Carandini, M. (2011). The detection of visual contrast in the behaving mouse. *J. Neurosci.* 31, 11351–11361.
49. Urai, A.E., de Gee, J.W., and Donner, T.H. (2018). Choice history biases subsequent evidence accumulation. *bioRxiv*. <https://doi.org/10.1101/251595>.
50. Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Curr. Biol.* 23, 981–986.
51. Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340, 95–98.
52. Goris, R.L.T., Ziemba, C.M., Stine, G.M., Simoncelli, E.P., and Movshon, J.A. (2017). Dissociation of Choice Formation and Choice-Related Activity in Macaque Visual Cortex. *J. Neurosci.* 37, 5195–5203.
53. Nienborg, H., and Roelfsema, P.R. (2015). Belief states as a framework to explain extra-retinal influences in visual cortex. *Curr. Opin. Neurobiol.* 32, 45–52.
54. Haefner, R.M., Berkes, P., and Fiser, J. (2016). Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron* 90, 649–660.
55. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. *Neuron* 80, 791–806.



56. Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., and Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception* 36, 1–16.
57. Scase, M.O., Braddick, O.J., and Raymond, J.E. (1996). What is noise for the motion system? *Vision Res.* 36, 2579–2586.
58. Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.
59. Wichmann, F.A., and Hill, N.J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept. Psychophys.* 63, 1293–1313.
60. Jazayeri, M., and Movshon, J.A. (2006). Optimal representation of sensory information by neural populations. *Nat. Neurosci.* 9, 690–696.
61. Bogacz, R., and Cohen, J.D. (2004). Parameterization of connectionist models. *Behav. Res. Methods Instrum. Comput.* 36, 732–741.
62. Efron, B., and Tibshirani, R. (1986). Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. *Stat. Sci.* 1, 54–75.
63. Schwarz, G. (1978). Estimating the Dimension of a Model. *Ann. Stat.* 6, 461–464.
64. Kass, R.E., and Raftery, A.E. (1995). Bayes Factors. *J. Am. Stat. Assoc.* 90, 773–795.
65. Green, D.M., and Swets, J.A. (1966). *Signal Detection Theory and Psychophysics* (John Wiley and Sons).

## STAR★METHODS

### KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited Data		
Behavioral data: Perceptual task	This paper	<a href="https://doi.org/10.6084/m9.figshare.7048430">https://doi.org/10.6084/m9.figshare.7048430</a>
Behavioral data: Numerical task	[24]	<a href="https://datadryad.org/resource/doi:10.5061/dryad.40f6v">https://datadryad.org/resource/doi:10.5061/dryad.40f6v</a>
Software and Algorithms		
MATLAB	Mathworks	Matlab_R2015a, Matlab_R2016b
Custom Code (experiment, models, analyses)	This paper	<a href="https://github.com/BharathTalluri/postchoicebias">https://github.com/BharathTalluri/postchoicebias</a>
Psychtoolbox	[56]	<a href="http://psychtoolbox.org/">http://psychtoolbox.org/</a>

### CONTACT FOR REAGENT AND RESOURCE SHARING

All resources, including data and code used for the analyses on this paper, are publicly available (see Data and Software Availability). Further information and requests for resource sharing should be directed to and will be fulfilled by the Lead Contact, Tobias H. Donner ([t.donner@uke.de](mailto:t.donner@uke.de)).

### EXPERIMENTAL MODEL AND SUBJECT DETAILS

Data from sixteen participants (six men and ten women) between the ages of 18 and 29 were collected for this study. Two participants did not complete the full experiment and were discarded from all analyses. The estimations of some subjects did not increase monotonically, quantified by the slope of best-fitting line, as a function of mean direction (red boxes in Figure S1A). We excluded four participants for whom the slopes were  $< 0.3$  (Figure S1B), and the results in Figures 1 and 2 are based on the remaining 10 participants. All gave written informed consent prior to participation, and were naive to the aim of the experiment. The University of Amsterdam ethics review board approved the project. Each participant performed a total of 12 sessions, distributed across six days: One session to determine the motion coherence of the stimuli that corresponded to the individual psychophysical threshold and 11 sessions of the main experimental task. Each session of the main task consisted of 345 trials, divided into five experimental blocks of 69 trials. We used the first two sessions (690 trials) as training sessions to get participants acquainted to the task. We also re-analyzed previously collected data [24] from an additional 21 participants (age range: 21 to 29). In this data, after a short block of 20 practice trials, each participant completed 300 trials (5 blocks of 60 trials each).

