



Learning optimal decisions with confidence

Jan Drugowitsch^{a,1}, André G. Mendonça^b, Zachary F. Mainen^b, and Alexandre Pouget^c

^aDepartment of Neurobiology, Harvard Medical School, Boston, MA 02115; ^bChampalimaud Research, Champalimaud Centre for the Unknown, 1400-038 Lisbon, Portugal; and ^cDepartment of Basic Neuroscience, University of Geneva, CH-1211 Geneva, Switzerland

Edited by Paul W. Glimcher, New York University, New York, NY, and accepted by Editorial Board Member Thomas D. Albright October 18, 2019 (received for review April 19, 2019)

Diffusion decision models (DDMs) are immensely successful models for decision making under uncertainty and time pressure. In the context of perceptual decision making, these models typically start with two input units, organized in a neuron–antineuron pair. In contrast, in the brain, sensory inputs are encoded through the activity of large neuronal populations. Moreover, while DDMs are wired by hand, the nervous system must learn the weights of the network through trial and error. There is currently no normative theory of learning in DDMs and therefore no theory of how decision makers could learn to make optimal decisions in this context. Here, we derive such a rule for learning a near-optimal linear combination of DDM inputs based on trial-by-trial feedback. The rule is Bayesian in the sense that it learns not only the mean of the weights but also the uncertainty around this mean in the form of a covariance matrix. In this rule, the rate of learning is proportional (respectively, inversely proportional) to confidence for incorrect (respectively, correct) decisions. Furthermore, we show that, in volatile environments, the rule predicts a bias toward repeating the same choice after correct decisions, with a bias strength that is modulated by the previous choice's difficulty. Finally, we extend our learning rule to cases for which one of the choices is more likely a priori, which provides insights into how such biases modulate the mechanisms leading to optimal decisions in diffusion models.

decision making | diffusion models | optimality | confidence

Decisions are a ubiquitous component of everyday behavior. To be efficient, they require handling the uncertainty arising from the noisy and ambiguous information that the environment provides (1). This is reflected in the trade-off between speed and accuracy of decisions. Fast choices rely on little information and may therefore sacrifice accuracy. In contrast, slow choices provide more opportunity to accumulate evidence and thus may be more likely to be correct, but are more costly in terms of attention or effort and lost time and opportunity. Therefore, efficient decisions require not only a mechanism to accumulate evidence but also one to trigger a choice once enough evidence has been collected. Drift-diffusion models (or diffusion decision models) (DDMs) are a widely used model family (2) that provides both mechanisms. Not only do DDMs yield surprisingly good fits to human and animal behavior (3–5), but they are also known to achieve a Bayes-optimal decision strategy under a wide range of circumstances (4, 6–10).

DDMs assume a particle that drifts and diffuses until it reaches one of two boundaries, each triggering a different choice (Fig. 1A). The particle's drift reflects the net surplus of evidence toward one of two choices. This is exemplified by the random-dot motion task, in which the motion direction and coherence set the drift sign and magnitude. The particle's stochastic diffusion reflects the uncertainty in the momentary evidence and is responsible for the variability in decision times and choices widely observed in human and animal decisions (3, 5). A standard assumption underlying DDMs is that the noisy momentary evidence that is accumulated over time is one-dimensional—an abstraction of the momentary decision-related evidence of some stimulus. In reality, however, evidence would usually be distributed across a larger number of inputs, such as a neural population in the brain, rather than individual neurons (or neuron/antineuron pairs; Fig. 1A). Furthermore,

the brain would not know a priori how this distributed encoding provides information about the correctness of either choice. As a consequence, it needs to learn how to interpret neural population activity from the success and failure of previous choices. How such an interpretation can be efficiently learned over time, both normatively and mechanistically, is the focus of this work.

The multiple existing computational models for how humans and animals might learn to improve their decisions from feedback (e.g., refs. 11–14) do not address the question we are asking, as they all assume that all evidence for each choice is provided at once, without considering the temporal aspect of evidence accumulation. This is akin to fixed-duration experiments, in which the evidence accumulation time is determined by the environment rather than the decision maker. We, instead, address a more general and natural case in which decision times are under the decision maker's control. In this setting, commonly studied using “reaction time” paradigms, the temporal accumulation of evidence needs to be treated explicitly, and—as we will show—the time it took to accumulate this evidence impacts how the decision strategy is updated after feedback. Some models for both choice and reaction times have addressed the presence of high-dimensional inputs (e.g., refs. 15–17). However, they usually assumed as many choices as inputs, were mechanistic rather than normative, and did not consider how interpreting the input could be learned. We furthermore extend on previous work by considering the effect of a priori biases toward believing that one option is more correct than the other, and how such biases can be learned. This yields a theoretical understanding of how choice biases impact optimal decision making in diffusion models. Furthermore, it clarifies of how different implementations of this bias result in different diffusion model implementations, like the one proposed by Hanks et al. (18).

Significance

Popular models for the trade-off between speed and accuracy of everyday decisions usually assume fixed, low-dimensional sensory inputs. In contrast, in the brain, these inputs are distributed across larger populations of neurons, and their interpretation needs to be learned from feedback. We ask how such learning could occur and demonstrate that efficient learning is significantly modulated by decision confidence. This modulation predicts a particular dependency pattern between consecutive choices and provides insight into how a priori biases for particular choices modulate the mechanisms leading to efficient decisions in these models.

Author contributions: J.D., A.G.M., Z.F.M., and A.P. designed research; J.D. and A.P. performed research; J.D. analyzed data; and J.D., A.G.M., Z.F.M., and A.P. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. P.W.G. is a guest editor invited by the Editorial Board.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: jan_drugowitsch@hms.harvard.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1906787116/-DCSupplemental>.

First published November 15, 2019.

