

Perceptual decisions result from the continuous accumulation of memory and sensory evidence.

Aaron M. Bornstein^{1,*}, Mariam Aly¹, Samuel F. Feng³,
Nicholas B. Turk-Browne^{1,2}, Kenneth A. Norman^{1,2}, Jonathan D. Cohen^{1,2}

1. Neuroscience Institute, Princeton University, Princeton, NJ, USA

2. Department of Psychology, Princeton University, Princeton, NJ, USA

3. Department of Applied Mathematics and Sciences, Khalifa University, Abu Dhabi, UAE

* To whom correspondence should be addressed: aaronmb@princeton.edu.

Abstract

How do we combine memories with sensory input to make decisions? Previous research has shown that perceptual decisions can be made on the basis of prior expectations combined with sensory input. To date, these expectations have been treated as static, received quantities, fixed across decisions of the same type. Here, we tested the hypothesis that expectations can themselves be inferred using dynamic evidence accumulation, in a process continuous with that of sensory inference. In two experiments using a novel cue-guided perceptual decision task that independently varied memory and sensory evidence, we tested the degree to which decisions reflected accumulation of both kinds of information. In Experiment 1, we found that participants' response times and choices matched the qualitative and quantitative predictions of a two-stage evidence accumulation model. In Experiment 2, participants performed the same task while being scanned using fMRI. Using neural pattern analysis, we measured the expectations that participants formed in advance of a noisy visual stimulus on each trial, and found that these trial-specific expectations reliably predicted the speed of subsequent responses. These results demonstrate that perceptual decisions rely on a continuous process of evidence accumulation, that begins by dynamically inferring possible responses even before sensory information is available.

Good decisions should draw on all available useful information. Laboratory studies of decision-making tend to focus on choices made on the basis of a single kind of information – such as anticipated utility [1], sensory input [2], or mnemonic evidence [3, 4] – taken alone. But in the real world, our decisions depend on integrating across sources.

For instance, when traveling on an unfamiliar train route, I might miss my intended stop. How do I figure out where to make the transfer to get back on my desired route? I could rely solely on sight – as the train stops at each station, quickly scan the platform for helpful signs or markings. I could rely solely on my memories – which station is next? Will it have

the transfer I need? Both kinds of information can be unreliable: station platforms may look very similar, with distant or unhelpful signage, or my memories could be sparse and unclear. More likely, I will combine both kinds of information: query my memories about which stations might have transfers, and combine those with what I can see from a quick look out the door at each stop. By combining what I remember with what I see, I can improve my ability to figure out where I am, and thus where I need to go.

An open question in the study of perceptual decisions is how expectations should, and do, influence the inference process. Within the canonical evidence-accumulation framework [2, 3], expectations could be encoded as a change to either the starting point of accumulation, or the rate at which evidence is accumulated [5–8]. Another, related idea is that expectations can dynamically impact both the rate and direction of accumulation, with increasing influence as a decision takes longer to resolve [9]. However, all of these approaches assume that the content of expectations is fixed before the decision starts, whether by learning or by instruction. In the train analogy, the map is known with certainty, though the reliability of the visual cues vary from station to station (trial to trial).

Recently, we and others have shown that decisions can be made on the basis of sampled memories, similar to the way in which samples of visual input are used to guide perceptual decisions [3, 10–13]. Building on these results, we test the hypothesis that these two types of sampling are actually components of a single, continuous, inference process, where actions are selected on the basis of the combined predictions of both kinds of evidence.

Our hypothesis yields two main predictions. First, evidence accumulation should begin before the onset of sensory information, with dynamics that change when the stimulus is presented. Specifically, before the visual stimulus, accumulation should reflect the contents of memory, and their reliability; after, the rate of accumulation should be inflected according to the coherence and content of visual information. Second, the hypothesis predicts that the accumulation process is integrative across modalities. Specifically, decisions made after the onset of a visual stimulus should reflect memory samples collected before the onset of the visual stimulus – if the memory samples concord with the visual samples, then the decision should be faster.

To test our predictions, we developed a memory-guided perceptual inference task. In the task, two distinct kinds of information – memory and sensory – indicate the correct response for that trial, and are made available at separate times. More specifically, a fractal cue probabilistically signals the likely identity and reliability of a noisy visual stimulus – a rapidly alternating stream of face or house images – which follows a few seconds later. Critically, participants can select a response whenever they wish, including before the image stimulus appears. Therefore, their responses can reflect the influence of memory or sensory information alone, or some combination of the two.

We formalized our predictions using a new, two-stage evidence-accumulation model [14]. The first stage of the model samples evidence from memories triggered by the fractal cue. The second stage carries forward the evidence accumulated from stage one, while incorporating new samples of evidence, this time visual input from the noisy image stimulus. This approach differs from previous models of expectation-guided perceptual inference in that it constructs

expectations dynamically for each trial. As a result, what the model expects will vary between decisions, depending on what evidence was sampled during the first stage.

Experiment 1 is a behavioral study that tests the first prediction of our model: that choices and response times in this task reflect a continuous, two-stage evidence-accumulation process whose rate of accumulation changes with the onset of visual information. We fit our hypothesized model to these data, and contrast its fit with that of more standard models. Experiment 2 is an fMRI study that tests the second prediction: that evidence accumulated from memory is carried forward and affects the sensory inference process. We use Multivariate Pattern Analysis (MVPA) to measure, on a trial-by-trial basis, neural evidence for accumulation from memory, and test its relationship to responses made after the onset of the image stimulus.

Taken together, the results of these experiments provide a new account of perceptual decisions, by demonstrating a critical role for integrated, dynamic inference from mnemonic, as well as sensory, information.

Results

Participants performed a cue-guided perceptual inference task (Figure 1a,b), in which fractal visual cues could be used to anticipate the content of a noisy perceptual stimulus that appeared after a short inter-stimulus-interval (ISI). In the *Learning* phase (Figure 1a), participants learned, by experience, how often each cue was followed by each photograph. In this phase, photographs were presented one per trial, without any perceptual noise. Photographs were from one of two categories: faces and scenes. Each cue was associated with photographs from one category only. In the *Test* phase (Figure 1b), the perceptual stimulus consisted of a rapidly, stochastically, alternating stream of two photographs, one of which, the *target* photograph, appeared more often than the other. In this phase, each of four fractal cues signaled two different quantities: the probability that a given photograph would be the target (the *cue probability*), and the relative prevalence of that target in the stream (the *coherence*). The coherence of the stimulus was fixed by category – on a given experimental run, either faces or scenes would be easier to discriminate (higher coherence), and the other would be difficult (lower coherence).

In both phases, participants were instructed to press only the key identifying the target photograph. (Key-photograph pairings were learned, and coherence levels calibrated, in separate training phases preceding the Learning phase.) Importantly, participants could use the cued information to respond before the onset of the flickering stream – though there was, by design, no explicit external incentive to respond early, or even correctly.

Most participants performed the task twice, in *blocks* with different photographs, fractals, cue-photograph probabilities, and mappings from category to coherence level. Due to time limitations or technical problems, some participants only completed one block (see *Methods* for details). In Experiment 1, 30 participants completed 56 blocks in total. In Experiment 2, a separate set of 31 participants performed the same task in an fMRI scanner, completing 52 blocks in total. Because participants completed different, pseudorandomly-assigned, subsets

of the task conditions (the four cue and coherence levels), statistics on the correlations between cue level and dependent measures were computed by bootstrap across participants. When comparing how these correlations varied, across participants, between pairs of different trial types – e.g. low versus high coherence, or pre-stimulus versus post-stimulus responses – we used paired bootstrap analyses for each level, and measured the difference between the bootstrapped distributions using Cohen’s d . For within-participant comparisons (e.g. between paired coherence levels) we used standard parametric tests.

Experiment 1

We first tested whether choices and response times reflected the influence of both cue probability and perceptual coherence. By our hypothesized mechanism, participants should respond more quickly and accurately when the cue predictions are more reliable, when the observed perceptual evidence is more coherent, and when the cue predictions match the perceptual evidence.

Response times and accuracy

Responses reflect cue predictions and perceptual evidence. Participants responded accurately, matching the target photograph on 75.20% (SEM 0.085%) of trials (including only trials for which there was a “correct” response possible before stimulus onset – i.e. for cue levels 60%, 70%, 80%) This proportion was reliably greater than chance for all blocks individually (all $p \leq 0.0465$ by binomial test of the proportion of correct responses within each block against the 50% chance level). Accuracy increased with both cue predictiveness ($R = 0.1947$, $P = 0.009$ by bootstrap across participants; Figure 2a) and target coherence ($t(27) = -4.4297$, $P = 0.0001$; defined only for the 28 participants who performed at least one block in which there was a correct response in both coherence conditions).

Participants used the fractal cue to decide whether or not to respond “early”, before the onset of the perceptual stimulus. The proportion of early responses increased with the predictiveness of the fractal cue ($R = 0.2224$, bootstrap $P < 0.0001$; Figure 2b). This relationship was driven by trials on which the cue signaled that the perceptual stimulus would be of low coherence (for low coherence trials, the correlation between cue predictiveness and early responses was $R = 0.3658$, bootstrap $P < 0.0001$; for high-coherence trials it was $R = 0.0868$, bootstrap $P = 0.1610$; difference between paired bootstrap iterations reliable with Cohen’s $d = 3.9071$).

Reflecting the two-stage nature of the integration task, response times were distinctly bimodal, with separate peaks following the onsets of the fractal cue and the flickering stream (Figure 3a; RT distributions multi-modal within each ISI condition by Hartigan’s Dip Test [15]: all $HDS \geq 0.0281$, all $P < 0.0001$).