### METHOD DETAILS

#### Perceptual task

##### Stimuli

Stimuli were presented using PsychToolbox-3 [56] in MATLAB and were viewed in a dark, quiet room on a CRT monitor with a resolution of 1024 pixels x 768 pixels and a refresh rate of 60 Hz. Participants placed their heads on a chinrest with a viewing distance of 50 cm from the screen. Dynamical random dot stimuli were presented in a central circle (outer radius  $12^\circ$ , inner radius  $2^\circ$ ) around fixation. A field of dots with a density of 1.7 dots/degrees<sup>2</sup> defined the annulus. Dots were  $0.2^\circ$  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  $11.5^\circ$ /second in the signal direction. Signal dots that left the annulus wrapped around and reappeared on the other side. Moreover, signal dots had a limited “lifetime,” and were re-plotted in a random location after being on the screen for four consecutive frames. Noise dots were assigned a random location within the annulus on each frame, resulting in ‘random position’ (white) noise with a ‘different’ rule [57]. Additionally, to avoid participants tracking individual signal dots as they move through the annulus, three independent motion sequences were interleaved on subsequent frames [58], making the effective speed of dots  $3.8^\circ$ /second.

##### Procedure: Determining individual motion coherence thresholds

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

$$P(\text{correct} | c) = \delta + (1 - \delta - \gamma) \left( 1 - e^{(-c/\alpha)^\beta} \right) \quad (\text{Equation 1})$$

where  $\delta$  was the guess rate (chance performance),  $\gamma$  was the lapse rate, and  $\alpha$  and  $\beta$  were the threshold and slope of the psychometric Weibull function, respectively. While keeping the guess rate  $\delta$  fixed at 50% correct, we fit the parameters  $\gamma$ ,  $\alpha$  and  $\beta$  maximizing the likelihood function [59] using a Nelder-Mead simplex optimization algorithm. The individual threshold was taken as the stimulus difficulty corresponding to an 80% correct fit of the cumulative Weibull. Across participants, motion coherence thresholds ranged from 11% to 28% (mean 18%).

### Procedure: Main experiment

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

For each participant, the reference position was constrained to be either within the top (0°–180°) or the bottom half (180°–360°) of the stimulus unit circle (balanced across participants) in order to keep the mapping between CW/CCW choices and left/right button presses constant within each participant. There were five possible directions (–20°, –10°, 0°, 10°, 20°) of each dot motion stimulus, yielding 25 possible combinations of directions across both subsequent stimuli. Of those, only 23 were used, excluding the two most obviously conflicting combinations (–20°/20° and 20°/–20°). The resulting distribution of mean directions was non-uniform and bimodal (Figure 1C). Feedback about their estimation performance was given at the end of each block as the mean deviation across trials of their estimation reports from the physical stimulus directions. A video demonstration of the task can be found in Video S1 accompanying the paper.

In total, 90 trials for each combination of first and second directions were presented per participant (45 in Choice and 45 in No-Choice trials). Trials were excluded from analysis according to the following criteria: (i) participants did not comply with the instructions (i.e., pressing the mouse wheel on Choice trials or a choice key on No-Choice trials); (ii) binary choice reaction time was below 200 ms (i.e., shorter than regular reaction times on two-choice tasks); and (iii) estimation outliers (defined as estimations beyond 1.5 times the interquartile range, above upper or below lower quartile). In total ~7% of the total trials across the 10 participants were excluded. In addition, we excluded all No-Choice trials from our analyses as we focus only on Choice trials here. The distributions of the remaining trials used for analysis are shown in Figures S1C and S1D.

### Task instructions

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

1. Every block consists of several trials of the same visual motion task. Always keep your gaze on the red fixation point in the center of the screen.
2. Blank screen: Each trial will begin with a black screen. The red fixation appearance indicates that the trial is about to start.
3. Interval 1: You will see a cloud of dots moving, with some of the dots moving together in a particular direction. Your task is to determine whether the dots are moving to the *left* or to the *right* of the reference mark.
4. Binary Response: Once the dots stop moving, you will hear an auditory prompt to report your decision about the direction by clicking the corresponding mouse key (left or right). Try your best to make this decision as quickly and accurately as possible.
5. After your response, trials will continue with either:
  - a) Feedback: Once you've pressed a mouse key, you will hear feedback about your response in some trials. A correct choice will be followed by a high beep, and an incorrect choice will be followed by a low beep. Following feedback, you will move on to the next trial, or
  - b) Interval 2: After your response, you will see a second cloud of dots moving, with some of the dots moving together in a particular direction. These dots may have a different angle of motion from the first stimulus. Your task is to determine and estimate what the *average* overall angle of motion is from this cloud and the first one combined.
6. Estimation Response: When the dots stop moving and the reference mark will turn red, you must complete an estimation task. Move the mouse to align the cursor to the *average* angle of motion you saw in the two trials. Once you are satisfied with your estimate, click the mouse to confirm your response.
7. When you see a blank screen, the trial is over and you will have the opportunity for a break. Rest your eyes for a moment to help you keep them open and fixated during the experimental trials. When you want to continue the experiment, click the mouse to continue.