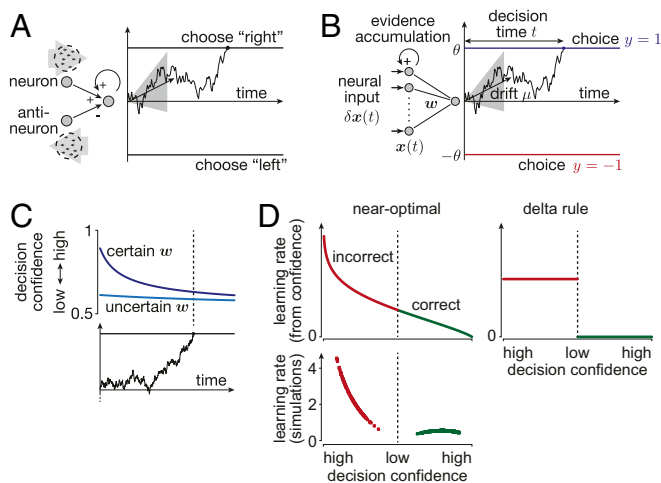


Fig. 1. Learning the input weights from feedback in diffusion models. In diffusion models, the input(s) provide at each point in time noisy evidence about the world's true state, here given by the drift μ . The decision maker accumulates this evidence over time (e.g., black example traces) to form a belief about μ . Bayes-optimal decisions choose according to the sign of the accumulated evidence, justifying the two decision boundaries that trigger opposing choices. (A) In standard diffusion models, the momentary evidence either arises directly from noisy samples of μ , or, as illustrated here, from a neuron/anti-neuron pair that codes for opposing directions of evidence. The illustrated example assumes a random-dot task, in which the decision maker needs to identify whether most of the dots that compose the stimulus are moving either to the left or to the right. The two neurons (or neural pools) are assumed to extract motion energy of this stimulus toward the right (*Top*) and left (*Bottom*), such that their difference forms the momentary evidence toward rightward motion. A decision is made once the accumulated momentary evidence reaches one of two decision boundaries, triggering opposing choices. (B) Our setup differs from that in A in that we assume the input information $\delta x(t)$ to be encoded in a larger neural population whose activity is linearly combined with weights w to yield the one-dimensional momentary evidence, and that the decision maker aims to learn these weights from feedback about the correctness of her choices. (C) Decision confidence (i.e., the belief that the made choice was correct) in this kind of diffusion model drops as a function of time (horizontal axis) and with increased uncertainty about the input weights (different shades of blue). (D) For near-optimal learning, the learning rate (the term ξ_w in Eq. 6) is modulated by decision confidence (*Top Left*). High-confidence decisions lead to little learning if correct (green, *Right*), and strong learning if incorrect (red, *Left*). Low-confidence decisions result in a moderate confidence-related learning rate term (*Top and Center*). The learning rate in 1,000 simulated trials (*Bottom*) shows that the overall learning rate preserves this trend, with an additional suppression of learning for low-confidence decisions. Other learning heuristics (e.g., the delta rule, *Right*) do not modulate their learning by confidence.

Results

Bayes-Optimal Decision Making with Diffusion Models. A standard way (8, 10, 19) to interpret diffusion models as mechanistic implementations of Bayes-optimal decision making is to assume that, in each trial, an unobservable latent state μ (called “drift rate” in diffusion models) is drawn from a prior distribution, $\mu \sim N(0, \sigma_\mu^2)$, with zero mean and variance σ_μ^2 . The decision maker's aim is to infer whether this latent state is positive or negative (e.g., rightward vs. leftward motion in the random-dot motion task), irrespective of its magnitude (e.g., the dot coherence level). The latent state itself is not directly observed but is indirectly conveyed via a stream of noisy, momentary evidence values $\delta z_1, \delta z_2, \dots$, that, in each small time step of size δt , provide independent and identically distributed noisy information about μ through $\delta z_i | \mu \sim N(\mu \delta t, \delta t)$. Here, we have chosen a unit variance, scaled by δt . Any rescaling of this variance by an additional parameter would result in a global rescaling of the evidence that

can be factored out (4, 8, 20), thus making such a rescaling unnecessary.

Having after some time $t \equiv n \delta t$ observed n pieces of such evidence, $\delta z_{1:n}$, the decision maker's posterior belief about μ , $p(\mu | \delta z_{1:n})$, turns out to be fully determined by the accumulated evidence, $z(t) = \sum_{i=1}^n \delta z_i$, and time t (*Materials and Methods*). Then, the posterior belief about μ being positive (e.g., leftward motion) results in the following (8):

$$p(\mu \geq 0 | z(t), t) = \int_0^\infty p(\mu | \delta z_{1:n}) d\mu = \Phi\left(\frac{z(t)}{\sqrt{t + \sigma_\mu^2}}\right), \quad [1]$$

where $\Phi(\cdot)$ is the cumulative function of a standard Gaussian. The opposite belief about μ being negative is simply $p(\mu < 0 | z(t), t) = 1 - p(\mu \geq 0 | z(t), t)$ (see Fig. 3A). The accumulated evidence follows a diffusion process, $z(t) | \mu \sim N(\mu t, t)$, and thus can be interpreted as the location of a drifting and diffusing particle with drift μ and unit diffusion variance (Fig. 1A). By Eq. 1, the posterior belief about $\mu \geq 0$ is $>1/2$ for positive $z(t)$, and $<1/2$ for negative $z(t)$. To make Bayes-optimal decisions, Bayesian decision theory (21) requires that these decisions are chosen to maximize the expected associated reward (or, more formally, to minimize the expected loss). Assuming equally rewarding correct choices, this implies choosing the option that is considered more likely correct. Given the above posterior belief, this makes $y = \text{sign}(z(t)) \in \{-1, 1\}$ the Bayes-optimal choice, which can be implemented mechanistically by (possibly time-varying) boundaries $\pm\theta(t)$ on $z(t)$, associated with the two choices. At these boundaries, the posterior belief about having made the correct choice, or decision confidence (22), is then given by Eq. 1 with $z(t)$ replaced by $\theta(t)$. The sufficient statistics, $z(t)$ and t , of this posterior remain unchanged by the introduction of such decision boundaries, such that Eq. 1 remains valid even in the presence of these boundaries (8). Thus, under the above assumptions of prior and evidence, diffusion models implement the Bayes-optimal decision strategy (see Fig. 3B).

Note that $|\mu|$ (i.e., the momentary evidence's signal-to-noise ratio) controls the amount of information provided about the sign of μ , and thus the difficulty of individual decisions. Thus, the used prior $\mu \sim N(0, \sigma_\mu^2)$, which has more mass on small $|\mu|$, reflects that the difficulty of decisions varies across trials and that harder decisions are more frequent than easier ones. The prior width, σ_μ^2 , determines the spread of μ values across trials, and therefore the overall difficulty of the task (larger σ_μ^2 = overall easier task). We chose a Gaussian prior for mathematical convenience, and also because hard trials are more frequent than easy ones in many experiments (e.g., ref. 20), even though they do not commonly use Gaussian priors. In general, the important assumption is that the difficulty varies across trials, but not exactly how it does so, which is to say that the shape of the prior distribution is not critical (8). Different prior choice will not qualitatively change our results but would make it hard or impossible to derive interpretable closed-form expressions. Model predictions would change qualitatively if we assume the difficulty to be fixed, or known a priori (8), but we will not consider this case, as it rarely if ever occurs in the real world.

Using High-Dimensional Diffusion Model Inputs. To extend diffusion models to multidimensional momentary evidence, we assume it to be given the k -dimensional vector δx_i . This evidence might represent inputs from multiple sensors, or the (abstract) activity of a neuronal population (Fig. 1B). As the activity of neurons in a population that encodes limited information about the latent state μ is likely correlated across neurons (23, 24), we chose the

momentary evidence statistics to also feature such correlations (*Materials and Methods*). In general, we choose these statistics such that $\mathbf{w}^T \delta \mathbf{x}_i = \delta z_i$, where the vector \mathbf{w} denotes the k input weights (for now assumed known). Defining the high-dimensional accumulated evidence by $\mathbf{x}(t) = \sum_{i=1}^n \delta \mathbf{x}_i$, this implies $z(t) = \mathbf{w}^T \mathbf{x}(t)$, such that it is

again Bayes-optimal to trigger decisions as soon as $\mathbf{w}^T \mathbf{x}(t)$ equals one of two decision boundaries $\pm \theta(t)$. Furthermore, the posterior belief about $\mu \geq 0$ is, similar to Eq. 1, given by the following:

$$p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t) = \Phi(\mathbf{w}^T \tilde{\mathbf{x}}(t)), \quad [2]$$

where we have defined the time-attenuated accumulated evidence $\tilde{\mathbf{x}}(t) = \mathbf{x}(t) / \sqrt{t + \sigma_\mu^{-2}}$. As a consequence, the decision-confidence for either choice, is, as before, given by Eq. 2, with $\mathbf{w}^T \tilde{\mathbf{x}}(t)$ replaced by $\theta(t) / \sqrt{t + \sigma_\mu^{-2}}$. For time-independent decision bounds, $\theta(t) = \theta$, this confidence decreases over time (Fig. 1C), reflecting the uncertainty about μ , and that late choices are likely due to a low μ , which is associated with a hard trial, and thus low decision confidence. This counterintuitive drop in confidence with time has been previously described for diffusion models with one-dimensional inputs (8, 25) and is a consequence of a trial difficulty that varies across trials. Specifically, it arises from a mixture of easy trials associated with large $|\mu|$ that lead to rapid, high-confidence choices, and hard trials associated with small $|\mu|$ that lead to slow, low-confidence choices. Therefore, it does not depend on our choice of Gaussian prior but is present for any choice of symmetric prior over μ (*SI Appendix*). The confidence remains constant over time only when the difficulty is fixed across trials (i.e., $\mu \in \{-\mu_0, \mu_0\}$ for some fixed μ_0).