For early responses, RTs showed a trend towards being faster as the fractal was more predictive ($R = -0.0351$, $P = 0.086$ by bootstrap across participants). Consistent with the normative prediction that uninformative cues discourage deliberation [16], this relationship

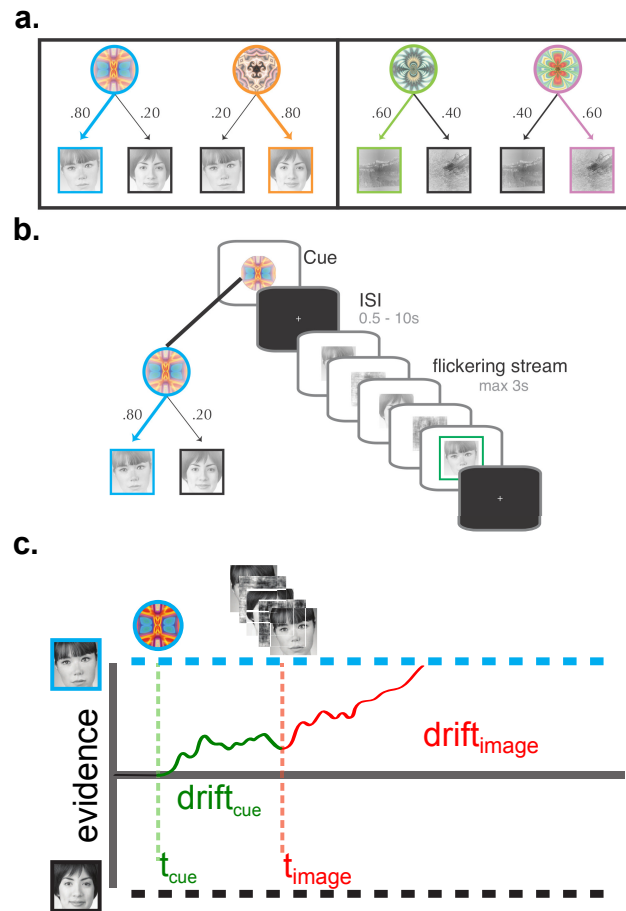


Figure 1: Cue-guided perceptual inference task. (a) In the *Learning* phase, participants performed 100 trials to learn which face or scene stimulus was predicted by each of four fractal cues. Each cue was shown 25 times, at each presentation followed by a picture of either a face or a scene. Cues were linked to one photograph category only, and to each picture within that category according to complementary proportions (50%/50%, 60%/40%, 70%/30%, 80%/20%). (b) In the *Test* phase, participants were again shown a fractal cue, but in this case the cue was followed by a “flickering” series of rapidly alternating (60Hz) pictures. Each frame of the series contained one of the two pictures from the cued category, or a phase-scrambled version of the previously-shown picture that served as a perceptual mask. One picture, the *target*, was shown in the stream more often than the other. The fraction of frames that contained the target picture was calibrated to elicit either higher (85%) or lower (65%) accuracy. In this way, the flickering stimulus parametrically varied the extent to which the visual evidence favored the given stimulus. Participants were asked to respond by pressing the key associated with the target picture. Critically, they were allowed to respond before the flickering stream appeared on screen. (c) Two-stage continuous-accumulation model. The *Multi-Stage Drift-Diffusion Model* (MSDDM) describes an evidence-accumulation process with time-varying drift rate [14]. The first drift rate corresponds to the period following onset of the fractal cue and preceding the onset of the flickering stream, while the second corresponds to the period after the onset of the stream. The end point of the first stage sets the starting point of the second.

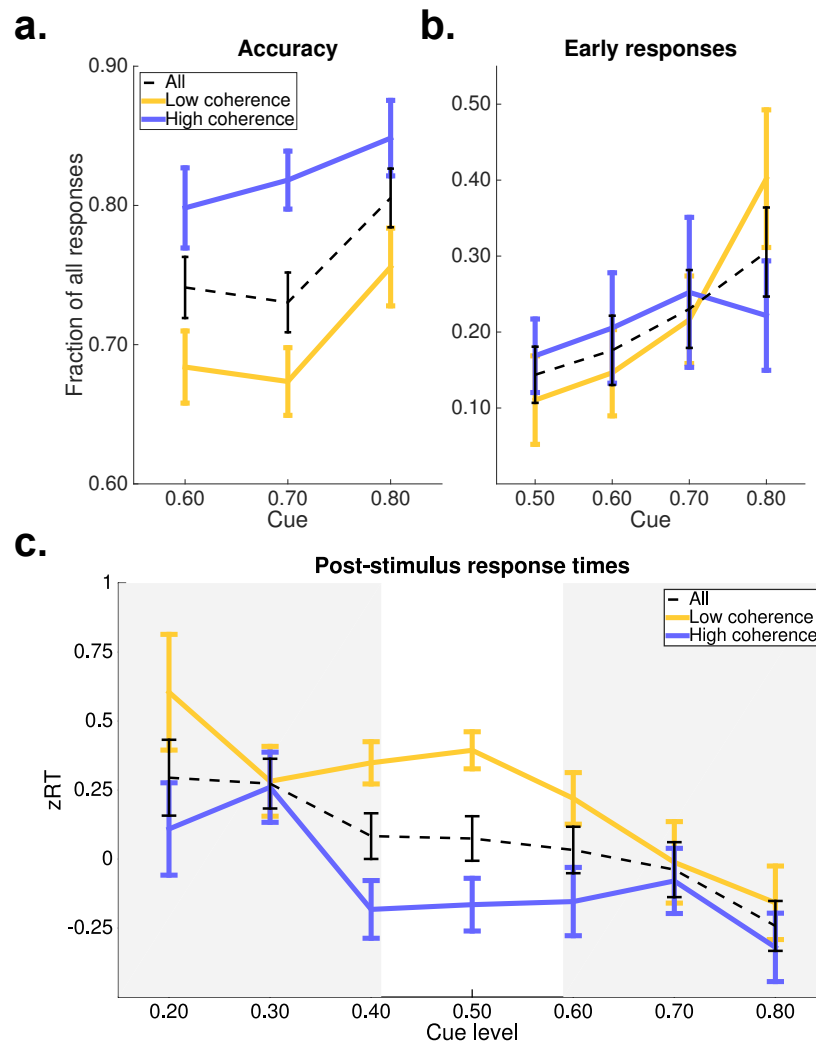


Figure 2: Early responses and accuracy scale with cue predictiveness and stimulus coherence. (a) Across all post-stimulus accuracy increased with both cue reliability and the coherence of the flickering stream. (b) Participants were more likely to respond early when the fractal cue made a more reliable prediction about the identity of the upcoming stimulus. This relationship was stronger when participants knew that the upcoming stimulus would be of low coherence (low: $R = 0.3658$, bootstrap $P < 0.0001$; high: $R = 0.0868$, bootstrap $P = 0.161$; Cohen's $d = 2.3841$). (c) When perceptual evidence was of high coherence, participant responses after the onset of the flickering stream were faster than on low-coherence trials ($t(29) = 3.145$, $P = 0.0038$). Across all such responses, response time decreased with the cued probability of the target photograph. This relationship was stronger for low-coherence trials, and for trials where the target photograph matched the one predicted by the fractal cue (right shaded region), as opposed to when the predicted image did not match the target (left shaded region).