8. After each block of around 10 min, you will see a screen indicating your performance on the last block and telling you to take a break.

### Numerical task

The task was identical to the perceptual task described above, with the following exceptions. Participants viewed two sequences of eight two-digit numbers each and reported the mean of all 16 samples as a continuous measure. Each sequence lasted for 4 s, and each numerical sample was displayed for 500 ms. In 75% of all trials, prompted by a visual cue, subjects made a binary choice about the mean of the first sequence of numbers- whether the mean was greater than or less than 50 by pressing the corresponding key. In a proportion of these trials (25% of all trials), the binary response was followed by a second sequence of eight numbers after which subjects made the estimation judgment by vertically sliding a mouse-controlled bar set on a number ruler between 0 and 100. Numbers were generated from pre-defined distributions ranged between 10 and 90. The data we analyzed in this paper constituted 25% of trials from each subject (~75 trials). Please see the original report on this dataset ([24]) for a more detailed description of the task, and <https://datadryad.org/resource/doi:10.5061/dryad.40f6v> for the behavioral data.

### Modeling discrimination judgments

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

$$P(\text{Choice} = \text{CW}) = \Phi(\delta + \alpha\phi_1) \quad (\text{Equation 2})$$

where  $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt$  was the cumulative Gaussian function,  $\alpha$  was the slope of the psychometric function (i.e., perceptual sensitivity), and  $\delta$  was the horizontal shift of the psychometric function (i.e., systematic bias toward one of the two choice options; see Figure 1B). The inverse of the parameter  $\alpha$  quantifies the internal sensory noise. The free parameters  $\alpha$  and  $\delta$  were estimated by maximum likelihood optimization [59].

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

### Modeling estimation reports

We used a statistical modeling approach to estimate the relative contributions of the evidence conveyed by both successive dot motion stimuli, or number streams, to participants' trial-by-trial estimation reports. All models described in this section were fit exclusively to the Choice trials. A comparison of Choice and No-choice trials was beyond the scope of this study; it can be found in [24] for the numerical task data and will be subject to a subsequent report for the perceptual task data.

#### Baseline model

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

$$\mathbf{y} = w_1 \mathbf{X}_1 + w_2 \mathbf{X}_2 + \mathcal{N}(0, \xi) \quad (\text{Equation 3})$$

where  $\mathbf{y}$  was the vector of single-trial estimations expressed as angular deviation from the reference mark,  $\mathbf{X}_1$  and  $\mathbf{X}_2$  were the noisy representation of stimulus direction 1 and 2 respectively (see below),  $w_1$  and  $w_2$  were the weights assigned to the corresponding evidence, and  $\mathcal{N}(0, \xi)$  was zero-mean Gaussian estimation noise with variance  $\xi$ . Because estimations  $\mathbf{y}$  were expressed relative to the reference, so were the internal representations  $\mathbf{X}_i$ . We used this format of internal representation and estimation reports because (i) in our design, the cursor movements used for estimation report were always initiated at the reference and (ii) recent work on post-decision biases has highlighted the importance of the reference [21]. The reference-dependent format of internal representation assumed for  $\mathbf{X}_i$  thus did not describe the sensory representation of motion direction, but rather a statistic extracted from that sensory representation in a task-dependent fashion (e.g., the direction with the largest posterior probability relative to reference [60]).

Here and below,  $\mathbf{X}_1$  and  $\mathbf{X}_2$  were computed by replacing the angular deviation of the physical stimulus from the reference,  $\phi$ , with:

$$X_i = \phi_i + \mathcal{N}(\delta, \sigma) \quad (\text{Equation 4})$$

where  $i \in (1, 2)$ , and  $\delta$  and  $\sigma$  were each observer's individual overall bias and sensory noise parameters taken from the psychometric function fit to the binary choice data (see Equation 2 above).  $\mathbf{X}_1$  and  $\mathbf{X}_2$  thus approximated the noisy internal representation that governed observers' estimations. Specifying these two parameters in this fashion avoided adding additional free parameters to Equation 3. The approach was based on the assumption that a substantial portion of biases was shared between the choice and

estimation judgments. We validated this assumption by confirming that biases were strongly correlated between binary choices (quantified as the horizontal shift of psychometric function estimations) and estimation reports (mean angular estimation error), across the 10 observers (Spearman's  $\rho = 0.79$ ;  $p = 0.0098$ ). In sum, the Baseline model had three free parameters,  $w_1$ ,  $w_2$  and  $\xi$ .