Using Feedback to Find the Posterior Weights. So far, we have assumed the decision maker knows the linear input weights \mathbf{w} to make Bayes-optimal choices. If they were not known, how could they be learned? Traditionally, learning has been considered an optimization problem, in which the decision maker tunes some decision-making parameters (here, the input weights \mathbf{w}) to maximize their performance. Here, we will instead consider it as an inference problem in which the decision maker aims to identify the decision-making parameters that are most compatible with the provided observations. These two views are not necessarily incompatible. For example, minimizing the mean squared error of a linear model (an optimization problem) yields the same solution as sequential Bayesian linear regression (an inference problem) (26). In fact, as we show in *SI Appendix*, our learning problem can also be formulated as an optimization problem. Nonetheless, we here take the learning-by-inference route, as it provides a statistical interpretation of the involved quantities, which provides additional insights. Specifically, we focus on learning the weights while keeping the diffusion model boundaries fixed. The decision maker's reward rate (i.e., average number of correct choices per unit time), which we use as our performance measure, depends on both weights and the chosen decision boundaries. However, to isolate the problem of weight learning, we fix the boundaries such that a particular set of optimal weights \mathbf{w}^* maximize this reward rate. The aim of weight learning is to find these weights. Weight learning is a problem that needs to be solved even if the decision boundaries are optimized at the same time. We have addressed how to best tune these boundaries elsewhere (8, 27).

To see how learning can be treated as inference, consider the following scenario. Before having observed any evidence, the decision maker has some belief, $p(\mathbf{w})$, about the input weights, either as a prior or formed through previous experience. They now observe new evidence, $\delta \mathbf{x}_1, \delta \mathbf{x}_2, \dots$ and use the mean of the belief over weights, $\langle \mathbf{w} \rangle$ (or any other statistics), to combine this

evidence and to trigger a choice y once the combined evidence reaches one of the decision boundaries. Upon this choice, they receive feedback y^* about which choice was the correct one. Then, the best way to update the belief about \mathbf{w} in light of this feedback is by Bayes' rule:

$$p(\mathbf{w} | \mathbf{x}(t), t, y^*) \propto p(y^* | \mathbf{w}, \mathbf{x}(t), t) p(\mathbf{w}), \quad [3]$$

where we have replaced the stream of evidence $\delta \mathbf{x}_1, \delta \mathbf{x}_2, \dots$ by the previously established sufficient statistics $\mathbf{x}(t)$ and t .

The likelihood $p(y^* | \mathbf{w}, \mathbf{x}(t), t)$ expresses for any hypothetical weight vector \mathbf{w} the probability that the observed evidence makes y^* the correct choice. To find its functional form, consider that, for a known weight vector, we have shown that $p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t)$, given by Eq. 2, expresses the probability that $y = 1$ (associated with $\mu \geq 0$) is the correct choice. Therefore, $1 - p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t)$ corresponds to the probability that $y = -1$ (associated with $\mu < 0$) is the correct choice. Therefore, it can act as the above likelihood function, which, by Eq. 2, is given by $p(y^* | \mathbf{w}, \mathbf{x}(t), t) = \Phi(y^* \mathbf{w}^T \tilde{\mathbf{x}}(t))$, where we have used $1 - \Phi(a) = \Phi(-a)$. In summary, the decision maker's belief is optimally updated after each choice by the following:

$$p(\mathbf{w} | \mathbf{x}(t), t, y^*) \propto \Phi(y^* \mathbf{w}^T \tilde{\mathbf{x}}(t)) p(\mathbf{w}). \quad [4]$$

This update equation only requires knowing the accumulated evidence $\mathbf{x}(t)$, decision time t , and feedback y^* , but is independent of the chosen option y , and how the decision maker came to this choice. As a matter of fact, the decision maker could make random choices, irrespective of the accumulated evidence, and still learn \mathbf{w} according to the above update equation, as long as they keep track of $\mathbf{x}(t)$ and t , and acknowledge the feedback y^* . Therefore, learning and decision making are not necessarily coupled. Nonetheless, we assume for all simulations that decision makers perform decisions by using the mean estimate $\langle \mathbf{w} \rangle$, which is an intuitively sensible choice if the decision maker's aim is to maximize their reward rate (*SI Appendix*).

As in Eq. 4, the likelihood parameters, \mathbf{w} , are linear within a cumulative Gaussian function, such problems are known as "probit regression" and do not have a closed-form expression for the posterior. We could proceed by sampling from the posterior by Markov chain Monte Carlo methods, but that would not provide much insight into the different factors that modulate learning the posterior weights. Instead, we proceed by deriving a closed-form approximation to this posterior to provide such insight, as well as a potential mechanistic implementation.

Confidence Controls the Learning Rate. To find an approximation to the posterior in Eq. 4, let us assume the prior to be given by the Gaussian distribution, $p(\mathbf{w}) = N(\mathbf{w} | \boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$, with mean $\boldsymbol{\mu}_w$ and covariance $\boldsymbol{\Sigma}_w$, which is the maximum entropy distribution that specifies the mean and covariance (28). First, we investigated how knowing \mathbf{w} with limited certainty, as specified by $\boldsymbol{\Sigma}_w$, impacts the decision confidence. Marginalizing over all possible \mathbf{w} values (*Materials and Methods*) resulted in the choice confidence to be given by the following:

$$p(y | \mathbf{x}(t), t) = \Phi\left(\frac{y \boldsymbol{\mu}_w^T \tilde{\mathbf{x}}}{\sqrt{1 + \tilde{\mathbf{x}}^T \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}}}\right). \quad [5]$$

Compared to Eq. 2, the choice confidence is additionally attenuated by $\boldsymbol{\Sigma}_w$. Specifically, higher weight uncertainty (i.e., an overall larger covariance $\boldsymbol{\Sigma}_w$) results in a lower decision confidence, as one would intuitively expect (Fig. 1C).

