Why Decisions Bias Perception: An Amortised Sequential Sampling Account

Jian-Qiao Zhu (J.Zhu@warwick.ac.uk)
Warwick Business School, University of Warwick
Adam N. Sanborn (A.N.Sanborn@warwick.ac.uk)
Department of Psychology, University of Warwick
Nick Chater (Nick.Chater@wbs.ac.uk)
Warwick Business School, University of Warwick
Coventry, UK CV4 7AL

Abstract

The judgments that people make are not independent – initial decisions can bias later perception. This has been shown in tasks in which participants first decide whether the direction of moving dots is to one side or the other of a reference line: their subsequent estimates are biased away from this reference line. This interesting bias has been explained in past work as either a consequence of weighting sensory neurons, or as a consequence of participants adjusting their estimate to match their decision. We propose a new explanation: that people sequentially sample evidence to make their decision, and reuse these samples to make their estimate (i.e., amortised inference). Because optimal stopping leads to samples that strongly favor one or another decision alternative, the subsequent estimates are also biased away from the reference line. We introduce a sequential sampling model for posterior samples that does not assume constant thresholds, and provide evidence for our explanation in a new experiment that generalizes the perceptual bias to a new domain.

Keywords: decision biases, adaptive sampling, amortised inference.

Introduction

Experiments in motion perception show that making a perceptual decision biases subsequent perception. As illustrated in Figure 1A, participants in these random-dot-motion experiments are first asked whether the motion was clockwise (CW) or counter-clockwise (CCW) of a decision boundary. After making this decision, participants are then asked to estimate the direction of motion. While participants' estimates are unsurprisingly consistent with their decision, these estimates also show a surprising perceptual bias: estimates are biased *away from* the decision boundary (Jazayeri & Movshon, 2007; Luu & Stocker, 2018; Zamboni, Ledgeway, McGraw, & Schluppeck, 2016).

Two main theories have been proposed to explain this perceptual bias: (a) the optimal weighting of outputs of orientation-tuned neurons used in the decision task is also used in the estimation task (Jazayeri & Movshon, 2007), or (b) people employ self-consistent reasoning by only considering hypotheses consistent with their initial decision when making an estimate (Luu & Stocker, 2018). Existing comparisons favor the self-consistency account (Luu & Stocker, 2018; Zamboni et al., 2016).

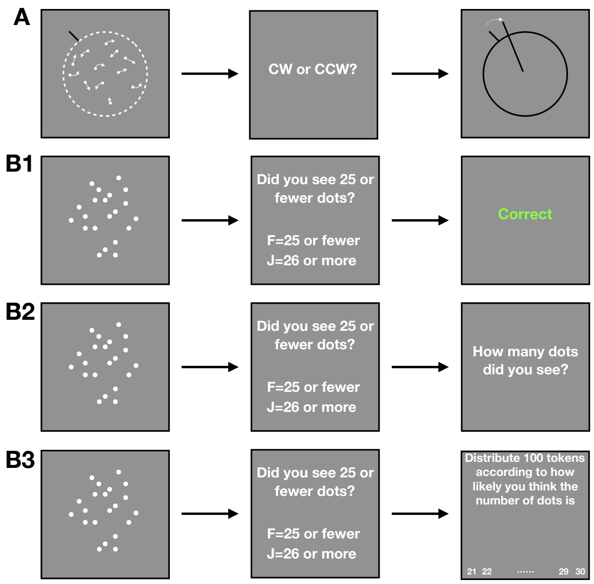

Figure 1: Illustration of experimental tasks. **(A)** Decision-estimation (D-E) task for random-dot-motion. **(B)** Numerosity experiment. **(B1)** Decision with feedback (D-F) task. **(B2)** D-E task. **(B3)** Decision-histogram (D-H) task.

However, it may be that self-consistency is unnecessary to explain the perceptual bias. We take a sequential sampling approach to modelling this task, following a long history of models in human decision making that sequentially draw perceptual or posterior samples and optimally accumulate evidence (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Vul, Goodman, Griffiths, & Tenenbaum, 2014). We introduce a sequential sampling model that optimally accumulates posterior samples, and then demonstrate that the perceptual bias in estimation is produced by simply averaging samples that were optimally accumulated for the initial decision. The intuition for why the bias is produced is simple: because sequential sampling models stop when the samples favor one of the alternatives, the estimate (i.e., the average) also favors one of the alternative. Interestingly, similar Bayesian analysis have arisen in a different experimental domain: the studies of probability estimation from sequential samples, where optimal stopping has shown to predict a distorting effects on subsequent judgments (Coenen & Gureckis, 2016). However, when the sampling process is external, such distortion did not receive empirical supports (Coenen & Gureckis, 2016).