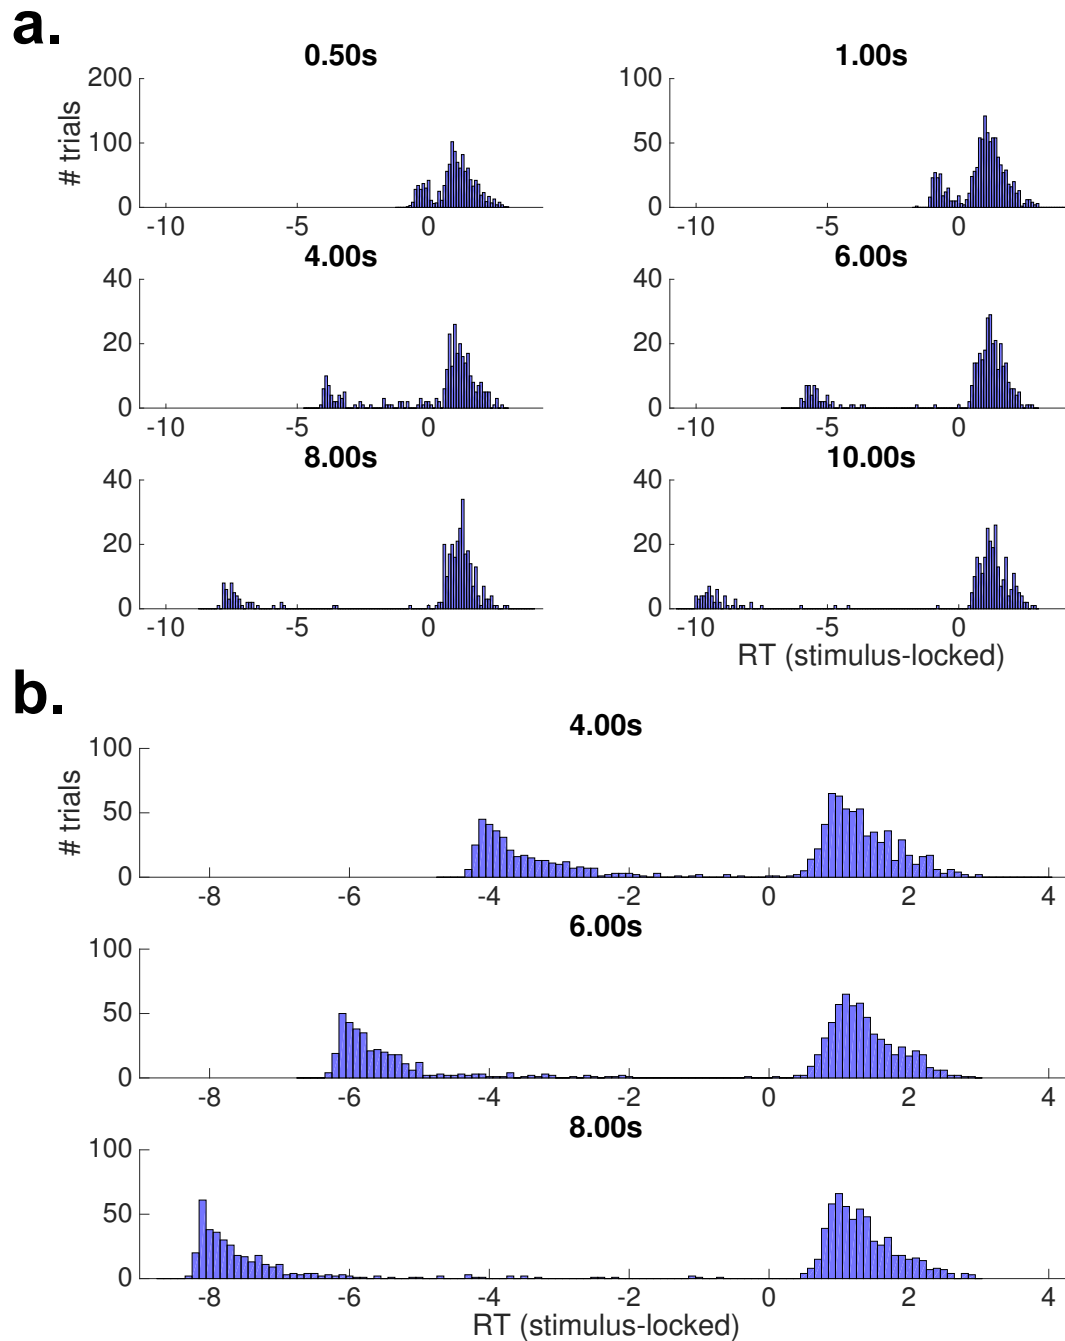


Figure 3: **Histogram of (stimulus-locked) response times on test-phase trials.** RT counts are aggregated across participants, and plotted separately for each ISI condition, binned in increments of 100ms. Separate peaks follow the onset of the fractal cue and the onset of the flickering stream, reflecting the fact that participants made responses on the basis of both types of information. **(a)** Experiment 1, ISIs of: 500ms, 1s, 4s, 6s, 8s, 10s. **(b)** Experiment 2, ISIs of: 4s, 6s, 8s.

was more pronounced, and reliable, when including only informative cues (60% and higher; $R = -0.0501$, bootstrap $P = 0.008$).

For post-stimulus responses, RTs were faster when the cue was more predictive ($R = -0.017$, bootstrap $P = 0.044$; Figure 2c), and when the flickering stream was higher coherence (mean RTs – stimulus-locked, log-transformed, Z-scored within participant: low: 0.1581 SEM 0.0612 high: -0.1283 SEM 0.0387 mean difference between low and high coherence RTs within-participant 0.2864 SEM 0.0911; $t(29) = 3.145$, $P = 0.0038$). These factors interacted: participants were more speeded by cue predictiveness when coherence was lower (low: $R = -0.147$, bootstrap $P < 0.0001$; high: $R = -0.065$, bootstrap $P = 0.033$; Cohen’s $d = 1.675$), and only when the target photograph matched the cue’s prediction (incongruent: $R = 0.0419$, bootstrap $P = 0.141$; congruent: $R = -0.063$, bootstrap $P < 0.001$; Cohen’s $d = 1.7385$).

Taken together, these results confirm that participants’ responses reflected the integration of mnemonic and sensory information.

Model comparison

We used model comparison to formally test the hypothesis that responses resulted from continuous accumulation of mnemonic and sensory evidence. Our primary model of interest implemented a continuous, two-stage evidence-accumulation process (hereafter: MSDDM; [14]; Figure 1c), – the first on the basis of the cue, preceding the flickering stream, and the second following the onset of the flickering stream – with different accumulation rates at each stage.

The MSDDM is distinguished from other models by two key features: first, that the drift rate changes at the time of flickering stream onset and second, that accumulation in the second stage proceeds from the evidence accumulated during the first stage. Therefore, our comparisons were against sampling models that selectively disabled each of those features. The first comparison model was a single DDM, which had continuous accumulation until the time of response, but no change in drift rate across the entire trial. We refer to this model as *1DDM*. The second comparison model was two unconnected DDMs, mirroring the change in drift rate found in MSDDM, but with the second-stage starting point set independently of the behavior of the first model. We refer to this model as *2DDM*. Each model was fit separately to responses aggregated, across participants, by condition – cue, coherence, and ISI.

Against the second-best model, the 2DDM model, MSDDM was superior by a BIC of 30 (constituting “Strong” evidence in favor of the model, according to [17]). This was the case across all conditions, and for nearly every condition individually (Figure 4a).

Experiment 2

Experiment 1 showed that behavior in this task reflects a dynamic integration of memory and sensory evidence, yielding patterns of choices and response times that are best captured by the MSDDM, across all trials. However, the key test of the hypothesis is that evidence

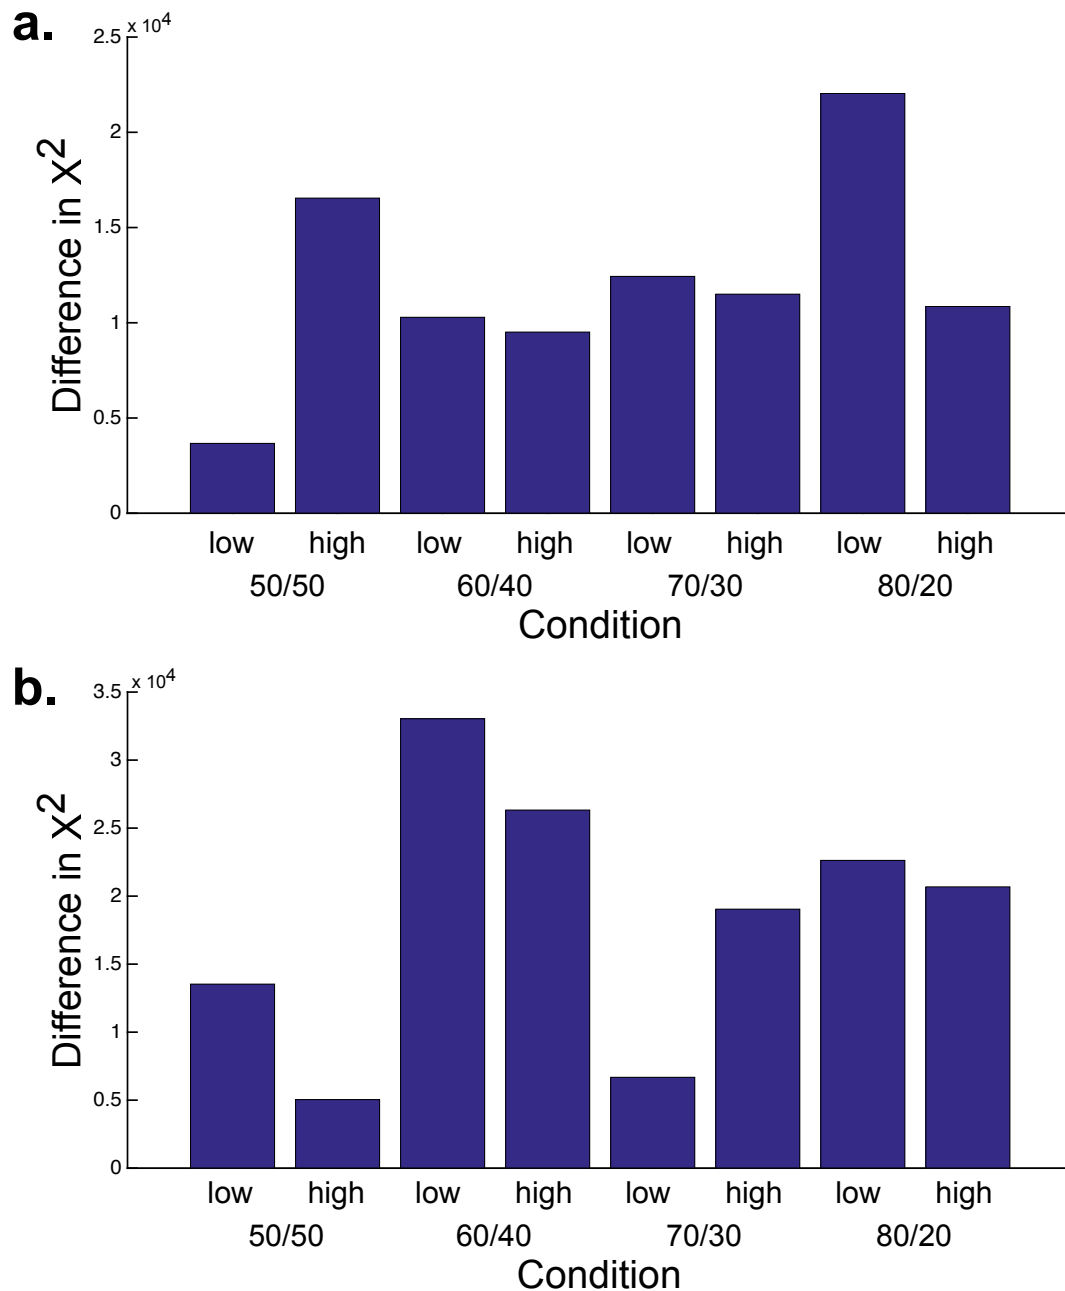


Figure 4: **Model comparison.** Models were compared for their fit to each bin of trials with the same combination of cue predictiveness, stimulus coherence, and ISI. Shown here is the χ^2 difference in favor of the MSDDM model (higher bar = greater evidence in favor MSDDM), for conditions aggregated by cue predictiveness and coherence level. MSDDM was favored over 2DDM for every condition individually, and across all conditions as a whole. (a) Experiment 1 (b) Experiment 2

accumulated during the first stage *on a specific trial* affects responses made during the second stage of that trial. In Experiment 2, we used multivariate pattern analyses (MVPA) of fMRI data to measure memory evidence accumulated following the fractal cue on each trial, and used this measure to predict responses after the onset of the flickering stream on that same trial. For this experiment, 31 additional participants completed the task from Experiment 1, while being scanned.