Next, we found a closed-form approximation to the posterior (Eq. 4). For repeated learning across consecutive decisions, the posterior over the weights after the previous decision becomes the prior for the new decision. Unfortunately, a direct application of this principle would lead to a posterior that changes its functional form after each update, making it intractable. We instead used assumed density filtering (ADF) (26, 29) that posits a fixed functional form $q(\mathbf{w}|y^*, \mathbf{x}(t), t) = \mathcal{N}(\mathbf{w}|\boldsymbol{\mu}_w^*, \boldsymbol{\Sigma}_w^*)$ of the posterior density—in our case, Gaussian for consistency with the prior—and then finds the posterior parameters $\boldsymbol{\mu}_w^*$ and $\boldsymbol{\Sigma}_w^*$ that make this approximate posterior best match the “true” posterior $p(\mathbf{w}|y^*, \mathbf{x}(t), t)$ (Eq. 4). Performing this match by minimizing the Kullback–Leiber divergence $\text{KL}(p||q)$ results in the posterior mean (30, 31):

$$\boldsymbol{\mu}_w^* = \boldsymbol{\mu}_w + \frac{\xi_w}{\sqrt{1 + \tilde{\mathbf{x}}^T \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}}} y^* \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}, \quad [6]$$

and a similar expression for the posterior covariance (*Materials and Methods*). Choosing $\text{KL}(p||q)$ to measure the distance between p and q is to some degree arbitrary, but has beneficial properties, such as that it causes the first two moments of q to match those of p (*SI Appendix*). In Eq. 6, the factor ξ_w modulates how strongly this mean is updated toward $y^* \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}$, and turns out to be a monotonically decreasing function of decision confidence (Fig. 1*D, Top*; see *Materials and Methods* for mathematical expression). For incorrect choices, for which the decision confidence is $p(y^*|\tilde{\mathbf{x}}(t), t) < 1/2$, ξ_w is largest for choices made with high confidence, promoting significant weight adjustments. For low-confidence choices, it only promotes moderate adjustments, notably irrespective of whether the choice was correct or incorrect. High-confidence, correct choices yield a low ξ_w , and thus an intuitively minor strategy update. The update of the posterior covariance follows a similar confidence-weighted learning rate modulation (*SI Appendix, Fig. S1*, and *Materials and Methods*).

Decision confidence is not the only factor that impacts the learning rate in Eq. 6. For instance, $\tilde{\mathbf{x}}$ shrinks for longer, less confidence choices (because it is inversely proportional to time) and results in overall less learning. Less certain weights, associated with larger magnitudes of $\boldsymbol{\Sigma}_w$, have a similar effect. To investigate the overall impact of all of these factors combined on the learning rate, we simulated a long sequence of consecutive choices and plotted the learning rate for a random subset of these trials against the decision confidence (Fig. 1*D, Bottom*). This plot revealed a slight down-weighting of the learning rate for low-confidence choices when compared to ξ_w , but left the overall dependency on ξ_w otherwise unchanged.

Performance Comparison to Optimal Inference and to Simpler Heuristics. The intuitions provided by near-optimal ADF learning are only informative if its approximations do not cause a significant performance drop. We quantified this drop by comparing ADF performance to that of the Bayes-optimal rule, as found by Gibbs sampling (*Materials and Methods*). Gibbs sampling is biologically implausible as it requires a complete memory of inputs and feedbacks for past decisions and is intractable for longer decision sequences, but nonetheless provides an optimal baseline to compare against. We furthermore tested the performance of two additional approximations. One was an ADF variant that assumes a diagonal covariance matrix $\boldsymbol{\Sigma}_w$, yielding a local learning rule that could be implemented by the nervous system. This variant furthermore reduced the number of parameters from quadratic to linear in the size of \mathbf{w} . The second was a second-order Taylor expansion of the log-posterior, resulting in a learning rule similar to ADF, but with a lower impact of weight uncertainty on the learning rate (*Materials and Methods*).

Furthermore, we tested whether simpler learning heuristics can match ADF performance. We focused on three rules of increasing complexity. The delta rule, which can be considered a variant of temporal-difference learning, or reinforcement learning (32), updates its weight estimate after the n th decision by the following:

$$\mathbf{w}_{n+1} = \mathbf{w}_n + \frac{\alpha}{2\theta(0)} \left(y_n^* \theta(t) - \mathbf{x}_n(t)^T \mathbf{w}_n \right) \mathbf{x}_n(t), \quad [7]$$

where $y_n^* \in \{-1, 1\}$ is the feedback about the correct choice provided after this decision, and we have chosen to normalize the learning rate α by the initial bound height $\theta(0)$ to make it less sensitive to this chosen height. As decisions are triggered at one of the two boundaries, $\mathbf{x}_n(t)^T \mathbf{w}_n \in \{-\theta(t), \theta(t)\}$, the residual in brackets is zero for correct choices, and $\pm 2\theta(t)$ for incorrect choices. As a result, and in contrast to ADF, weight adjustments are only performed after incorrect choices, and with a fixed learning rate α rather than one modulated by confidence (Fig. 1*D, Right*). Our simulations revealed that the delta rule excessively and suboptimally decrease in the weight size $\|\mathbf{w}\|$ over time, leading to unrealistically long reaction times and equally unrealistic near-zero weights. To counteract this problem, we designed a normalized delta rule, that updates the weight estimates as the delta rule, but thereafter normalizes them by $\mathbf{w} \leftarrow \mathbf{w} \|\mathbf{w}^*\| / \|\mathbf{w}\|$ to ensure that its size matches that of the true weights \mathbf{w}^* . Access to these true weights, \mathbf{w}^* , makes it an omniscient learning rule that cannot be implemented by a decision maker in practice. Last, we tested a learning rule that performs stochastic gradient ascent on the feedback log-likelihood:

$$\mathbf{w}_{n+1} = \mathbf{w}_n + \alpha \nabla_{\mathbf{w}} \log p(y_n^* | \mathbf{w}_n, \mathbf{x}_n(t), t) = \mathbf{w}_n + \alpha y_n^* \xi_w \tilde{\mathbf{x}}_n(t). \quad [8]$$

This rule introduces decision confidence weighting through ξ_w , but differs from ADF in that it does not take the weight uncertainty ($\boldsymbol{\Sigma}_w$ in ADF) into account, and requires tuning of the learning rate parameter α .

We evaluated the performance of these learning rules by simulating weight learning across 1,000 consecutive decisions (called “trials”; see *Materials and Methods* for details) in a task in which use of the optimal weight vector maximizes the reward rate. This reward rate was the average reward for correct choices minus some small cost for accumulating evidence over the average time across consecutive trials and is a measure we would expect rational decision makers to optimize. For each learning rule, we found its reward rate relative to random behavior and optimal choices.

Fig. 2*A* shows this relative reward rate for all learning rules and different numbers of inputs. As can be seen, the performance of ADF and the other probabilistic learning rules is indistinguishable from Bayes-optimal weight learning for all tested numbers of inputs. Surprisingly, the ADF variant that ignores off-diagonal covariance entries even outperformed Bayes-optimal learning for a large number of inputs (Fig. 2*A*, yellow line for 50 inputs). The reason that a simpler learning rule could outperform the rule deemed optimal by Bayesian decision theory is that this simpler rule has less parameters and a simpler underlying model that was nonetheless good enough to learn the required weights. Learning fewer parameters with the same data resulted in initially better parameter estimates, and better associated performance. Conceptually, this is similar to a linear model outperforming a quadratic model when fitting a quadratic function if little data are available, and if the function is sufficiently close to linear (as illustrated in *SI Appendix, Fig. S2*). Once more data are available, the quadratic model will outperform the linear one. Similarly, the Bayes-optimal learning rule will outperform the simpler one once more feedback has been observed. In our simulation, however, this does not occur within the 1,000 simulated trials.