### Correlated Noise model

This model assumed shared noise in the internal representations  $\mathbf{X}_1$  and  $\mathbf{X}_2$ . Specifically, the model assumed that additive noise corrupting the transformation of physical stimulus directions  $\phi_i$  into  $\mathbf{X}_i$  was correlated, to some degree, across the two intervals, inducing correlations between  $\mathbf{X}_1$  and  $\mathbf{X}_2$ , as follows:

$$\mathbf{X}_1 = \phi_1 + \delta + \varepsilon \quad (\text{Equation 5.1})$$

$$\mathbf{X}_2 = \phi_2 + \delta + (1 - c) \cdot \mathcal{N}(0, \sigma) + c\varepsilon \quad (\text{Equation 5.2})$$

where  $\varepsilon$  was the noise in  $\mathbf{X}_1$  in a given trial, drawn from the distribution  $\mathcal{N}(0, \sigma)$ , and  $c \in [0, 1]$ . Parameter  $c$  governed the degree of correlation among the two internal representations. Thus, the noise in  $\mathbf{X}_2$  was made up of a portion correlated with the noise in  $\mathbf{X}_1$  (i.e.,  $c\varepsilon$ ) and another portion independent of  $\mathbf{X}_1$  (i.e.,  $(1 - c) \cdot \mathcal{N}(0, \sigma)$ ). The estimations were modeled by Equation 3 above.

### Shift model

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

$$\mathbf{y} = w_1 \mathbf{X}_1 + w_2 \mathbf{X}_2 + k \mathbf{D} + \mathcal{N}(0, \xi) \quad (\text{Equation 6})$$

where  $k$  was the additive shift parameter, and  $\mathbf{D}$  was the vector of intermediate binary choices, taking the values  $[1, -1]$ .

### Selective Gain models

The Selective Gain model enabled testing for a selective change in sensitivity to Consistent versus Inconsistent evidence conveyed by the direction of the second stimulus  $\phi_2$ . Consistency of that direction could be defined with respect to the initial choice or with respect to the first stimulus direction  $\phi_1$ . This led to two alternative versions of the Selective Gain model, specified next.

**Choice-based Selective Gain.** This model was as the Baseline model, except that the weights were allowed to vary depending on whether  $\phi_2$  was consistent or inconsistent with the initial choice:

$$\mathbf{y} = w_{1c} \mathbf{X}_1 + w_{2c} \mathbf{X}_2 + \mathcal{N}(0, \xi_c) \text{ if } \text{sign}(\phi_2) = \mathbf{D} \quad (\text{Equation 7.1})$$

$$\mathbf{y} = w_{1i} \mathbf{X}_1 + w_{2i} \mathbf{X}_2 + \mathcal{N}(0, \xi_i) \text{ if } \text{sign}(\phi_2) \neq \mathbf{D} \quad (\text{Equation 7.2})$$

where  $w_{1c}$  ( $w_{2c}$ ) and  $w_{1i}$  ( $w_{2i}$ ) were the weights for Consistent and Inconsistent trials, respectively,  $\phi_1$  and  $\phi_2$  were the physical stimulus directions from both intervals, and  $\mathbf{D}$  was the vector of intermediate binary choice (values: 1 or  $-1$  for CCW and CW reports, respectively). Since consistency could not be defined in trials where  $\phi_2$  was  $0^\circ$ , we excluded this subset of trials before fitting Equations 7.1 and 7.2.

**Extended Choice-based Selective Gain.** We also tested whether there was an additive shift, over and above the multiplicative gain modulation described by the Choice-based Selective Gain model. To this end, we extended the model from Equations 7.1 and 7.2 by means of the shift parameter from Equation 6, as follows:

$$\mathbf{y} = w_{1c} \mathbf{X}_1 + w_{2c} \mathbf{X}_2 + k \cdot \mathbf{D} + \mathcal{N}(0, \xi_c) \text{ if } \text{sign}(\phi_2) = \mathbf{D} \quad (\text{Equation 8.1})$$

$$\mathbf{y} = w_{1i} \mathbf{X}_1 + w_{2i} \mathbf{X}_2 + k \cdot \mathbf{D} + \mathcal{N}(0, \xi_i) \text{ if } \text{sign}(\phi_2) \neq \mathbf{D} \quad (\text{Equation 8.2})$$

This Extended Choice-based Selective Gain model was fit by constraining all parameters to take the values estimated by the basic version of the model (i.e., Equations 7.1 and 7.2), with the shift as the only free parameter. Parameter recovery indicated that leaving all parameter free to vary in the fit made the model too complex given the limited amount of data (see *Parameter Recovery* below).

**Stimulus-based Selective Gain.** This version of the model was as the previous one, except that consistency depended on the direction of the first stimulus (specifically: the sign of its difference from the reference), not the initial choice:

$$\mathbf{y} = w_{1c} \mathbf{X}_1 + w_{2c} \mathbf{X}_2 + \mathcal{N}(0, \xi_c) \text{ if } \text{sign}(\phi_2) = \text{sign}(\phi_1) \quad (\text{Equation 9.1})$$

$$\mathbf{y} = w_{1i} \mathbf{X}_1 + w_{2i} \mathbf{X}_2 + \mathcal{N}(0, \xi_i) \text{ if } \text{sign}(\phi_2) \neq \text{sign}(\phi_1) \quad (\text{Equation 9.2})$$

Since consistency could not be defined in trials where  $\phi_1$  or  $\phi_2$  was  $0^\circ$ , we excluded this subset of trials before fitting Equations 9.1 and 9.2).