Behavior

Response times and accuracy.

Response behavior replicated the patterns observed in Experiment 1. Accuracy was again high overall: 70.24% correct responses (SEM 1.18%); and significant or trending for 49/52 blocks individually (all $p \leq 0.0728$).

Accuracy again increased with cue predictiveness ($R = 0.2474$, bootstrap $P = 0.005$) and coherence ($t(26) = -4.3017$, $P = 0.0002$). RTs were bimodal (all $HDS \geq 0.1017$, all $P < 0.0001$; Figure 3b). Higher cue predictiveness resulted in a greater tendency to respond early ($R = 0.1494$, bootstrap $P = 0.012$), though, in contrast to experiment 1, the effect was driven by high-coherence trials (low: $R = -0.1022$, bootstrap $P = 0.111$; high: $R = 0.4187$, bootstrap $P < 0.0001$; Cohens $d = 4.2134$), perhaps reflecting that early responding was at ceiling when participants anticipated low coherence stimuli. Cue predictiveness also caused early responses to be faster ($R = -0.0768$, bootstrap $P < 0.0001$).

Responses after the onset of the flickering stream were again speeded by coherence (low: 0.2268 SEM 0.0421; high: -0.0981 SEM 0.0782; mean difference 0.3248 SEM 0.0977; $t(30) = 3.3255$, $P = 0.0023$), and by cue predictiveness, in both coherence conditions (low: $R = -0.0921$, bootstrap $P = 0.021$; high: $R = -0.0599$, bootstrap $P = 0.013$), though moreso when coherence was lower (Cohen's $d = 0.7604$), and when the target photograph matched the cue's prediction: (incongruent: $R = 0.0249$, bootstrap $P = 0.303$; congruent: $R = -0.0165$, bootstrap $P = 0.191$; Cohen's $d = 0.9447$).

Finally, model comparison again favored the MSDDM over the alternative candidate models (all BIC > 21 ; Figure 4b).

Neuroimaging

Our analyses of choices and RTs in this task supported the hypothesis that perceptual decisions can be made on the basis of both cue-driven expectations and sensory input. However, because they measure only the final response, behavioral data cannot in principle reveal a relationship between the actual memory evidence accumulated on each trial and responses made to the ensuing flickering stream. Here, we used neural pattern similarity to measure the influence of accumulated memory evidence on responses. For each participant, we localized regions in the ventral visual stream that were more active for face versus scene processing (FFA; [18]) and for scene versus face processing (PPA; [19]) (Figure 5a). We next computed activity patterns corresponding to each photograph, in the appropriate category-preferring

region (faces in FFA, scenes in PPA). We refer to these picture-specific patterns as the *target patterns*. The target patterns were defined on the basis of data from an early training phase of the task, in which participants learned which keys were mapped to each picture (see *Methods*). Critically, because this response-training phase preceded the introduction of the fractal cues, these patterns were decoupled from the fractals that would come to predict them.

We next computed, for each trial from the Test phase, a *trial pattern* – the average pattern in these regions over the period following the onset of the fractal cue, up to either the participant’s response, or one TR before the onset of the flickering stream, whichever came first. Hereafter, we define the trial-by-trial *reinstatement index* as the correlation between these trial patterns and the target pattern corresponding to the photograph predicted by the fractal cue. (Note that on 50/50 trials this value is not defined, and so these trials were excluded from neuroimaging analysis.)

Pre-stimulus reinstatement scales with task conditions.

As in a previous study of memory sampling [10], we expected that anticipatory stimulus reinstatement would scale negatively with cue probability. The reason for this is that in evidence-accumulation models, samples are accumulated until a threshold is reached in favor of a given response. At that point, no more sampling – and thus no more activity linked to memory reinstatement – should follow. Because memory sampling likely occurs at a timescale below the resolution of the BOLD response [20], more predictive cues should lead to a shorter period of reinstatement, and thus less consistent trial patterns which would be less correlated with the template pattern. Consistent with the prediction, reinstatement index decreased as cue probability increased ($R = -0.0341$, bootstrap $P = 0.039$; Figure 5b).

Following the same logic, this relationship should only hold for early responses and not on trials where the participant responded following stimulus onset. This is because early responses are by definition those trials on which accumulation terminates in the time period we are measuring, preceding the onset of the flickering stream. Conversely, responses to sensory input are those on which memory accumulation continues throughout the measurement period. Indeed, although the correlation between cue probability and reinstatement index was reliably negative for early response trials ($R = -0.0718$, bootstrap $P = 0.004$), no such correlation was observed for responses after the onset of the flickering stream ($R = 0.0203$, bootstrap $P = 0.225$; difference between paired bootstrap iterations Cohen’s $d = 1.7417$).

Pre-stimulus reinstatement predicts 2nd stage response times. The fact that reinstatement index is related to cue probability is consistent with evidence-accumulation models, but does not differentiate between MSDDM versus 1DDM or 2DDM. The key test of the MSDDM is whether responses to the flickering stream are influenced by evidence-accumulation in anticipation of sensory input, at a trial-by-trial level. Therefore, we tested whether reinstatement index could predict response times after the onset of the flickering stream on each trial.

Supporting our hypothesis, reinstatement index was indeed associated with faster post-stimulus response times ($R = -0.07$, bootstrap $P = 0.001$), a relationship that held after

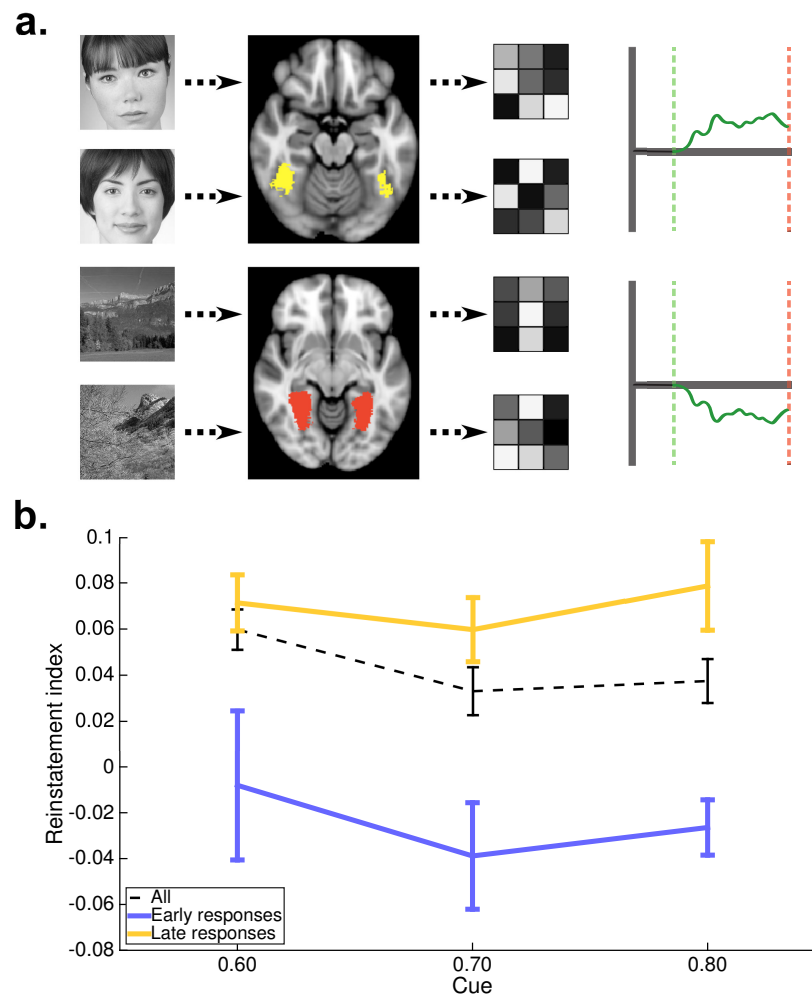


Figure 5: Reinstatement pattern analysis. (a) We defined, for each participant, the Fusiform Face Area (FFA) and Parahippocampal Place Area (PPA), using a localizer task that followed the main experiment. The resulting mask defined the region across which we calculated picture-specific templates. Next, for each trial on which participants responded after the onset of the flickering stream, we computed the average pattern of activity in the corresponding ROI over the period following the onset of the fractal cue, but preceding the onset of that flickering stream. We then calculated a *reinstatement index* as the correlation between this trial-specific pattern and the template pattern for the picture predicted by the fractal cue. (b) Reinstatement index related to cue predictiveness. Across all trials, as the cue led to more specific predictions, the reinstatement index decreased ($R = -0.0341$, bootstrap $P = 0.039$). However, on trials where the response followed the onset of the flickering stream, reinstatement index was unchanged across cue levels ($R = 0.0203$, bootstrap $P = 0.225$). This observation can be interpreted as consistent with an evidence-accumulation process that reached bound quickly, under the temporal resolution of the BOLD response. However, on trials where participants responded after the onset of the perceptual stimulus, accumulation did not, by definition, terminate.