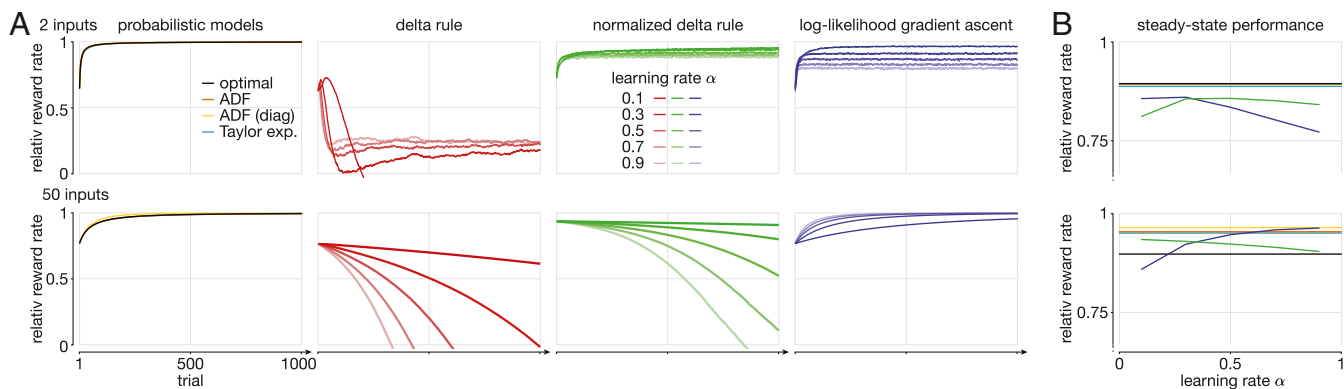


Fig. 2. Input weight learning and tracking performance of different learning rules. All plots show the relative reward rate (0 = immediate, random choices; 1 = optimal) averaged over 5,000 simulations with different true, underlying weights, and for 2 (*Top*) and 50 (*Bottom*) inputs. (A) The relative reward rate for probabilistic and heuristic learning rules. The probabilistic learning rules include the optimal rule (Gibbs sampling), assumed density filtering (ADF), ADF with a diagonal covariance matrix (ADF [diag]), and a learning rule based on a second-order Taylor expansion of the log-posterior (Taylor exp.). For both 2 and 50 inputs, all rules perform roughly equally. For the heuristic rules, different color shadings indicate different learning rates. The initial performance shown is that after the first application of the learning rule, such that initial performances can differ across learning rules. (B) The steady-state performance across different heuristic rule learning rates. Steady-state performance was measured as an average across 5,000 simulations, averaging over the last 100 of 1,000 simulated trials in which the true weights slowly change across consecutive trials. An optimal relative reward rate of 1 corresponds to knowing the true weight in each trial, which, due to the changing weight, is not achievable in this setup. The color scheme is the same as in A, but the vertical axis has a different scale. The delta rule did not converge and was not included in B.

All other learning heuristics performed significantly worse. For low-dimensional input, the delta rule initially improved its reward rate but worsens it again at a later stage across all learning rates. The normalized delta rule avoided such performance drops for low-dimensional input, but both delta rule variants were unable to cope with high-dimensional inputs. Only stochastic gradient ascent on the log-likelihood provided a stable learning heuristic for high-dimensional inputs, but with the downside of having to choose a learning rate. Small learning rates lead to slow learning, and an associated slower drop in angular error. Overall, the probabilistic learning rules significantly outperformed all tested heuristic learning rules and matched (and in one case even exceeded) the weight learning performance of the Bayes-optimal estimator.

Tracking Nonstationary Input Weights. So far, we have tested how well our weight learning rule is able to learn the true, underlying weights from binary feedback about the correctness of the decision maker's choices. For this, we assumed that the true weights remained constant across decisions. What would happen if these weights change slowly over time? Such a scenario could occur if, for example, the world around us changes slowly, or if the neural representation of this world changes slowly through neural plasticity or similar. In this case, the true weights would become a moving target that we would never be able to learn perfectly. Instead, we would after some initial transient expect to reach steady-state performance that remains roughly constant across consecutive decisions. We compared this steady-state performance of Bayes-optimal learning (now implemented by a particle filter) to that of the probabilistic and heuristic learning rules introduced in the previous section. The probabilistic rules were updated to take into account such a trial-by-trial weight change, as modeled by a first-order autoregressive process (*Materials and Methods*). The heuristic rules remained unmodified, as their use of a constant learning rate already encapsulates the assumption that the true weights change across decisions.

Fig. 2B illustrates the performance of the different learning rules. First, it shows that, for low-dimensional inputs the various probabilistic models yield comparable performances, but for high-dimensional inputs the approximate probabilistic learning rules outperform Bayes-optimal learning. In case of the latter,

these approximations were not actually harmful, but instead beneficial, for the same reason discussed further above. In particular, the more neurally realistic ADF variant that only tracked the diagonal of the covariance matrix again outperformed all other probabilistic models. Second, only the heuristic learning rule that performed gradient ascent on the log-likelihood achieved steady-state performance comparable to the approximate probabilistic rules, and then only for high input dimensionality and a specific choice of learning rate. This should come as no surprise, as its use of the likelihood function introduces more task structure information than the other heuristics use. The delta rule did not converge and therefore never achieved steady-state performance. Overall, the ADF variant that focused only on the diagonal covariance matrix achieved the best overall performance.

Learning Both Weights and a Latent State Prior Bias. Our learning rule can be generalized to learn prior biases in addition to the input weights. The prior we have used so far for the latent variable, $\mu \sim N(0, \sigma_\mu^2)$, is unbiased, as both $\mu \geq 0$ and $\mu < 0$ are equally likely. To introduce a prior bias, we instead used $\mu \sim N(m, \sigma_\mu^2)$, where m controls the bias through $P^+ \equiv p(\mu \geq 0) = \Phi(m/\sigma_\mu)$. A positive (or negative) m causes $P^+ > 1/2$ (or $< 1/2$), thus making $y = 1$ (or $y = -1$) the more likely correct choice even before evidence is accumulated. After evidence accumulation, such a prior results in the posterior:

$$p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t) = \Phi\left(\frac{\mathbf{w}^T \mathbf{x}(t) + \sigma_\mu^{-2} m}{\sqrt{t + \sigma_\mu^{-2}}}\right). \quad [9]$$

Comparing this to the unbiased posterior, Eq. 2, reveals the additional term $\sigma_\mu^{-2} m$ whose relative influence wanes over time.