**Choice-based Selective Gain model after matching evidence disparity.** As control for the differences in disparity between motion directions in first and second interval, we randomly subsampled trials such that the absolute distance between  $\phi_1$  and  $\phi_2$  in Consistent and Inconsistent trials is matched. This was done before fitting the Choice-based selective gain model described above.

**Choice-based Selective Gain model for numerical task.** The data from the numerical task (numerical integration, Figure 3A) were also fit with the Choice-based Selective Gain model, but with the following differences due the nature of the task design and the smaller number of trials per individual than available for the perceptual task. As in the perceptual task, the mean evidence also exhibited a small bias after splitting by choice-consistency. Different from the perceptual task, the group average estimations exhibited a small opposite trend, i.e., there was an interaction between mean evidence versus estimations and (Consistent, Inconsistent) condition. The model as in Equation 7 could not capture this interaction, and indeed we found that fitting the model without accounting for it yielded poor fits. (Please note that the model-free analysis of sensitivity described below was unaffected by this issue.) To account for this, we introduced two additional free parameters  $\theta_c$  and  $\theta_i$ , as follows:

$$\mathbf{y} = w_{1c}\mathbf{X}_1 + w_{2c}\mathbf{X}_2 + \mathcal{N}(\theta_c, \xi_c) \text{ if } \text{sign}(\phi_2) = \mathbf{D} \quad (\text{Equation 10.1})$$

$$\mathbf{y} = w_{1i}\mathbf{X}_1 + w_{2i}\mathbf{X}_2 + \mathcal{N}(\theta_i, \xi_i) \text{ if } \text{sign}(\phi_2) \neq \mathbf{D} \quad (\text{Equation 10.2})$$

where  $\theta_c$  and  $\theta_i$  accounted for the above interaction,  $\phi_i$  was the mean of 8 samples (again relative to the reference, i.e., 50) in each interval ( $i = 1, 2$ ),  $\mathbf{X}_i$  was the noise-corrupted and biased internal representation of the mean value in each interval computed as in Equation 4. For the results shown in Figure 3, we constrained all weights to be positive based on the assumption that weights should not be negative. We also constrained the possible values of  $\theta_c$  and  $\theta_i$  to be within a reasonable range,  $[-10, 10]$ , still far larger than the magnitude of the interaction observed in the group average estimation data. The above constraints were introduced in order to obtain reliable model fits in the face of limited data (trials). Results were qualitatively similar (especially, significant difference between Consistent and Inconsistent weights for the second interval) when fitting the model without those constraints.

Further, as an additional control, we also fitted the model without constraints and without  $\theta$  (i.e., Equation 7), after randomly sampling trials from Consistent and Inconsistent conditions, so as to minimize the above interaction through matching the mean evidence between both conditions. Because this procedure substantially reduced the number of trials (~23%), we only fitted the model on the remaining data after pooling trials across all participants. We repeated this ‘mean-matching’ procedure 500 times and re-fitted the model for each random trial selection. The median across iterations of the difference in weights for Consistent versus Inconsistent was 0.065, with a 95% confidence range that excluded zero (0.00001, 0.133). In sum, also this second approach for fitting the data from the numerical task supported the re-allocation of the weights for the second evidence dependent on choice-consistency as observed in the first model-based approach (Figure 3B) as well as in the model-free analysis (Figure 3C).

### Likelihood computation

We used maximum likelihood estimates to estimate parameters and the goodness of fit of different models. For any unique combination of experimental variables (first and second stimuli, and choice), we numerically derived the estimation distribution of each model for a given parameter set and used this estimation distribution to assess the likelihood of the estimation reported by participants on a given trial with the corresponding experimental variables. All models described above assume that the stimuli on each trial are represented in the form of scalar values. Thus, the estimation distribution represents a distribution of estimations over trials. Specifically, the estimation distribution was the expected distribution of estimations for a given set of experimental variables, if the model was simulated several times using the same set of parameters. Using this numerical method avoided the need to rely on large number of stochastic simulations in order to compute the likelihoods and made the fitting procedure less prone to converging to local minima. For all models except for the Correlated Noise model (see below), we numerically derived each model’s estimation distribution for each experimental condition, by first generating Gaussian distributions centered at  $w_1X_1$  and  $w_2X_2$  with standard deviation  $|w_1| \sigma$  and  $|w_2| \sigma$  for intervals 1 and 2 respectively. Then, we set the probability in the non-chosen side to zero in the interval 1 distribution (i.e., we truncated the distribution to only have density in the chosen side) and normalized it so as the integral of the resulting distribution is equal to 1. We combined the probability distributions corresponding to stimulus 1 (truncated distribution) and 2 using convolution and renormalized the resulting distribution. Note that different weights applied to stimulus 1 and stimulus 2 distributions (see Equations 3, 6–10) in different models. We then generated a zero-mean Gaussian probability distribution with variance  $\xi$  and convolved this distribution with the distribution from the previous step, renormalizing the resulting distribution to obtain the estimation probability distribution for that trial. We used this probability distribution to calculate the likelihood of the reported estimation in the trial. Finally, we summed the logarithm of likelihood values over all trials to obtain the final log-likelihood value for a given set of parameters.