To discriminate between these accounts, we first show that weighted decoding, self-consistency, and sequential sampling make qualitatively different predictions about the perceptual beliefs that people will have about individual stimuli. We then test these predictions in a new experiment which generalizes the perceptual bias from a simple random-dot-motion task to a perceptually more complex numerosity task (see Figure 1B), finding that the perceptual bias is best explained by sequential sampling.

Computational Models

In this section, we introduce and compare computational models of the perceptual bias.

Weighted Decoding

The Weighted Decoding (WD) model argues that post-decision bias is a result of optimally tuning the sensory representation to boost responses for the initial decision (Jazayeri & Movshon, 2007; Zamboni et al., 2016). For example, to discriminate whether a random-dot-motion stimulus is coherently moving clock-wise or counter-clockwise of a reference line, the neurons that respond maximally to motion directions that are slightly different from the reference line are the most informative. This leads to a optimal weighting profile that is bimodal: emphasizing directions that are slightly away from the reference line.

WD assumes this weighting profile is also used in the estimation task, and that the mode of the weighted sensory distribution is taken as the estimate. As a result, a post-decision bias naturally emerges (Figure 2A).

Self-Consistency

To predict the post-decision bias, WD must assume that the selective read-out of sensory information in the decision task is carried over to the subsequent estimation task. However, more recent work has demonstrated that the perceptual bias is actually a late decision-related bias, rather than a sensory bias (Luu & Stocker, 2018; Zamboni et al., 2016).

The Self-Consistency (SC) model is a Bayesian model that makes the initial decision according to which option has highest posterior probability, given noisy sensory evidence. However, because the estimate is made after the decision, this model assumes that the quality of the sensory evidence has decayed by the time participants are asked to make an estimate. Instead of relying on the low-quality sensory evidence alone, SC assumes that the participant treats their initial decision (which was made with high-quality sensory evidence) as information as well (cf. Fleming & Daw, 2017), only considering hypotheses that are consistent with the initial decision. SC's estimate is then the mean of the posterior distribution over hypotheses consistent with the initial choice. As shown in Figure $2B_1$, SC also produces estimates that are biased away from the decision boundary.

Our implementation of SC also predicts a bias toward the decision boundary for true stimuli that are far away from the boundary. This is for an uninteresting reason: in the task we will describe below the response range was restricted, and so we also truncated the posterior at the edges of the allowable response range – this leads presentations of extreme values to be biased toward the center of the range.

Simple Amortised Sampling

Because perfectly storing and representing probability distributions can easily become computationally daunting, sample-based approximations have been proposed as a way for the brain to approximate Bayesian inference (Sanborn & Chater, 2016; Zhu, Sanborn, & Chater, 2018). On each trial, the Simple Amortised Sampling (SAS) model generates a set of N samples from the posterior distribution:

$$x_i \overset{\text{i.i.d}}{\sim} P(X|S), \quad i = 1, 2, ..., N \tag{1}$$

where $P(X|S)$ is the posterior sensory representation of number of dots given the stimulus. For the decision task, SAS chooses the alternative that attracts the larger number of samples, which introduces a natural stochasticity into the decision.

In the later estimation task, it makes little sense to draw a new set of samples, as an average of the samples drawn to make the decision can serve as the estimate. This effort-saving strategy is a form of *amortised inference* (Gershman & Goodman, 2014). Reusing samples in this way ensures a high degree of consistency between the decision and the estimate SAS make. However, SAS will not produce a perceptual bias away from the decision boundary – the average of a fixed number of samples is unbiased (Figure $2C_3$), and it only shows a bias toward the center for extreme stimuli. Thus, this model is not actually a candidate for explaining the perceptual bias, but instead serves to illustrate why the following model does.

Bayesian Amortised Sequential Sampling

If samples are drawn sequentially and require effort to generate, it often makes no sense to continue sampling until a fixed number are obtained. Instead, it is more efficient to stop sampling when it is no longer worthwhile.