This additional term has two consequences. First, appending the elements m and σ_μ^{-2} to the vectors \mathbf{w} and $\mathbf{x}(t)$, respectively, shows that \mathbf{w} and m can be learned jointly by the same learning rule we have derived before (*Materials and Methods*). Second, the term requires us to rethink the association between decision boundaries and choices. As Fig. 3C illustrates, such a prior causes a time-invariant shift in the association between the

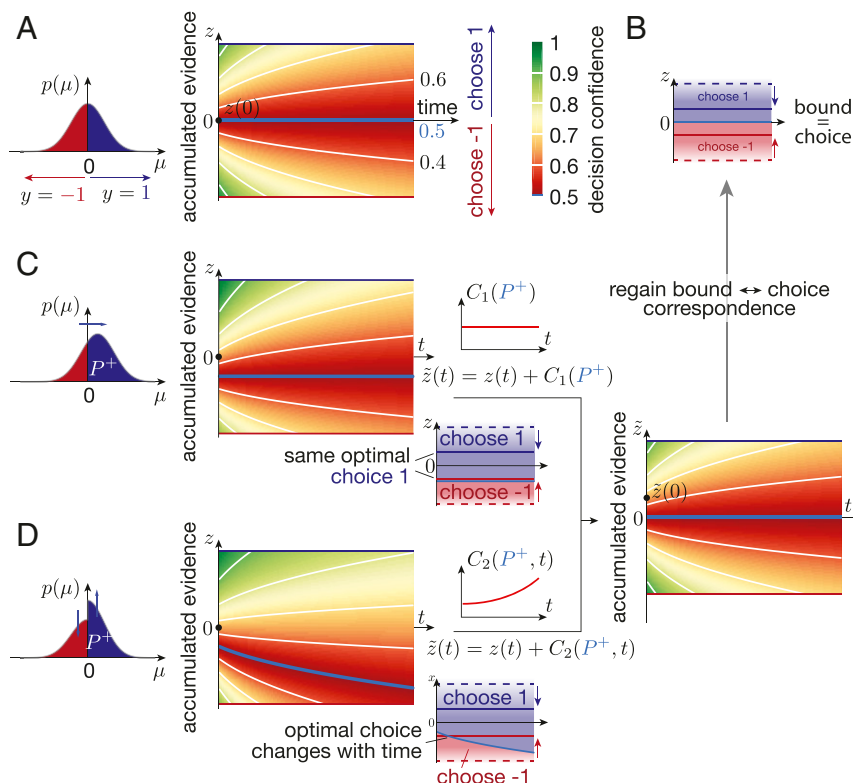


Fig. 3. Decision confidence, prior biases, and the relation between decision boundary and choice. (A) For an unbiased prior [i.e., $P^+ \equiv p(\mu \geq 0) = 1/2$], the decision confidence (color gradient) is symmetric around $z=0$ for each fixed time t . The associated posterior belief $p(\mu \geq 0|z(t), t)$ (numbers above/below “time” axis label; constant along white lines; $1/2$ along light blue line) promote choosing $y = 1$ and $y = -1$ above (blue area in B) and below (red area in B) $z=0$. (B) As a result, different choices are Bayes-optimal at the blue/red decision boundaries, as long as they are separated by $z=0$, irrespective of the boundary separation (solid vs. dashed blue red lines). (C) If the prior is biased by an overall shift, the decision confidence is countershifted by the same constant across all t . In this case, both decision boundaries might promote the same choice, which can be counteracted by a time-invariant shift of z by $C_1(P^+)$. (D) If the prior is biased by boosting one side while suppressing the other, the decision confidence shift becomes time dependent, such that the optimal choice at a time-invariant boundary might change over time. Counteracting this effect requires a time-dependent shift of z by $C_2(P^+, t)$. In both C and B, we have chosen $P^+ = 0.6$, for illustration.

accumulated evidence, $z(t) = \mathbf{w}^T \mathbf{x}(t)$, and the posterior belief of $\mu \geq 0$ and corresponding decision confidence. This shift makes it possible to have the same Bayes-optimal choice at both decision boundaries (Fig. 3C, blue/red decision areas). Hence, we have lost the mechanically convenient unique association between decision boundaries and choices. We recover this association by a boundary counter shift, such that these boundaries come to lie at the same decision confidence levels for opposite choices, making them asymmetric around $z = 0$. Mathematically, this is equivalent to shifting the evidence accumulation starting point, $\tilde{z}(0)$ away from zero in the opposite direction [Fig. 3C, shift by $C_1(P^+) = \sigma_\mu^{-2} m$; *SI Appendix*]. Therefore, a prior bias is implemented by a bias-dependent simple shift of the accumulation starting point, leading to a mechanically straightforward implementation of Bayes-optimal decision making with biased priors.

A consequence of the shifted accumulation starting point is that, for some fixed decision time t , the decision confidence at both boundaries is the same (Fig. 3C, Right). This seems at odds with the intuition that a biased prior ought to bias the decision confidence in favor of the more likely option. However, this mechanism does end up assigning higher average confidence to the more likely option because of reaction times. As the starting point is now further away from the less likely correct boundary, it will on average take longer to reach this boundary, which lowers the decision confidence since confidence decreases with elapsed time. Therefore, even though the decision confidence at both boundaries is the same for the given decision time, it will on

average across decision times be lower for the a priori non-preferred boundary, faithfully implementing this prior (see *SI Appendix* for a mathematical demonstration).

Our finding that a simple shift in the accumulation starting point is the Bayes-optimal strategy appears at odds with previous work that suggested that the optimal shift of the accumulator variable $z(t)$ varies with time (18). This difference stems from a different implementation of the bias. While we have chosen an overall shift in the prior by its mean (Fig. 3C), an alternative implementation is to multiply $p(\mu \geq 0)$ by P^+ , and $p(\mu < 0)$ by $1 - P^+$ (Fig. 3D), again resulting in $P^+ = p(\mu \geq 0)$. A consequence of this difference is that the associated shift of the posterior belief of $\mu \geq 0$ in the evidence accumulation space becomes time dependent. Then, the optimal choice at a time-invariant boundary in that space might change over time (Fig. 3D). Furthermore, undoing this shift to regain a unique association between boundaries and choices not only requires a shifted accumulation starting point, but additionally a time-dependent additive signal [$C_2(P^+, t)$ in Fig. 3D; *SI Appendix*], as was proposed in ref. 18. Which of the two approaches is more adequate depends on how well it matches the prior implicit in the task design. Our approach has the advantage of a simpler mechanistic implementation, as well as yielding a simple extension to the previously derived learning rule. How learning prior biases in the framework of ref. 18 could be achieved remains unclear (but see ref. 33).

Sequential Choice Dependencies due to Continuous Weight Tracking.

In everyday situations, no two decisions are made under the exact same circumstances. Nonetheless, we need to be able to learn from the outcome of past choices to improve future ones. A common assumption is that past choices become increasingly less informative about future choices over time. One way to express this formally is to assume that the world changes slowly over time—and that our aim is to track these changes. By “slow,” we mean that we can consider it constant over a single trial but that it is unstable over the course of an hour-long session. We implemented this tracking of the moving world, as in Fig. 2B, by slowly allowing the weights mapping evidence to decisions to change. With such continuously changing weights, weight learning never ends. Rather, the input weights are continuously adjusted to make correct choices more likely in the close future. After correct choices, this means that weights will be adjusted to repeat the same choice upon observing a similar input in the future. After incorrect choices, the aim is to adjust the weights to perform the opposite choice, instead. Our model predicts that, after an easy correct choice, in which confidence can be expected to be high, the weight adjustments are lower than after hard correct choices (Fig. 1D, Top, green line). As a consequence, we would expect the model to be more likely to repeat the same choices after correct and hard, than after correct and easy trials.

To test this prediction, we relied on the same simulation to generate Fig. 2B to measure how likely the model repeated the same choice after correct decisions. Fig. 4A illustrates that this repetition bias manifests itself in a shift of the psychometric curve that makes it more likely to repeat the previous choice. Furthermore, and as predicted, this shift is modulated by the difficulty of the previous choice and is stronger if the previous choice was easy (i.e., associated with a large $|\mu|$; Fig. 4B). Therefore, if the decision maker expects to operate in a volatile, slowly changing world, our model predicts a repetition bias to repeat the same choices after correct decisions, and that this bias is stronger if the previous choice was easy.