For the Correlated Noise Model, we used Monte Carlo techniques to simulate a probability distribution of estimations over trials. For each combination of experimental variables (combination of first and second stimuli, and choice), we generated a set of 10,000 normally distributed noisy representations ( $\mathcal{N}(0, \sigma)$ ) or noisy samples for interval 1 ( $X_1$ , Equation 5.1). From these noisy samples, we discarded those where the sign of  $X_1$  did not match the binary choice of the subject. These maintained samples featured as variable  $\varepsilon$  in interval 2 (Equation 5.2). Another set of noisy samples was generated afresh for interval 2 ( $\mathcal{N}(0, \sigma)$  in Equation 5.2). Note that the number of new noisy samples and thus of the simulated representations  $X_2$  was less than 10,000 because of the sub-selection described above. We combined  $X_1$  and  $X_2$  using Equation 3, to obtain a distribution of estimations for this trial. Smoothing kernels

were obtained from the simulated estimations in each trial in order to identify the underlying distributions, which were then used to calculate the likelihood. The kernels were defined using non-parametric *Epanechnikov* function. Finally, we summed the logarithm of likelihood values of all trials to obtain the final log-likelihood value for a given set of parameters ( $w_1$ ,  $w_2$ ,  $\xi$  and  $c$ ). In some trials, the likelihood of the estimations was zero regardless of the values of the parameters (possibly because these estimations were motor lapses), resulting in an optimization function that never converged. To address this, we added one simulated estimation trial in the response range ( $-180^\circ$  to  $180^\circ$  in steps of 1) to the distribution of estimations obtained from Equation 3, before obtaining the estimation kernels. This did not influence the maximum likelihood fitting procedure in other trials, but just gave an estimation kernel with non-zero probability value for the whole range of estimations.

#### Comparison of fitting procedure to Bronfman et al., 2015

In our previous report on analyses of the data from the numerical task [24], we had also fitted a so-called Selective Gain model to the data from the numerical estimation task and compared that to a model without such selective gain modulation. Here, we applied a model fitting procedure that differed from the previous one in two important respects. First, in the previous study, models were fitted to the across participants aggregated data (i.e., a ‘fixed effects’ approach) [24], while we here fitted models to each participant’s data individually. The second difference concerns the calculation of the likelihood in [24]: In the previous study, model estimations were derived for each trial (using 1000 simulated trials) and then the likelihood of the reported estimation was computed under the simplifying assumption that the model’s estimation distribution is Gaussian. This was done in order to avoid kernel-based smoothing of the simulated estimations, which could significantly slow down the fitting procedures. In the current study, however, we did not make this assumption since in all models the predicted distribution, albeit symmetric looking under most parameter sets, always had non-zero skew. We thus computed likelihoods from the actual non-Gaussian distributions that the models predict.

#### Fitting procedure and computation of confidence intervals

To obtain the best fitting parameters that maximize the likelihood function of each model, we used Subplex algorithm [61], a generalization of the Nelder-Mead simplex method, which is well suited to optimize high dimensional noisy objective functions. Subplex starts at a specified starting point of the objective function and works by dividing the parameter space into subspaces. It then performs a simplex search in each of these subspaces before converging on the set of parameters that maximize the function. The starting points were randomly chosen from the interval [0,20] for  $\xi$  and [0 1] for  $w_1$  and  $w_2$ .

We used bootstrapping [62] to obtain confidence intervals for the fitted parameters for each individual. Specifically, we randomly selected trials with replacement and fit the selective gain model to these resampled datasets. We repeated this procedure 500 times, each time using Subplex optimization with starting points at the best-fitting parameters of the actual data. We then obtained confidence intervals from the distribution of estimated parameters.

#### Parameter recovery

We simulated data with different sets of parameters using the number of trials as in a typical dataset. We then fit the simulated data using the Choice-based Selective gain model (Equations 7.1 and 7.2). Overall, the model recovered parameters well: The Spearman correlations between actual and recovered parameters was 0.8 for Noise parameters ( $p < 10^{-10}$ ), and it ranged between 0.91 to 0.94 for all weight parameters ( $p < 10^{-10}$ ). Importantly, the model also did not introduce any spurious correlations between the recovered parameters. Inter-parameter correlations for the actual parameters ranged between  $-0.03$  to  $0.03$  ( $p > 0.39$ ), and between  $-0.05$  to  $0.03$  ( $p > 0.23$ ) for the recovered parameters. This allowed us to confirm that our fitting procedures were able to recover the parameters when the ground truth of the data was known.