Many different kinds of sequential sampling models have been proposed, including those that accumulate sensory information (Bogacz et al., 2006), and those that accumulate the kind of posterior samples similar to SAS (Vul et al., 2014). We take as a starting point the sequential model introduced in Vul et al. (2014), which accumulates samples until there are a threshold T more in favor of one alternative than the other. This scheme has the advantages of producing a fixed probability of choosing the better alternative regardless of the number of samples, and it is possible to find the optimal threshold for maximizing utility.

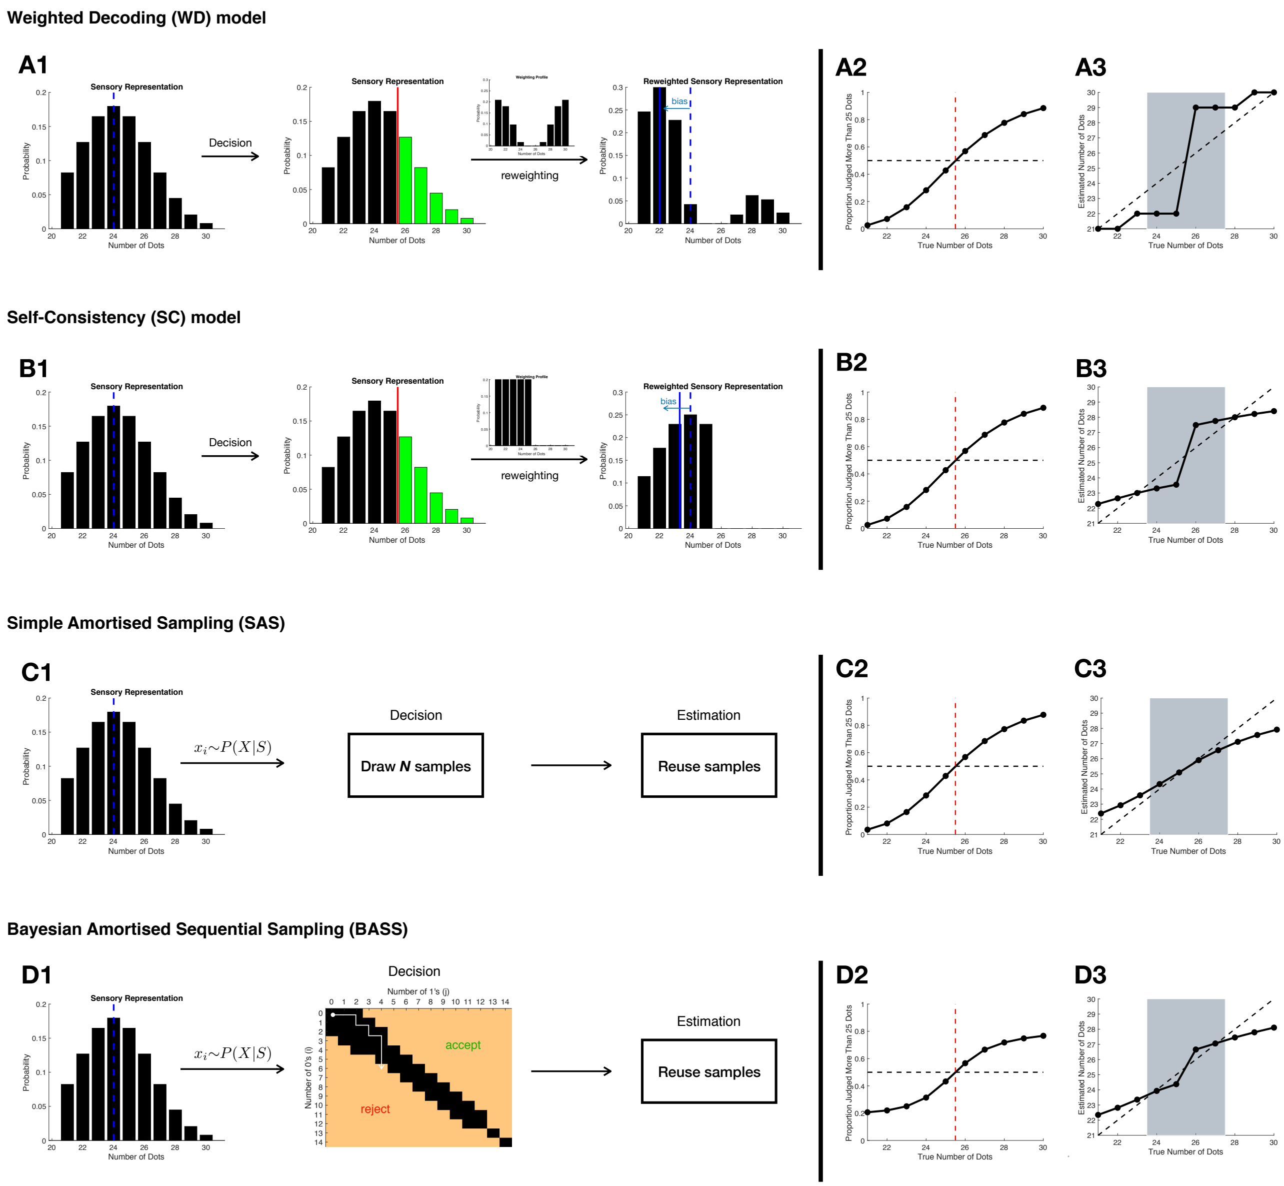

Figure 2: A comparison of model mechanisms and predictions for the numerosity task. For illustrative purposes, we assume a Gaussian likelihood of dot numbers. Then, the posterior distribution combines this likelihood with the prior (i.e., a uniform distribution from 21 to 30). A posterior distribution with its mean at 24 dots (dashed blue line) is shown for each model. **(A)** Schematic illustration of the WD model (Jazayeri & Movshon, 2007). The sensory representation is reweighted according to the optimal bimodal weighting profile. The initial decision is based on the relative probability above and below the decision boundary. The estimate is then the mode of this same reweighted sensory representation. **(B)** Schematic illustration of the SC model (Luu & Stocker, 2018). The initial decision is based on whether there is more probability above or below the decision boundary. The estimate is the mean of the portion of the posterior distribution that is consistent with the initial decision. **(C)** Schematic illustration of the SAS model. A fixed number of posterior samples are drawn and the alternative that attracts the larger number of samples is chosen in the initial decision. The average of these samples is the estimate. **(D)** Schematic illustration of the BASS model. Samples are sequentially drawn from the posterior distribution until it is no longer worthwhile to continue. The alternative that attracts the larger number of samples is chosen in the initial decision. The average of these samples is the estimate.