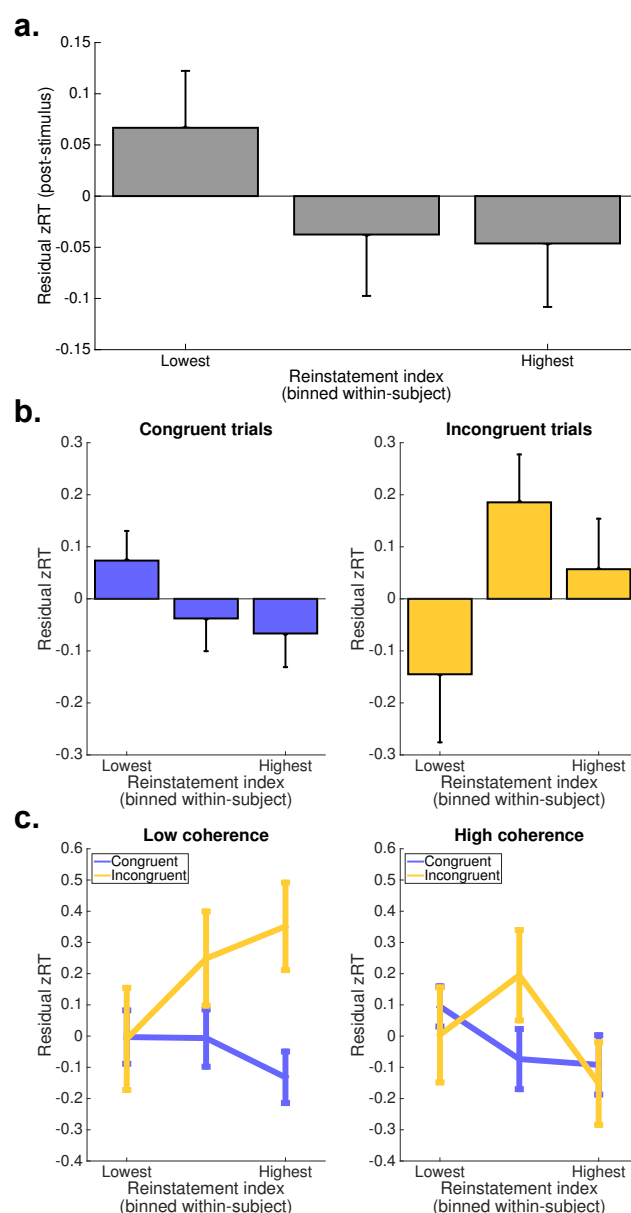


Figure 6: Accumulated memory evidence is incorporated into perceptual evidence-accumulation. (a) Across all trials, a higher reinstatement index predicted faster responding after stimulus onset. (b) Reinstatement index speeds responses to cue-congruent, but not cue-incongruent, sensory information. (c) The MSDDM predicts that the accumulated memory evidence should set the starting point for visual evidence accumulation. Thus, if the visual stimulus doesn't match the anticipated stimulus, responses should be slowed, relative to when evidence is congruent between stages. Consistent with this model, greater accumulated memory evidence slows responses to incongruent visual stimuli on low-coherence trials, but not high-coherence trials. (Error bars are ± 1 SEM, across participants.)

controlling for other factors each of which also modulate post-stimulus response times (cue predictiveness, coherence, ISI; $R = -0.0337$, bootstrap $P = 0.039$; Figure 6a). If accumulated memory evidence sets the starting point for sensory evidence accumulation, then reinstatement index should only predict faster responses when memory and sensory evidence are in agreement – and, indeed, RTs were uniquely speeded on such trials (cue-incongruent: $R = 0.0096$, bootstrap $P = 0.404$; cue-congruent: $R = -0.0527$, bootstrap $P = 0.029$; Cohen’s $d = -1.2037$; Figure 6b). Finally, if memory and sensory evidence are integrated, memory evidence should show correspondingly less influence when sensory evidence is stronger; this should be reflected as a greater speeding of matching, relative to non-matching, trials. Consistent with this prediction, the divergence in the relationship between memory evidence and response times to cue-congruent, versus cue-incongruent, flickering streams was more pronounced in the low-coherence condition: (low coherence, cue-incongruent: $R = 0.127$, bootstrap $P = 0.087$; low coherence, cue-congruent: $R = -0.0384$, bootstrap $P = 0.161$; Cohen’s $d = -1.586$; high coherence, cue-incongruent: $R = -0.0886$, $P = 0.117$; high coherence, cue-congruent: $R = -0.0633$, $P = 0.066$; Cohen’s $d = 0.4229$; Figure 6c).

Together, these results support a role for the dynamic accumulation of memory evidence in perceptual decisions, via a continuous perceptual inference process linking memory, sensation, and action.

Discussion

Humans [6, 21], animals [9], and even intelligent machines [22] rely on expectations, derived from experience, in order to act quickly and accurately [16]. While important empirical and theoretical work has described ways in which expectations influence dynamic, deliberative decisions [6, 7], these investigations have set aside the question of how the expectations themselves are set. Our study examined a corner of the space left unaddressed by previous work – that in which the dynamics of expectation setting could be separately manipulated, and measured, on a trial-by-trial basis. In doing so, we reveal that perceptual inference begins before sensory evidence is available.

The evidence accumulation framework used to model perceptual decisions also has roots in the study of recognition memory [3] – behavior that is, by definition, an inference over past experiences. Recently, we showed that decisions for reward are informed in part by “samples” of relevant past experiences [10, 12, 13]. At each new choice, decision-makers brings to mind similar previous choices, and take the outcomes of those choices as evidence for one option or the other. We reasoned that a similar dynamic inference process could underlie expectations formed before the onset of perceptual decisions.

We tested this hypothesis in two experiments using a novel cue-guided perceptual inference task that separately varied the content and reliability of expectations derived from previous experience, as well as the content and reliability of sensory input. Participants viewed a fractal cue followed, after a variable-length delay, by a “flickering” stream where two photographs were presented in rapid alternation, at a trial-specific ratio. Their task was to press a key corresponding to the photograph that appeared most often in the stream.

Participants learned to use the cues to anticipate which photograph would be the target, and also how difficult it would be to decide on the basis of sensory information alone. Accordingly, their responses reflected the influence of associations drawn from memory and online sensory evidence.

Response times were best-fit by a two-stage evidence accumulation model. This model describes decisions as arising from a continuous inference process that samples evidence at one drift rate until a given deadline, and then samples evidence at a second drift rate. The critical prediction of this model is that evidence accumulated during the first stage should influence responses made during the second stage. To test this directly, we used fMRI to observe, on each trial, participants' neural responses to the cue, in anticipation of the flickering stream. This measure was indeed a reliable predictor of response times – the more participants anticipated the photograph predicted by the cue, the faster they were to respond after the onset of the perceptual stimulus. Taken together, our results demonstrate that the dynamics of expectation-setting combine with those of sensory inference to affect responses.

Previous theoretical and empirical work has outlined the normative conditions under which expectations should bias inference [6, 7], and the mechanisms by which expectations are incorporated into the decision itself [9]. Several possible mechanisms have been proposed, and tested, across a range of tasks. These approaches differ on the form that expectations should have: a shift in the starting point of accumulation, the drift rate, or a dynamic bias on the accumulation process. In the case that memories are reliable, and also consistent with sensory information, then – to a model fit across many trials – accumulated memory evidence would appear as a change to the starting point, if memories are reliably predictive, or drift rate, if less predictive. The degree of this effect would be proportional to the average congruence of memory and sensory evidence: if memory and sensory evidence are reliably consistent with each other, then the effect of expectations on starting point and drift rate should be correspondingly higher. However, in cases where memory evidence is reliable, but sensory evidence is less decisive, the same process could appear as a dynamic bias signal, with steadily increasing influence over the course of a given decision. The rate of increase of this influence would be proportional to the coherence of memory evidence.

Continuous accumulation of memory and perceptual evidence also predicts that these effects vary, trial-by-trial, as a function of the relative reliability and consistency of the two kinds of evidence. In perceptual sampling models, this sort of variability in decision parameters tends to be modeled as noise [23]. Here, we attempt to unpack that noise, providing both an instrumental justification for, and a process-level explanation of, time-varying properties of the decision process.

These findings fit with a burgeoning literature on the influence of memory on decisions. Empirical work, by ourselves and others [4, 10, 12, 24, 25] has outlined a process by which decisions can be made on the basis of evidence accumulated from individual memories of past decisions (see Shohamy and Shadlen [11] for a proposal framing these findings in terms of evidence accumulation). Though we were motivated by our work on sampling from episodic memory, the exact content of the internal representation used to set expectations in this task remains unclear (e.g., individual episodes, semantic knowledge, motor responses).

A key differentiating factor of our study relative to previous work is that expectations were based on a relatively small number of direct, but variable, experiences – rather than explicit instruction or overtraining, which can diminish the need to rely on memory accumulation. Of particular relevance, Hanks and colleagues [9] evaluated how perceptual decisions can be aided by statistical regularities. They observed that the influence of priors grew as a function of the elapsed decision time on a given trial. On this basis they inferred that decision time was used as an online estimate of decision difficulty. Interestingly, the investigators report that neural “ramping” activity preceding stimulus onset reflected the prior expectations participants were trained to develop. While this activity did not predict eventual decision outcome in their task, the effect may have been impeded by the fact that the task was not designed to probe memory evidence accumulation – specifically, that the task offered limited time before the onset of sensory information, a reliable correspondence between expectations and sensory information, and expectations set either by overtraining or instruction, thus limiting the likelihood that subjects would need to dynamically infer expectations.