Unreliable Feedback Reduces Learning. What would occur if choice feedback is less-than-perfectly reliable? For example, the feedback itself might not be completely trustworthy, or hard to interpret. We simulated this situation by assuming that the feedback is inverted with probability β . Here, $\beta = 0$ implies the so far assumed perfectly reliable feedback, and $\beta = 1/2$ makes the feedback completely uninformative. This change impacts how decision confidence modulates the learning rate (Fig. 4C) as follows. First, it reduces the overall magnitude of the correction, with weaker learning for higher feedback noise. Second, it results

in no learning for highly confident choices that we are told are incorrect. In this case, one’s decision confidence overrules the unreliable feedback. This stands in stark contrast to the optimal learning rule for perfectly reliable feedback, in which case the strongest change to the current strategy ought to occur.

Discussion

Diffusion models are applicable to model decisions that require some accumulation of evidence over time, which is almost always the case in natural decisions. We extended previous work on the normative foundations of these models to more realistic situations in which the sensory evidence is encoded by a population of neurons, as opposed to just two neurons, as has been typically assumed in previous studies. We have focused on normative and mechanistic models for learning the weights from the sensory neurons to the decision integrator without additionally adjusting the decision boundaries, as weight learning is a problem that needs to be solved even if the decision boundaries are optimized at the same time.

From the Bayesian perspective, weight learning corresponds to finding the weight posterior given the provided feedback, and resulted in an approximate learning rule whose learning rate was strongly modulated by decision confidence. It suppressed learning after high-confidence correct decisions, supported learning for uncertain decisions irrespective of their correctness, and promoted strong change of the combination weights after wrong decisions that were made with high confidence (Fig. 1D). Evidence for such confidence-based learning has already been identified in human experiments (34), but not in a task that required the temporal accumulation of evidence in individual trials. Indeed, as we have previously suggested (22), such a modulation by decision confidence should arise in all scenarios of Bayesian learning in N -AFC tasks in which the decision maker only receives feedback about the correctness of their choices, rather than being told which choice would have been correct. In the 2-AFC task we have considered, being told that one’s choice was incorrect automatically reveals that the other choice was correct, making the two cases coincide. Moving from one-dimensional to higher-dimensional inputs requires performing the accumulation of evidence for each input dimension separately [Fig. 1B; Eqs. 6 and 12 require $x(t)$ rather than only $w^T x(t)$], even if triggering choices only requires a linear combination of $x(t)$. This is because uncertain input weights require keeping track of how each input dimension contributed to the particle crossing the decision boundary in order to correctly improve these weights upon feedback (i.e., proper credit assignment). The multidimensional evidence accumulation

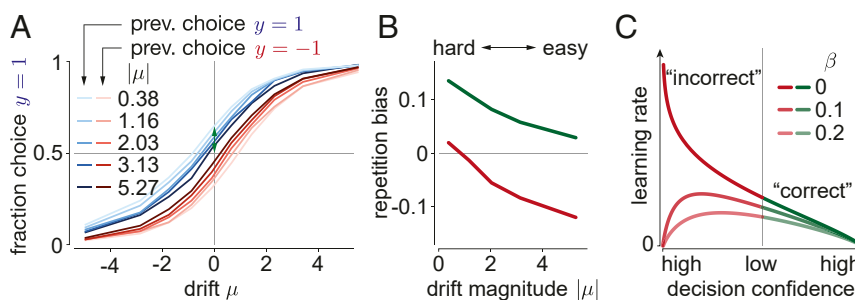


Fig. 4. Sequential choice dependencies due to continuous learning, and effects of noisy feedback. Bayes-optimal learning in a slowly changing environment predicts sequential choice dependencies with the following pattern. (A) After hard, correct choices (low prev. $|\mu|$; light colors), the psychometric curve is shifted toward repeating the same choice (blue/red = choice $y = 1/-1$). This shift decreases after easier, correct choices (high prev. $|\mu|$; dark colors). (B) We summarize these tuning curve shifts in the repetition bias, which is the probability of repeating the same choice to a $\mu = 0$ stimulus (example green arrow for $\mu = -0.38$ in A). After correct/incorrect choices (green/red curve), this leads to a win–stay/lose–switch strategy. Only the win–stay strategy is shown in A. (C) If choice feedback is noisy (inverted with probability β), the learning rate becomes overall lower. In particular for high-confidence choices with “incorrect” feedback, the learning rate becomes zero, as the learners trust their choice more than the feedback.

predicted by our work arises naturally if inputs encode full distributions across the task-relevant variables, such as in linear probabilistic population codes (35) that trigger decisions by bounding the pooled activity of all units that represent the accumulated evidence (36).

Continual weight learning predicts sequential choice dependencies that make the repetition of a previous, correct choice more likely, in particular if this choice was difficult (Fig. 4). Thus, based on assuming a volatile environment that promotes a continual adjustment of the decision-making strategy, we provide a rational explanation for sequential choice dependencies that are frequently observed in both humans and animals (e.g., refs. 37 and 38). In rodents making decisions in response to olfactory cues, we have furthermore confirmed that these sequential dependencies are modulated by choice difficulty, and that the exact pattern of this modulation depends on the stimulus statistics, as predicted by our theory (39) (but consistency with ref. 40 is unclear).

Last, we have clarified how prior biases ought to impact Bayes-optimal decision making in diffusion models. Extending the work of Hanks et al. (18), we have demonstrated that the exact mechanisms to handle these biases depend on the specifics of how these biases are introduced through the task design. Specifically, we have suggested a variant that simplifies these mechanisms and the learning of this bias. This variant predicts that the evidence accumulation offset, that has previously been suggested to be time-dependent, to become independent of time, and it would be interesting to see whether the lateral intraparietal cortex activity of monkeys performing the random-dot motion task, as recorded by Hanks et al. (but see ref. 41), would change accordingly.

Materials and Methods

We here provide an outline of the framework and its results. Detailed derivations are provided in *SI Appendix*.

Bayesian Decision Making with One and Multidimensional Diffusion Models.

We assume the latent state to be drawn from $\mu \sim N(m, \sigma_\mu^2)$, and the momentary evidence in each time step δt to provide information about this latent state by $\delta z_i | \mu \sim N(\mu \delta t, \delta t)$. The aim is to infer the sign of μ , and choose $y = 1$ if $\mu \geq 0$, and $y = -1$ otherwise. After having observed this evidence for some time $t \equiv n \delta t$, the posterior μ given all observed evidence $\delta z_{1:n}$ is by Bayes' rule given by the following:

$$p(\mu | \delta z_{1:n}) \propto N(\mu | m, \sigma_\mu^2) \prod_{i=1}^n N(\delta z_i | \mu \delta t, \delta t) \propto N\left(\mu \mid \frac{\sigma_\mu^2 m + z(t)}{\sigma_\mu^2 + t}, \frac{1}{\sigma_\mu^2 + t}\right). \quad [11]$$

In the above, all proportionalities are with respect to μ , and we have defined $z(t) = \sum_{i=1}^n \delta z_i$ and have used $t = \sum_{i=1}^n \delta t$. How to find the posterior belief for μ 's sign with $m = 0$ is described around Eq. 1.