Simulations of the most complex model assessed here, the Extended Choice-based Selective Gain (Equations 8.1 and 8.2) with all parameters left free to vary, also showed decent overall recovery of parameters (Spearman correlations ranged between 0.78 to 0.94,  $p < 10^{-10}$ ). However, for special cases, the fits exhibited significant spurious correlations between parameter estimates. Specifically, we introduced a few iterations where the actual shift parameter ( $k_{actual}$ ) was 0, or the consistency parameter for weights of interval 2 ( $\Delta w_{2, actual} = w_{2c, actual} - w_{2i, actual}$ ) was 0. When  $k_{actual}$  was 0, the mean recovered shift ( $k_{recovered}$ ) in these iterations was  $-1.63$  ( $p = 0.0043$ ). This spurious shift was introduced by the model at the expense of  $\Delta w_2$  i.e., the correlation between  $k_{recovered}$  and  $\Delta w_{2, actual} - \Delta w_{2, recovered}$  was  $-0.85$  ( $p < 10^{-5}$ ). Likewise, in iterations where  $\Delta w_{2, actual} = 0$ , mean  $\Delta w_{2, recovered} = 0.115$  ( $p = 0.046$ ). This spurious consistency effect was introduced by the model at the expense of the shift parameter i.e., the correlation between  $\Delta w_{2, recovered}$  and  $k_{actual} - k_{recovered}$  was  $-0.65$  ( $p < 10^{-3}$ ). Because these spurious correlations rendered fits of this complex model generally hard to interpret, we did not report any parameter estimates from this model.

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

#### Model comparison

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

$$BIC = -2 \cdot \ln(\mathcal{L}) + m \cdot \ln(n) \quad (\text{Equation 11})$$

where  $\mathcal{L}$  was the likelihood value,  $m$  was the number of free parameters in the model and  $n$  was the number of observations that are used to fit the model [63]. 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 [64]. Since BIC values depended on the number of observations used to fit the model, we fit all models under comparison on the same subset of trials to enable us identify the model that best explains the data. We calculated BIC values for all individual model fits to identify the model that better explained the data for that participant.

### Model-free analysis of estimation reports

We also assessed the impact of the second stimuli to participants' estimations in a model-independent fashion. The rationale was to quantify the impact of small differences in the evidence values (stimulus directions or numerical means) on the estimation reports produced by the participant, depending on whether the two directions were consistent or inconsistent with the previous choice. Our analysis aimed to compare the separability of distributions of single-trial estimations from subsets of trials, between Consistent and Inconsistent conditions. We quantified the separability of estimation distributions by means of the receiver operating characteristic (ROC) from signal-detection theory [65]. The area under the ROC-curve, referred to as ROC-index, could range from 0 to 1. An index of 0.5 implied perfectly overlapping distributions (i.e., no sensitivity to the  $10^\circ$  difference) and any deviation from 0.5 implies some sensitivity to the evidence. An ROC-index of 1 (or 0) implied that the two distributions were completely separable.

We intended to use the ROC measure for specifically quantifying the sensitivity to the smallest presented difference ( $10^\circ$ ) in the direction of the second stimulus ( $\phi_2$ ), while eliminating the impact of the direction of the first stimulus ( $\phi_1$ ) on the final estimation. To this end, we used the following procedure for the perceptual task. All trials (except for those with  $\phi_2 = 0^\circ$  where choice-consistency was not defined) were first sorted by whether the direction of second interval was consistent or inconsistent with the initial choice. For each thus-defined condition (Consistent, Inconsistent), we further sorted trials by  $\phi_1$  (i.e.,  $-20^\circ$ ,  $-10^\circ$ ,  $0^\circ$ ,  $10^\circ$ ,  $20^\circ$ ). For each  $\phi_1$  we compared estimation distributions from trials with  $\phi_2$  differing by  $10^\circ$  (i.e.,  $-20^\circ$  versus  $-10^\circ$  and  $10^\circ$  versus  $20^\circ$ ). The resulting ROC-indices were then pooled across the different  $\phi_1$  directions, so as to yield a single pooled ROC index, separately for Consistent and Inconsistent conditions. We pooled the ROC indices by means of weighted averaging, whereby the weight of each ROC index was determined by the number of trials that went into the calculation of that ROC-index. That number differed substantially between ROC indices due to the uneven distribution of pairs of directions of the first and second stimulus (Figure S1D). The resulting ROC-indices were compared between both conditions by means of permutation tests (see next section). We obtained qualitatively identical results when simply discarding the trial pairs with small trial numbers ( $< 15$ ) and averaging the other ROC-indices without weighting (mean difference in ROC-index between Consistent and Inconsistent trials across subjects = 0.04).