The synthesis of our results with theirs implies an interesting path for future research. Specifically, our results suggest that the weighting process they observe could begin before decision onset, while theirs suggest that inference over experiences may continue in parallel, after stimulus onset. An ongoing mixture of expectations and sensory information should fit with their account when sensory information is largely congruent with expectations, but may qualitatively diverge when trial-varying expectations – observed before the onset of the sensory stimulus – are incongruent with sensory information. A conclusive test of this hypothesis will likely require a method with finer temporal resolution than fMRI, as well as wider spatial coverage than single-unit recording.

More broadly, this sort of cooperative evidence accumulation – whether in serial or in parallel – may be present in other neural computations. Though various kinds of decisions have been modeled as arising from the competition of one or more “racing” accumulators [26, 27], the idea that multiple types of information might be sampled in a cooperative fashion, towards a single decision resolution, has been less widely investigated. Most interesting is the possibility that these two approaches – competitive and cooperative accumulation – need not be at odds, or entirely distinct. Race models are supported by the observation that response time distributions correspond to the accrued evidence for one alternative or the other, rather than the net between. A unifying explanation could be that accumulation is competitive at the level of each sensory modality or internal representation, but integrative across such inputs. For instance, ongoing samples from past experience could amplify consistent sensory evidence samples, or weigh against inconsistent ones.

A related question is how exactly it is that information is transmitted downstream from memory. One possibility is that reinstated memories provide a feedback signal that shapes processing in earlier visual regions; Recalled memories, triggered by visual information and internal state, could serve as a “teaching signal” that propagates upstream, perhaps as far as primary visual cortex. A recent framework proposes that such a mechanism explains a wealth of observations in perceptual learning and inference [28] Another, not incompatible, possibility is that accumulated memory evidence persists as a form of visual working memory.

In some decisions, participants may choose to forego further sampling, even if they have accumulated only partially to bound [29], though the accumulated evidence may still carry forward to the later inference process. In a separate study, we examined how reinstated episodic memories persist across a delay by altering the existing contents of working memory [30]. In that study, because participants could not anticipate the exact decision they would be asked to make, cues provided information about the range of options that would be available, resulting in anticipatory memory samples that affected the drift rate, but not the starting point, of response selection. Task demands may thus modulate the form of the influence that memory sampling has on later evidence accumulation.

More broadly, this finding has important implications for the study of decisions in many fields. Economic decisions, for instance, require integrating learned and/or inferred information about multiple attributes of the choice options. Individuals may differ in how they combine these attributes. The mechanisms by which they do so have been proposed to underlie some of the observed “irrationalities” of behavioral economics [31, 32] and measurements of the variation in the attention paid to different options has been shown to improve predictions of choices [33]. Here, we show how sampled memories can combine with sampled features, improving choice predictions over those made on the basis of either type of information alone.

Because expectations are nearly omnipresent in decision-making, it is possible that previous investigations have obscured an important source of trial-by-trial variation. Decisions may often be biased by samples from internal information – memories, but also emotions, values, and rules – that give rise to expectations established in the moment, rather than fixed across time. Biases, derived from experience, are necessary for helpful, even necessary, for efficient decision-making – they help us take account of, and leverage, the statistics of our environment. But biases aren’t just learned regularities. The fact that they are constructed means that they can be changed, either by changing experiences, or by targeted interactions with the construction process itself. That biases can be altered in the service of better decisions may be a crucial adaptation, allowing organisms to adapt their behavior at a timescale faster than the long-run statistics of their environment. More broadly, it means that, when it comes to individual decisions, the link between past and present can be revisited, even changed, when the need arises.

Acknowledgements

The authors wish to thank Abigail Hoskin, Amitai Shenhav, Judith Fan, Phillip Holmes, Michael Waksom and Roozbeh Kiani for helpful conversations, Ghootae Kim for providing ranked face and scene stimuli, Nicholas Hindy for providing fractal stimuli, and Charlotte Townsend for extensive assistance with data collection. This publication was made possible through the support of funding from the Intel corporation, and a grant from the John Templeton Foundation (Grant ID #57876; K.A.N. and J.D.C.). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Contributions

A.M.B. and M.A. conceived experiment; A.M.B., M.A., and S.F.F. designed experiment and analyses, with input from N.B.T., K.A.N. and J.D.C.. A.M.B. and M.A. wrote the experiment code; A.M.B. and M.A. ran the experiment; A.M.B. and S.F.F. contributed analytic tools; A.M.B. and S.F.F. performed analyses; A.M.B. wrote the paper, with input from M.A., N.T.B., K.A.N., and J.D.C.. All authors approved the final manuscript.

Methods

Participants

33 participants (15 male, 30 right-handed; ages 18-50, mean 21.9) each performed two repetitions of the task in Experiment 1. Ten blocks were excluded for failing to meet one or more criteria: if the participant failed to respond on 10% of learn or test-phase trials (nine blocks); if the combined number of skipped trials and post-stimulus error trials during the test phase were greater than 30% (four blocks); if the difference between calibrated accuracies for any pair of stimuli was less than 5% (one block). Three participants failed to meet criteria for all blocks they performed; they were excluded entirely from analysis. In all, 30 participants and 56 blocks were included in the final analysis.

36 participants (10 male, 29 right-handed; ages 18-33, mean 23.19) each performed one (5) or two (31) repetitions of the task in Experiment 2. (Five blocks were excluded due to scanner malfunction (1), participant discomfort (1), or programming error (3).) 15 blocks were excluded for failing to meet one or more criteria: nine for failing to respond on enough learn or test-phase trials; one for failing to respond correctly or at all on enough test-phase trials; nine for failing the calibration accuracy threshold. Five participants failed to meet criteria for all blocks they performed; they were excluded entirely from analysis. In all, 31 participants and 52 blocks were included in the final analysis.

In Experiment 1, participants were compensated with course credit. In Experiment 2, participants were paid a flat fee of \$50. All participants reported themselves as free of neurological or psychiatric disease, and fully consented to participate. The study protocol was approved by the Institutional Review Board for Human Subjects at Princeton University.

Task

The experiment was controlled by a script written in Matlab (Mathworks, Natick, MA, USA), using the Psychophysics Toolbox [34]. Both Experiment 1 and Experiment 2 consisted of the following four phases, repeated for two blocks for each participant, with different stimuli and task conditions as detailed below. Experiment 2 was performed in an fMRI scanner, and consisted of an additional, fifth phase, a Localizer task described below.

In *Phase 1*, the *Response training* phase, participants learned to map response keys to stimuli. Four response keys – numbers one through four on a standard US keyboard – were

each associated with one of four stimuli – black and white photographs, two faces and two natural scenes.

Stimulus photographs were chosen from a set of four possible scenes and four possible faces. Each category was subdivided into two sets of two paired photographs. Each photograph was black and white, normalized for contrast and brightness, and chosen to be highly confusable with the paired face or scene.

Participants were first shown each photograph, centered on a black background, in order of the associated response keys, and asked to press the current key in the sequence. In all experiments, keys one and two corresponded to the faces, and keys three and four corresponded to scenes. Then, the photographs were shuffled, and presented one at a time for two seconds each. Participants were instructed to press the corresponding key. If they pressed the correct key, a green box appeared around the photograph. If they pressed the incorrect key, the photograph remained on the screen. Each photograph was displayed ten times. If participants pressed the incorrect key on the first try more than twice for any photograph, they were made to repeat the response training phase in its entirety.

Phase 2, the *Calibration* phase, measured the ability of participants to discriminate between each pair of photographs when they were presented in a noisy, “flickering” stream (Figure 1). On each trial, participants were shown a rapid stream of pictures, displayed for 1/60th of a second apiece. They were instructed to press the key corresponding to the *target* – the photograph shown most often. Each frame consisted of either the target photograph, the paired same-category photograph, or a perceptual mask consisting of a phase-scrambled version of the previously shown photograph. Perceptual masks were shown for between one and three frames, with mask display length chosen from a truncated, discretized exponential distribution of mean 2. Calibration trials lasted three seconds, regardless of response. When participants pressed a key, the stream stopped, and the target was shown for the remainder of the trial length. If the participant pressed the correct key, a green box appeared around the photograph. If the participant pressed the incorrect key, a red box appeared around the photograph. A one second inter-trial-interval (ITI) followed each trial. On each trial, the proportion of frames that contained the target photograph – the *coherence* – was updated according to a Quest algorithm [35], with the goal of calibrating participants responses to either 65% (low) or 85% (high) accuracy. Each block measured the coherence necessary to elicit either high or low accuracy for each photograph. In Experiment 1, the first 24 participants performed 60 calibration trials per photograph, while the last 9 participants performed 40 calibration trials per photograph. In Experiment 2, participants performed 30 calibration trials per photograph. Although Experiment 2 participants remained in the fMRI scanner for this phase, no scanner data was collected. This is the only phase for which scanner data was not collected.

In Experiment 1, for stimuli calibrated to low accuracy (65%), the average coherence (proportion of non-mask frames that contained the target photograph), across participants, blocks, and stimuli, was 60.98% (SEM 1.06%); whereas for the high-accuracy (85%) condition, the target photograph was shown on 75.88% (SEM 1.08%) of frames. In Experiment 2, these figures were 62.17% (SEM 1.03%) coherence in the low-accuracy condition, 77.66%

(SEM 1.15%) coherence in the high-accuracy condition.