However, it is possible to sequentially draw samples from a posterior distribution even more efficiently by allowing a non-constant threshold and determining before every sample whether it is better to continue or stop sampling. We term this scheme Bayesian Amortised Sequential Sampling (BASS). The problem of finding for the optimal changing threshold was solved by Wald (1950) for deciding when to stop drawing binomial samples from an external source, and a similar approach to external samples was investigated empirically by Coenen and Gureckis (2016). We simply adopt Walds approach to optimally drawing internal posterior samples.

The posterior probability p that one decision alternative is true is assumed to be unknown, but we assume that binomial samples can be sequentially drawn with probability p. We perform Bayesian inference using the obtained samples, by first placing a prior distribution over p and assume a fixed cost c of drawing a sample, reflecting the time and effort of doing so. After drawing j samples in favor of a decision alternative and i against, we denote p_{ij} as the posterior probability of a decision alternative given those samples, with p_{00} being the prior probability. The binary decision task essentially becomes a sequential test on whether $p_{ij} < 1/2$ and the optimal stopping rule for this test can then be derived from the following using dynamic programming:

$$F(i,j) = \min \begin{cases} F_0(i,j), \\ c + p_{ij}F(i,j+1) + (1-p_{ij})F(i+1,j). \end{cases} \tag{2}$$

where $F(i,j)$ and $F_0(i,j)$ are respectively the expected cost of sampling and expected cost of termination after i samples against and j in favor have been observed. The sampling process should terminate whenever $F(i,j) \geq F_0(i,j)$.

Because $F_0(i,j)$ represents the expected cost of stopping the sampling process when the posterior probability of an alternative is p_{ij}, if the punishment for an incorrect decision is one unit of utility and thus,

$$F_0(i,j) = \min \begin{cases} i/(i+j), \\ j/(i+j). \end{cases} \tag{3}$$

This is the expected cost of incorrectly choosing an alternative when the posterior probability of that alternative is $p_{ij} = \text{Beta}(i,j)$.