ROC indices for the numerical task (numerical integration, Figure 3C) were computed as for the perceptual task, with the following exceptions. We binarized the mean of the stimulus presented in the first interval into two bins ( $\phi_1 > 50^\circ$  and  $\phi_1 < 50^\circ$ ), and split the mean of the second stimulus into four bins with means at 40, 47, 53 and 60, two each on either side of the reference number 50. Those four bins were treated equivalently to the different  $\phi_2$  values in the description for the perceptual task above.

### Simulated estimations from the models

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

To simulate estimations in each trial, we first calculated the internal representations  $X_1$  and  $X_2$  using Equation 4 (Equation 5 for the Correlated Noise model). In addition, we ensured that the sign of  $X_1$  matches the binary decision made by the subject in the trial. We then combined these internal representations to obtain a simulated estimation for the trial, using the corresponding parameters and equations for each model.

#### Simulated estimations from the best fitting parameters

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

#### Simulated estimations using a range of parameters

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

#### Simulations of Extended Conditioned Perception model

For simulations, we extended the Conditioned Perception model described in [20, 22] for discrimination and estimation judgments on a single stimulus to our task with two successive stimuli, intermittent choice, and a total estimation judgment at the end. We simulated a version of this 'Extended Conditioned Perception' model, in which the posterior distribution over stimulus directions after both stimuli were conditioned on the intermittent choice. The resulting procedure was as described in the section *Likelihood computation*

for Choice-based Selective Gain model above, with the following differences. In the Conditioned Perception model, contrary to all other models presented, external stimuli on each trial are not represented by scalar variables but as posterior probability distributions over stimulus features (i.e., directions), given the sensory stimulus and the choice. Before conditioning, the mean of this posterior distribution for each stimulus was given by each stimulus' direction and its standard deviation was given by the individual psychometric noise (parameter  $\sigma$  from Equation 4, so-called 'input noise' capturing imprecise encoding of direction). Per trial two such distributions were obtained, one per stimulus. Both of these distributions were then conditioned on the choice: the probability was set to zero on all stimulus directions that were inconsistent with the choice, for the first and the second stimulus. The two resulting distributions were finally combined with equal weight, producing an overall estimation distribution for each trial. We extracted the mean from the resulting distribution as the model's estimate of direction. In different runs of the simulations we added different amounts of independent Gaussian (zero-mean) noise to these estimates ('output noise' capturing both imperfect memory of stimulus identity as well as motor noise). The resulting value was taken as the estimation report on a given trial. We simulated this model in order to assess if it would produce similar behavioral features as Choice-based Selective Gain. Please note that a more elaborate version of the Conditioned Perception model has been used to fit estimation data in [22].

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

In further simulations of this model, we used an average subject (pooling the trial distributions across all subjects), with input noise as mean parameter  $\sigma$  across all participants, and varied only the output noise. We computed the mean estimations as a function of average stimulus direction predicted by the model for several levels of output noise (Figure 1C). We performed an additional analysis to uncover subtler differences in the behavior of Extended Conditioned Perception and Choice-based Selective Gain (Figures S2G and S2H). The Extended Conditioned Perception model was essentially insensitive to new evidence inconsistent with the choice. Thus, we reasoned that the fraction of inconsistent estimations (i.e., estimations falling on the side of the reference opposite from the choice) predicted by this model should be lower than the fraction predicted by the Choice-based Selective Gain model. Specifically, the increase in this fraction as a function of inconsistent second stimulus should be higher for the Choice-based Selective Gain model. To test this prediction, we simulated estimations for the fixed-effects subject using the Conditioned Perception model for different levels of output noise, and the Choice-based Selective Gain model with the mean of the best fitting parameters across subjects. We then calculated the fraction of inconsistent estimations for correct and error trials, separately for positive and negative direction of the first stimulus ( $X_1 = -10, 10$ ), as well as for ambiguous first direction ( $X_1 = 0$ ), in the simulated estimations. We repeated this procedure for 100 iterations, and compared the mean fraction across the iterations in both models to that of the behavioral data (Figures S2G and S2H).

## QUANTIFICATION AND STATISTICAL ANALYSIS

Non-parametric permutation tests [62] were used to test for group-level significance of individual measures, unless otherwise specified. This was done by randomly switching the labels of individual observations either between two paired sets of values, or between one set of values and zero. After repeating this procedure 100,000 times, we computed the difference between the two group means on each permutation and obtained the p value as the fraction of permutations that exceeded the observed difference between the means. All p values reported were computed using two-sided tests.

## DATA AND SOFTWARE AVAILABILITY

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