Phase 3, the *Cue learning* phase, provided participants with a set of experiences that linked each of four fractal cues to the photographs (Figure 1). On each trial, participants were shown one of four fractal cues, displayed on the screen for 750ms. In Experiment 1, the cue was followed by a variable inter-stimulus-interval (ISI). For 24 participants, this ISI was either 500ms, 1000ms, or 4000 ms, selected pseudorandomly at each trial according to a uniform distribution. For the remaining 9 participants in Experiment 1, and all participants in Experiment 2, this ISI was a fixed length of one second. After the ISI, participants were shown either of two photographs linked to the cue, both from the same category (face or scene). The photographs that followed the cue were selected according to one of four binomial distributions – 50/50, 60/40, 70/30, or 80/20. The two cues in each category (face or scene) predicted their consequents using symmetric distributions – if one cue predicted Face A with 80% probability, the other cue predicted Face B with 80% probability. Participants were instructed to press the button corresponding to the displayed photograph. If the response was accurate, the photograph was surrounded by a green box. If the response was inaccurate, the photograph was surrounded by a red box. Regardless of response time or accuracy, the picture remained on the screen for two seconds. In Experiment 1, the trial was followed by an ITI of two seconds. In Experiment 2, the trial was followed by an ITI of between 500ms and 8000ms, chosen from a truncated exponential distribution, discretized in units of 500ms, with mean 2000ms. This phase consisted of 100 trials, 25 for each cue, ordered pseudorandomly.

Phase 4, the *Cued inference task*, was the primary test of our hypotheses. On each trial during this phase, participants first viewed a fractal cue that predicted the likelihood of the target photograph during the following flickering stream. Cues were presented for 750ms, and followed by an ISI of variable length, selected at each trial from a uniform distribution. For the first 24 participants of Experiment 1, this ISI was either 500ms, 1000ms, or 4000ms. For the remaining 9 participants of Experiment 1, this ISI was either six, eight, or ten seconds. In Experiment 2, this ISI was either four, six, or eight seconds. In both experiments, ISI durations were chosen from a uniform distribution over the possible values. The flickering stream used one of the two mixture proportions calibrated during Phase 2; mixture proportions were fixed for each category – e.g. faces might be set to low coherence, and scenes to high coherence. Thus, the fractal cue predicted both the likely identity of the target photograph, and also the coherence of the subsequent stream. The stream remained on the screen for three seconds. When a key was pressed, the target photograph appeared, and remained on the screen until the three seconds were finished. If the keypress was correct, the photograph was surrounded by a green box. If the keypress was incorrect, the photograph was surrounded by a red box. Participants were instructed to press the key corresponding to the identity of the target photograph. Critically, however, participants were allowed to respond early – during the ISI, before the flickering stream began. Participants were not given any explicit or implicit inducement to respond early or accurately – they were informed that, regardless of the speed or correctness of their response, all trials were of fixed length, modulo the ISI. This phase continued for 80 trials, 20 trials of each cue, ordered pseudorandomly.

Phases one through four were repeated as two blocks, each with different fractal cues and picture stimuli. Cue were selected pseudorandomly for each block, and the mapping from coherence level to category was counterbalanced between blocks.

After the two blocks, Experiment 2 participants completed a final phase, *Phase 5*, the *Localizer* task. We used the data collected in this phase to localize regions of cortex preferentially active during processing of face and scene images. Participants performed a 1-back image repeat detection task. Images were presented in mini-blocks of 10 trials each. Eight of the pictures in each block were trial-unique, and two were repeats of the picture on the immediately preceding trial. Repeats were inserted pseudorandomly according to a uniform distribution. Stimuli in each mini-block were chosen from a large stimulus set of pictures not used in the main experiment. The pictures belonged to one of four categories – faces, objects, scenes or phase-scrambled scenes. Pictures were each presented for 500ms, and separated by a 1.3s ISI. A total of 12 mini-blocks were presented (3 per category), with each mini-block separated by a 12 second inter-block interval.

Imaging methods

Experiment 2 was collected while participants were laying in the fMRI scanner. Data were acquired using a 3T Siemens Prisma scanner with a 64-channel volume head coil. We collected three functional runs with a T2*-weighted gradient-echo multi-band echo-planar sequence (44 slices oriented parallel to the long axis of the hippocampus, 2.5mm isotropic resolution, echo time 26 ms; TR 1000 ms; flip angle 50 deg; field of view 192 mm). To allow for T1 equilibration, we discarded the first six volumes of each functional run (6s). We also collected a high-resolution 3D T1-weighted MPRAGE sequence (1mm isotropic resolution) for registration across participants to standard space. Functional image preprocessing was performed using FSL (FMRIB Software Library version 5.0.8; [36, 37]). Anatomical images were coregistered to the standard MNI152 template image, then individual participant functional images were coregistered to the realigned anatomical images. The transformation matrices generated during this coregistration process were used to transform Region of Interest (ROI) images (described below, *ROI definition*). Functional images were motion corrected and spatially smoothed using a 5mm full-width half-maximum Gaussian kernel prior to analysis. Data were scaled to their global mean intensity and high-pass filtered with a cutoff period of 128s.

Behavioral analysis

Response time analyses

Bimodality. We tested whether response time distributions within each ISI condition were bimodal, using Hartington's Dip Test [15]. This test measures the relative spread between modes to the mean of the distribution – larger values indicate a higher likelihood of true bimodality in the tested data. P-values are estimated via bootstrap against distributions with the same summary statistics as the tested data, provided by the MATLAB function

HARTIGANSDIPTEST [38].

Permutation tests for across-condition correlations. Each participant performed a different subset of the task conditions (cue level, perceptual coherence). To provide a robust measure of the relationship between response times and conditions, we therefore performed a bootstrap analysis, across participants and conditions [39]. On each iteration, we sampled, with replacement, the number of participants in the study group (33 in Experiment 1, 33 in Experiment 2). We then computed, on this selected group, the correlation of interest. By repeating the process 10,000 times, we obtained a distribution of correlation values across shuffled permutations of the study group. The reported p-value is thus the fraction of correlation values with a different sign from the base effect size (the correlation across the entire original group). We also leveraged these shuffled permutations to measure the difference between correlations across conditions (coherence levels, early versus late responses). To do this, we performed the paired bootstrap analyses using the same set of 10,000 permutations – each permutation of participants was matched between the two correlations. We then measured the effect size using Cohen’s d [40]; by convention effect sizes measured in this way greater than 0.80 are “Large”, and likely reliable.

Model comparison

Multi-stage DDM. Our primary model of interest is an extension of the drift-diffusion model [3] to allow for a time-varying drift rate [14]. The model specifies drift rate as a time-varying, piecewise constant function – each shift in drift rate defines a separate stage of the accumulation process. Critically, the endpoint of one stage leads to the starting point of the next. Our instantiation used two stages. The free parameters were the drift rates, d_1 and d_2 , threshold a , non-decision time T_0 , and starting point x_0 . We refer to this model as *MSDDM*.

Our comparison models were matched to this MS-DDM in parameters, with selective disabling of each key feature – the time-varying drift rate, and the connection between stages. The first comparison model of interest was a single DDM, with continuous accumulation until the time of response, but no change in drift rate across the entire period between cue onset and response. We refer to this model as *1DDM*, with free parameters d_1 , a , T_0 , and x_0 . The second comparison model of interest was two DDMs, each fit to pre-stimulus and post-stimulus responses separately and thus mirroring the change in drift rate found in MS-DDM, but with the second starting point its own free parameter. We refer to this model as *2DDM*, with free parameters d_1 , d_2 , a , T_0 , and starting points $x_{0,1}$ and $x_{0,2}$. Each model was fit to participant responses aggregated according to cue, coherence, and ISI condition. The fitting procedure minimized the difference between the χ^2 of the distribution of RTs in each cue-coherence-ISI bin and the RT distribution generated by the chosen model at the given parameters. Fitting was performed using a genetic algorithm (MATLAB *ga*) that ran for 1,000 generations per parameter, at a population size of 50 per parameter. The summed residuals across conditions were used to calculate likelihood.

Imaging analysis

To identify neural markers of stimulus reinstatement, we first defined patterns of activity in ventral visual stream regions that indicated participants were processing “face” or “scene” photographs. We then analyzed the degree to which these patterns were present during the post-cue, pre-stimulus ISIs in Phase 4. Because no pictures were present on-screen during this period of interest, we reasoned that greater evidence of stimulus reinstatement would indicate that participants were recalling the cued photograph. We therefore predicted that this reinstatement evidence would be reflected in response accuracies, response times, and DDM model parameters.

ROI definition. We identified a region of interest consisting of voxels that (across the group) showed preferential activation to face or scene photographs, using the following procedure.

First, for each participant, we performed a GLM analysis of BOLD signal during the localizer task. We identified voxels that responded more to scenes or faces, relative to other categories (univariate contrasts: faces > scenes | scrambled_scenes | objects; scenes > faces | scrambled_scenes | objects). For each participant, we selected clusters in the posterior parahippocampal region (matching the reported Parahippocampal Place Area (PPA); [19]) and posterior fusiform gyrus (matching the reported Fusiform Face Area (FFA); [18]) that were significant at $p < 0.005$, uncorrected. Next, each per-participant voxel mask was binarized; all above-threshold voxels were set to 1. To regularize the ROIs and ensure they were consistent across participants, the resulting individual mask was then warped to match the group average anatomical; these group-space masks were added together and the summed image thresholded to include all voxels present in more than 90% of participants. This final group ROI was then warped back to the individual participant space, and the result used as the final mask for pattern analysis.

Stimulus-specific pattern analysis. We computed the pattern of activity for each target photograph, across the corresponding category-preferring ROI. For each photograph in each block, we took the average pattern of activity over the last five presentations of the photograph during Phase 1. (The first five presentations were excluded to allow repetition suppression and learning effects to stabilize.)

We next used these four three-dimensional patterns as a template for analyzing activity during the post-cue, pre-stimulus ISI in Phase 4. For each trial, we computed, within the ROI corresponding to the cued category, the pattern of activity between the time of cue onset and the time of response or one TR before the onset of the flickering stream, whichever came first. We then correlated this three-dimensional activity pattern with the corresponding target pattern, defined above. These correlation values, one for each Phase 4 trial, were then fisher-transformed and used as predictor variables in our analyses of interest. We refer to these values as the *reinstatement index*.