The expected cost of drawing another sample is the sum of (a) the cost of generating one sample c, (b) the expected cost if the new sample turns out to be in favor $p_{ij}F(i,j+1)$, and (c) the expected cost if the new sample turns out to be against $(1-p_{ij})F(i+1,j)$.

While the exact solution of the Bayesian optimal stopping problem is difficult to obtain, once computed it can also be reused across different cognitive tasks. For illustrative purpose, we set cost of collecting one sample $c = 0.006$ and prior probability to $\text{Beta}(1,1)$. This leads to the termination conditions for sequential sampling in Figure 2D$_1$, which shows a collapsing threshold.

We assume that BASS is performing amortised inference, and because the decision and estimation tasks are so similar, the samples drawn for decision are simply averaged to produce the estimate. Like SAS, BASS produces a high degree of consistency between decision and estimate and a biased toward the decision boundary for extreme stimuli. However, unlike SAS, BASS produces estimates that are biased away from the decision boundary for central stimuli (Figure 2D$_3$). The reason for the model's behavior can be seen in the termination conditions shown in Figure 2D$_1$. The sampling process is very unlikely to stop when there are an equal number of samples in favor of the two alternatives, instead waiting until there are more samples in favor of one of the alternatives. Then, after averaging the resulting samples to produce an estimate, these estimates are unlikely to be close to the decision boundary.

Table 1: Summary of model predictions on empirical effects

Effects	WD	SC
Decision bias	yes	yes
Self consistency	low	high
Belief distribution	bimodal	one-sided
	SAS	BASS
Decision bias	no	yes
Self consistency	high	high
Belief distribution	undistorted	favors one side

Note. WD=Weighted Decoding, SC=Self-Consistency, SAS=Simple Amortised Sampling, BASS=Bayesian Amortised Sequential Sampling.

Comparing the Models

As seen across Figure 2, qualitatively similar patterns of decision and estimation bias are predicted by the WD, SC, and BASS models. What distinguishes the models are the beliefs about the probability of each possible response in the estimation task. WD predicts that the optimal weighting will result in a bimodal belief distribution. SC predicts a one-sided belief distribution: that only estimates consistent with the decision will be considered. In contrast, BASS predicts that people will believe that several estimates are possible, including a low probability of those that are not consistent with the decision (see Table 1 for a summary).

We now test these predictions in a new experiment that includes a new type of trial that is used to elicit participants' belief distributions over possible estimates. We use a numerosity task in this experiment both to generalize the results, and because the discrete responses required in a numerosity experiment make it easier to elicit a belief distribution.

Experiment

Participants

Twenty-four participants (12 Males, ages between 18 and 35) were recruited through SONA system, University of Warwick. They received £4 for completing the experiment.

Materials

Participants were shown a briefly appearing number of dots (0.5 sec) on computer screen in a series of trials. The true number of dots was uniformly distributed between 21 and 30, and participants were explicitly told this at the beginning of the experiment. To generate a stimulus, dots were randomly positioned within a circular field subject to a minimum spacing between any two dots of four times the dot size. To encourage reliance on numerosity, rather than low-level visual features, the dot sizes varied uniformly between 3 and 9 pixels, and the radius of the dot field also varies uniformly between 150 and 450 pixels.

There were three different trial types (Figure 1B). For all trial types, participants made an initial decision as to whether there were 25 or fewer or 26 or more dots, so that the trials were identical until after this point. Following the decision, participants either immediately received feedback (D-F trials), were immediately asked to estimate the number of dots (D-E trials), or were immediately asked to state their beliefs about the number of dots using a histogram (D-H trials). When given a histogram, participants were asked to distribute 100 tokens among all of the possible numbers of dots according to how likely they believed these numbers were on that particular trial.

Procedure

Before the main experiment, participants received one practice example for each of the D-F, D-E, and D-H trials (Figure 1B). Participants additionally received feedback during practice, to introduce them to the point system used in the experiment that was used to encourage them to engage with the more demanding histogram trials. Correct decisions and estimates were both worth one point, while the number of points assigned to a histograms was $R = 100 \times [(1 - T_i/100)^2]$, where T_i was the number of tokens placed on the correct response. This formula is based on the Brier score, which incentivizes accurately reporting a belief distribution. Participants were also told, "If you had placed all the tokens on the correct number of dots, you would have scored 100 points. But if you had placed no tokens on the correct number of dots, you would have scored 0 points." Points were tallied throughout the experiment, but were only displayed to participants at the end of the experiment.