References

- [1] Antonio Rangel, Colin Camerer, and P Read Montague. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7):545–56, jul 2008. ISSN 1471-0048. doi: 10.1038/nrn2357. URL <http://www.ncbi.nlm.nih.gov/pubmed/18545266>.
- [2] Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–74, jan 2007. ISSN 0147-006X. doi: 10.1146/annurev.neuro.29.051605.113038. URL <http://www.ncbi.nlm.nih.gov/pubmed/17600525>.
- [3] Roger Ratcliff. A Theory of Memory Retrieval. *Psychological Review*, 85(2):59–108, 1978.
- [4] Alan M Gordon, Jesse Rissman, Roozbeh Kiani, and Anthony D Wagner. Cortical Reinstatement Mediates the Relationship Between Content-Specific Encoding Activity and Subsequent Recollection Decisions. *Cerebral Cortex*, 24(12):3350–3364, aug 2013. ISSN 1460-2199. doi: 10.1093/cercor/bht194. URL <http://www.ncbi.nlm.nih.gov/pubmed/23921785>.
- [5] Don van Ravenzwaaij, Martijn J Mulder, Francis Tuerlinckx, and Eric-Jan Wagenmakers. Do the dynamics of prior information depend on task context? An analysis of optimal performance and an empirical test. *Frontiers in psychology*, 3(May):132, jan 2012. ISSN 1664-1078. doi: 10.3389/fpsyg.2012.00132.
- [6] Jan Drugowitsch, Rubén Moreno-Bote, Anne K Churchland, Michael N Shadlen, and Alexandre Pouget. The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11):3612–28, mar 2012. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.4010-11.2012. URL <http://www.ncbi.nlm.nih.gov/pubmed/22423085>.
- [7] Rani Moran. Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, pages 38–53, 2015. doi: 10.3758/s13423-014-0669-3.
- [8] Don Van Ravenzwaaij, Martijn J Mulder, Francis Tuerlinckx, and Eric Jan Wagenmakers. Paradoxes of Optimal Decision Making : A Response to Moran (2015). *Psychonomic Bulletin & Review*, 22:307–308, 2015.
- [9] Timothy D. Hanks, Mark E. Mazurek, Roozbeh Kiani, Elizabeth Hopp, N Michael, and M. N. Shadlen. Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task. *Journal of Neuroscience*, 31(17):6339–6352, apr 2011. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.5613-10.2011. URL <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.5613-10.2011>.

- [10] Aaron M. Bornstein and Nathaniel D. Daw. Cortical and Hippocampal Correlates of Deliberation During Model-Based Decisions for Rewards in Humans. *PLoS Computational Biology*, 9(12):e1003387, dec 2013. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003387. URL <http://dx.plos.org/10.1371/journal.pcbi.1003387>.
- [11] Michael N Shadlen and Daphna Shohamy. Decision Making and Sequential Sampling from Memory. *Neuron*, 90(5):927–939, 2016. ISSN 0896-6273. doi: 10.1016/j.neuron.2016.04.036. URL <http://dx.doi.org/10.1016/j.neuron.2016.04.036>.
- [12] Aaron M Bornstein, Mel W Khaw, Daphna Shohamy, and Nathaniel D. Daw. Reminders of past choices bias decisions for reward in humans. *Nature Communications*, 8:15958, 2017. doi: 10.1038/ncomms15958.
- [13] Aaron M Bornstein and Kenneth A Norman. Reinstated episodic context guides sampling-based decisions for reward. *Nature Neuroscience*, 2017. doi: 10.1038/nn.4573.
- [14] Vaibhav Srivastava, Samuel F. Feng, Jonathan D. Cohen, Naomi Ehrich Leonard, and Amitai Shenhav. A martingale analysis of first passage times of time-dependent Wiener diffusion models. *Journal of Mathematical Psychology*, 2016. ISSN 00222496. doi: 10.1016/j.jmp.2016.10.001. URL <http://dx.doi.org/10.1016/j.jmp.2016.10.001>.
- [15] J. a. Hartigan and P. M. Hartigan. The Dip Test of Unimodality. *The Annals of Statistics*, 13(1):70–84, 1985. ISSN 0090-5364. doi: 10.1214/aos/1176346577.
- [16] Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700–65, oct 2006. ISSN 0033-295X. doi: 10.1037/0033-295X.113.4.700. URL <http://www.ncbi.nlm.nih.gov/pubmed/17014301>.
- [17] R E Kass and A E Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995.
- [18] Nancy Kanwisher, Josh Mcdermott, and Marvin M Chun. The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17(11):4302–4311, 1997.
- [19] R Epstein and N Kanwisher. A cortical representation of the local visual environment. *Nature*, 392(6676):598–601, apr 1998. ISSN 0028-0836. doi: 10.1038/33402. URL <http://www.ncbi.nlm.nih.gov/pubmed/9560155>.
- [20] Zeb Kurth-Nelson, Gareth Barnes, Dino Sejdinovic, Ray Dolan, and Peter Dayan. Temporal structure in associative retrieval. *eLife*, 4:1–18, 2015. ISSN 2050-084X. doi: 10.7554/eLife.04919. URL <http://elifesciences.org/lookup/doi/10.7554/eLife.04919>.

- [21] Rafal Bogacz, Peter T Hu, Philip J Holmes, and Jonathan D Cohen. Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology*, 63(5):863–91, 2010. ISSN 1747-0226. doi: 10.1080/17470210903091643.
- [22] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, London, 2006.
- [23] Roger Ratcliff. Parameter variability and distributional assumptions in the diffusion model. *Psychological review*, 120(1):281–92, jan 2013. ISSN 1939-1471. doi: 10.1037/a0030775. URL <http://www.ncbi.nlm.nih.gov/pubmed/23148742>.
- [24] Aaron M. Bornstein and Nathaniel D. Daw. Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience*, 35(7):1011–1023, apr 2012. ISSN 0953816X. doi: 10.1111/j.1460-9568.2011.07920.x. URL <http://doi.wiley.com/10.1111/j.1460-9568.2011.07920.x>.
- [25] Michael L Mack and Alison R Preston. Decisions about the past are guided by reinstatement of specific memories in the hippocampus and perirhinal cortex. *NeuroImage*, 127:144–157, 2016. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2015.12.015. URL <http://dx.doi.org/10.1016/j.neuroimage.2015.12.015>.
- [26] Scott D. Brown and Andrew Heathcote. The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–178, 2008. ISSN 00100285. doi: 10.1016/j.cogpsych.2007.12.002.
- [27] Roger Ratcliff and Jeffrey J Starns. Modeling confidence judgments, response times, and multiple choices in decision making: recognition memory and motion discrimination. *Psychological review*, 120(3):697–719, jul 2013. ISSN 1939-1471. doi: 10.1037/a0033152. URL <http://www.ncbi.nlm.nih.gov/pubmed/23915088>.
- [28] David J Heeger. Theory of cortical function. *Proceedings of the National Academy of Sciences*, 114(8), 2017. ISSN 0027-8424. doi: 10.1073/pnas.1619788114.
- [29] Rajesh P. N. Rao. Decision Making Under Uncertainty: A Neural Model Based on Partially Observable Markov Decision Processes. *Frontiers in Computational Neuroscience*, 4:1–18, 2010. ISSN 1662-5188. doi: 10.3389/fncom.2010.00146. URL http://www.frontiersin.org/Computational_Neuroscience/10.3389/fncom.2010.00146/abstract.
- [30] Abigail N. Hoskin, Aaron M. Bornstein, Kenneth A. Norman, and Jonathan D. Cohen. Resting reinstatements from episodic memory alter the content of working memory. *bioRxiv*, 2017.
- [31] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention, and choice. 2017.
- [32] Xavier Gabaix and David Laibson. Myopia and discounting. 2017.

- [33] Ian Krajbich, Carrie Armel, and Antonio Rangel. Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10):1292–8, oct 2010. ISSN 1546-1726. doi: 10.1038/nn.2635. URL <http://www.ncbi.nlm.nih.gov/pubmed/20835253>.
- [34] D H Brainard. The Psychophysics Toolbox. *Spatial Vision*, 10(4):433–6, jan 1997. ISSN 0169-1015. URL <http://www.ncbi.nlm.nih.gov/pubmed/9176952>.
- [35] A. B. Watson and Dennis G. Pelli. QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33(2):113–120, 1983.
- [36] Stephen M. Smith, Mark Jenkinson, Mark W. Woolrich, Christian F. Beckmann, Timothy E J Behrens, Heidi Johansen-Berg, Peter R. Bannister, Marilena De Luca, Ivana Drobnjak, David E. Flitney, Rami K. Niazy, James Saunders, John Vickers, Yongyue Zhang, Nicola De Stefano, J. Michael Brady, and Paul M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23 (SUPPL. 1):208–219, 2004. ISSN 10538119. doi: 10.1016/j.neuroimage.2004.07.051.
- [37] Mark Jenkinson, Christian F. Beckmann, Timothy E J Behrens, Mark W. Woolrich, and Stephen M. Smith. FSL. *NeuroImage*, 62(2):782–790, 2012. ISSN 10538119. doi: 10.1016/j.neuroimage.2011.09.015.
- [38] Nicholas Price and F. Mechler. HartigansDipTest, 2002. URL <http://www.nicprice.net/diptest/>.
- [39] Ghootae Kim, Jarrod A Lewis-peacock, Kenneth A Norman, and Nicholas B Turk-browne. Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences*, 111(24), 2014. doi: 10.1073/pnas.1319438111.
- [40] Jacob Cohen. *Statistical power analysis for the behavioural sciences*. NJ: Lawrence Earlbaum Associates, Hillside, NJ, 1988.