Results and Discussion

Decisions As shown in Figure 3A, participants were more likely than not to pick the correct answer for each true number of dots, but were never perfect.

Estimates Figure 3B shows the results of the estimates from the D-E trials for each true number of dots. True numbers that were close to the edge of the range showed an average bias towards the decision boundary, as participants tended not to respond outside the allowable range, and these responses outside the allowable range were excluded from further analysis (7.32%).

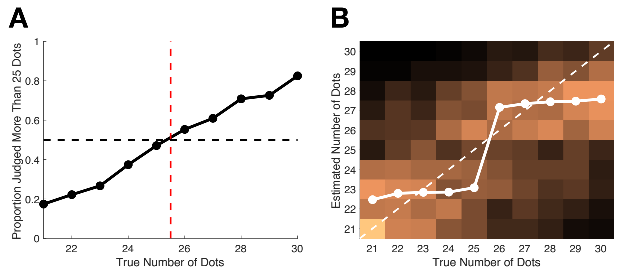

Figure 3: Behavioural data from the decision-estimation trials. **(A)** The psychometric curve describing the relationship between the true number of dots displayed and the average proportions of times participants judged the number of dots to be 26 or more. **(B)** The estimated number of dots (solid white line) systematically deviates from the true number of dots (dashed white line), constituting a perceptual bias.

For true numbers of dots near the decision boundary, particularly for numbers 25 and 26, participants were biased away from the decision boundary, in line with past work on the perceptual bias. When true number of dots was 25, estimates were on average smaller than 25, $t(23) = -14.97, p < .001$. When true number of dots was 26, estimates were on average larger than 26, $t(23) = 6.37, p < .001$. This estimation bias is, as expected, further qualitative evidence against the SAS model.

Histograms Figure 4A shows the average belief histogram following a decision of 25 or fewer, and the average belief histogram following a decision of 26 or more. These average histograms show greater mass on the side of the boundary consistent with the decision, indicating that overall participants engaged with these trials. They also show no evidence of bimodality as WD predicts, nor is the mass completely on one side of the boundary as SC predicts. Instead qualitatively these average histograms are most consistent with BASS.

To quantitatively test for whether there were the kind of bimodal histograms that WD predicts, we computed the average proportions, for each participant, that the mean token mass on the boundary numbers (i.e., 25 and 26) would be lower than either the mean token mass on 21-24 or the mean token mass on 27-30 (Figure 4B). If responses were random, we expect a $1/3$ chance that the mean token mass on 25 and 26 would be smallest. However, there were very few histograms that were bimodal, fewer than would be predicted by random responding, $t(23) = -7.68, p < .001$.

We then quantitatively tested whether all of the belief mass was on one side of the boundary, as SC predicts. We first estimated how much true responding there was, by calculating the proportion of trials on which a participant's histogram was consistent with their decision. Per participant, we calculated the proportion of trials on which the majority of token mass matched the decision. Next, we calculated,

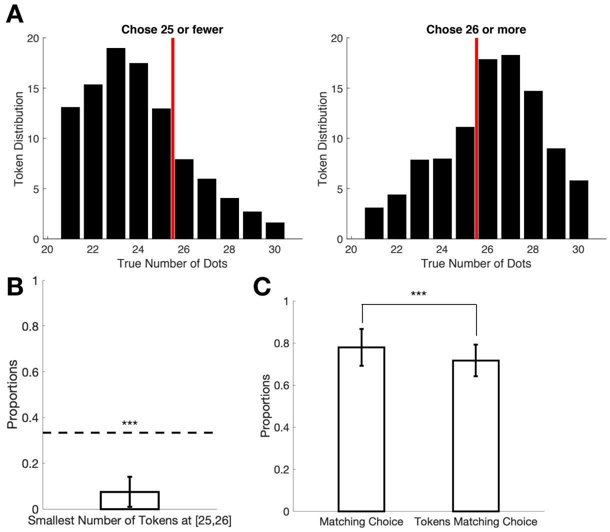

Figure 4: Results from the belief distributions reported in decision-histogram trials. **(A)** Average histograms when people chose "25 or fewer" (*left*) and when people chose "26 or more" (*right*). **(B)** The proportion of times when the token mass on central boundary numbers (25 and 26) is smallest. Error bar indicates 95% confidence interval across participants. **(C)** The proportion of times when preponderance of tokens matches the choice (*left*) and the percentages of token mass consistent with the choice (*right*). Error bars indicate 95% confidence interval across participants.

per participant, the proportion of tokens on the same side of boundary as the decision, as a measure of whether inconsistent estimates were considered. These two quantities, shown in Figure 4C, were different, $t(23) = 3.46, p = .002$, showing that tokens were certainly placed on numbers that disagreed with the decision, a number that exceeded what was expected from our estimate of noisy responding.

The results of the histogram trials are in line with the predictions of BASS, in which samples from both sides of the boundary are expected to be reused for the estimate. According to BASS, the amount of tokens placed on the opposite side of choice is a consequence of stochastic samples from the posterior distribution. Due to the termination conditions, there are always more samples that match the decision than those that mismatch the decision, but there are often samples on both sides of the boundary.

Conclusions

We proposed a new explanation for the decision bias in perception. To make a decision, we assume that participants sequentially sample hypotheses about the true nature of the environment, and stop when they have strong enough evidence in favor of one alternative over the other. As an application of amortised inference, we assumed that participants then reuse the samples to save cognitive effort, averaging them to produce their estimate. The bias in the estimate occurs because the samples that were sequentially obtained are never balanced between the options, and so estimates tend to be biased away from the decision boundary.

We generalized the bias from orientation tasks to a numerosity task, showing that it also occurs when participants give discrete responses to these perceptually complex stimuli. Using a novel type of trial in which we elicited participants' beliefs about which numbers were likely on a single trial, we found evidence for sequential sampling over other explanations of the decision bias in perception.

The sequential sampling model we evaluated here, BASS, is a novel application of the work of Wald (1950) for optimally deciding when to stop sampling from a binomial distribution with unknown probability. Of course it is almost certain that other sequential sampling approaches, such as those by Vul et al. (2014) and Bogacz et al. (2006), would predict the same qualitative results. Discriminating between these sequential sampling approaches will likely require quantitative comparisons, which is an interesting avenue for future work.

Additionally, a soft version of the self-consistency model, which relaxes the assumption that self-consistent estimates are made by every participant on every trial, could reproduce the qualitative results here. To distinguish the BASS model from a soft self-consistency account, we could test whether the belief distributions are a mixture of self-consistent sensory representations and unmodified sensory representations. However, to properly answer this question, we will need a further experiment that characterizes unmodified sensory representations for comparison.

Another avenue for future work is to explore the extent to which amortised inference and sequential sampling can explain other psychological biases. One interesting possibility is the anchoring bias (Tversky & Kahneman, 1974). In anchoring, participants are first asked to make a decision about whether a number, such as the percentage African countries in the UN, is smaller or larger than an often transparently irrelevant number, such as a number that results from the spin of wheel of fortune. Then participants are asked to make their estimate. While the task participants engage in is almost identical to the one that we used here, the effect that is found is the opposite: participants estimates are biased *toward* the anchor (Tversky & Kahneman, 1974). A unified explanation of these similar perceptual and cognitive biases will need to account for both the push and pull that decisions can exert on subsequent estimates.

Acknowledgements

J.Q.Z, A.N.S, and N.C. were supported by a grant from the National Institute of Economic and Social Research from their program Rebuilding Macroeconomics. A.N.S was supported by a European Research Council consolidator grant (817492-SAMPLING). N.C. was supported by the Economic and Social Research Council Network for Integrated

Behavioural Science (ES/P008976/1) and the Leverhulme Trust (RP2012-V-022).

References

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700.

Coenen, A., & Gureckis, T. M. (2016). The distorting effect of deciding to stop sampling. In *Proceedings of the annual meeting of the cognitive science society.*

Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general bayesian framework for metacognitive computation. *Psychological Review*, *124*(1), 91.

Gershman, S., & Goodman, N. (2014). Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).

Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, *446*(7138), 912.

Luu, L., & Stocker, A. A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife*, *7*, e33334.

Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, *20*(12), 883–893.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive Science*, *38*(4), 599–637.

Wald, A. (1950). Statistical decision functions.

Zamboni, E., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? decision-biases in motion perception. *Proceedings of the Royal Society B*, *283*(1833), 20160263.

Zhu, J.-Q., Sanborn, A., & Chater, N. (2018). Mental sampling in multimodal representations. In *Advances in Neural Information Processing Systems* (pp. 5753–5764).