We extend diffusion models to multidimensional inputs with momentary evidence $\delta \mathbf{x}_i | \mu, \mathbf{w} \sim N((\mathbf{a}\mu + \mathbf{b})\delta t, \Sigma \delta t)$, with \mathbf{a} , \mathbf{b} , and Σ chosen such that $\mathbf{w}^T \mathbf{x}(t) | \mu = z(t) | \mu \sim N(\mu t, t)$, as before. The posterior over μ and $\mu \geq 0$ is the same as for the one-dimensional case, with $z(t)$ replaced by $\mathbf{w}^T \mathbf{x}(t)$. Defining $\tilde{\mathbf{x}}(t) = \mathbf{x}(t) / \sqrt{\sigma_\mu^2 + t}$, we find $p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t) = \Phi(\mathbf{w}^T \tilde{\mathbf{x}}(t))$. As $y = 1$ and $y = -1$ correspond to $\mu \geq 0$ and $\mu < 0$, and $y = 1$ is only chosen if $p(\mu \geq 0 | \mathbf{w}, \mathbf{x}(t), t) \geq 1/2$, the decision confidence for $m = 0$ at some boundary $\mathbf{w}^T \mathbf{x}(t) = \pm \theta(t)$ is given by $\Phi(\theta(t) / \sqrt{\sigma_\mu^2 + t})$. If input weights are unknown, and the decision maker holds belief $\mathbf{w} \sim N(\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$ about these weights, the decision confidence needs to additionally account for weight uncertainty by marginalizing over \mathbf{w} , resulting in Eq. 5.

Probabilistic and Heuristic Learning Rules. We find the approximate posterior $q(\mathbf{w}) = N(\mathbf{w} | \boldsymbol{\mu}_w^*, \boldsymbol{\Sigma}_w^*)$ that approximates the target posterior p Eq. 4 by ADF.

This requires minimizing the Kullback–Leiber divergence $KL(p|q)$ (26, 29), resulting in Eq. 6 for the posterior mean, and the following:

$$\boldsymbol{\Sigma}_w^* = \boldsymbol{\Sigma}_w + \xi_{\text{cov}}(\gamma) \left((\boldsymbol{\Sigma}_w^{-1} + \tilde{\mathbf{x}} \tilde{\mathbf{x}}^T)^{-1} - \boldsymbol{\Sigma}_w \right), \quad [12]$$

with learning rate modulators $\xi_w(\gamma) = N(\gamma | 0, 1) / \Phi(\gamma)$ and $\xi_{\text{cov}}(\gamma) = \xi_w(\gamma)^2 + \xi_w(\gamma)\gamma$, and where we have defined $\gamma \equiv y^* \boldsymbol{\mu}_w^T \tilde{\mathbf{x}} / \sqrt{1 + \tilde{\mathbf{x}}^T \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}}$, which is monotonic in the decision confidence (Eq. 5). Noisy choice feedback (Fig. 4C) changes the likelihood to assume reversed feedback with probability β , and follow the same procedure as above to derive the posterior moments (*SI Appendix*). The ADF variant that only tracks the diagonal covariance elements assumes $\boldsymbol{\Sigma}_w$ to be diagonal, and only computes the diagonal elements of $\boldsymbol{\Sigma}_w^*$. A second-order Taylor expansion of the log of Eq. 4 leads to update equations similar to Eqs. 6 and 12, but without the normalization by weight uncertainty (see *SI Appendix* for details). All heuristic learning rules are described in the main text.

We modeled nonstationary input weights by $\mathbf{w}_n | \mathbf{w}_{n-1} \sim N(\mathbf{A} \mathbf{w}_{n-1} + \mathbf{b}, \boldsymbol{\Sigma}_d)$ after a decision in trial $n - 1$. This weight transition is taken into account by the probabilistic learning rules by setting the parameter priors to $\boldsymbol{\mu}_{w,n} = \mathbf{A} \boldsymbol{\mu}_{w,n-1}^* + \mathbf{b}$ and $\boldsymbol{\Sigma}_{w,n} = \mathbf{A} \boldsymbol{\Sigma}_{w,n-1}^* \mathbf{A}^T + \boldsymbol{\Sigma}_d$. For stationary weights, we have $\mathbf{A} = \mathbf{I}$, $\mathbf{b} = 0$, and $\boldsymbol{\Sigma}_d = 0$.

Bayes-optimal weight inference was for stationary weights performed by Gibbs sampling for probit models, and for nonstationary weights by particle filtering (*SI Appendix*).

Simulation Details. We used parameters $\mathbf{a} = \mathbf{w} / \|\mathbf{w}\|^2$ and $\mathbf{b} = 0$ for the momentary evidence $\delta \mathbf{x}$. Its covariance Σ was generated to feature eigenvalues that drop exponentially from $\sigma_x^2 = 2 / \|\mathbf{w}\|^2$ to zero until it reaches a constant $\sigma_0^2 = 0.001 / \|\mathbf{w}\|^2$ noise baseline, as qualitatively observed in neural populations. It additionally contains an eigenvector \mathbf{w} with eigenvalue set to guarantee $\mathbf{w}^T \Sigma \mathbf{w} = 1$, limiting the information that $\delta \mathbf{x}$ provides about μ . For nonstationary weights, all momentary evidence parameters are adjusted after each weight change (*SI Appendix*). The diffusion model bounds $\pm \theta$ were time-invariant and tuned to maximize the reward rate when using the correct weights. The reward rate is given by $(p(\text{correct}) - c_{\text{accum}} t) / (t_{\text{tr}} + t)$, where averages were across trials, and we used evidence accumulation cost $c_{\text{accum}} = 0.01$ and intertrial interval $t_{\text{tr}} = 2s$. We used $\sigma_\mu^2 = 3^2$ to draw μ in each trial, and drew \mathbf{w} from $\mathbf{w} \sim N(1, \mathbf{I})$ before each trial sequence. For nonstationary weights, we resampled weights after each trial according to $\mathbf{w}_n | \mathbf{w}_{n-1} \sim N(\lambda \mathbf{w}_{n-1} + (1 - \lambda) \sigma_\mu^2 \mathbf{I}, \sigma_\mu^2 \mathbf{I})$, with decay factor $\lambda = 1 - 0.01$ and $\sigma_\mu^2 = 1 - \lambda^2$ to achieve steady-state mean $\mathbf{1}$ and identity covariance.

To compare the weight learning performance of ADF to alternative models (Fig. 2A), we simulated 1,000 learning trials 5,000 times, and reported the reward rate per trial averaged across these 5,000 repetitions. To assess steady-state performance (Fig. 2B), we performed the same procedure with nonstationary weights and reported reward rate averaged over the last 100 trials, and over 5,000 repetitions. The same 100 trials were used to compute the sequential choice dependencies in Fig. 4A and B. To simulate decision making with diffusion models and uncertain weights, we used the current mean estimate $\langle \mathbf{w} \rangle$ of the input weights to linearly combine the momentary evidence. The probabilistic learning rules were all independent of the specific choice of this estimate. The learning rate in Fig. 1D shows the prefactor to $y^* \boldsymbol{\Sigma}_w \tilde{\mathbf{x}}$ in Eq. 6 over decision confidence for a subsample of the last 10,000 trials of a single 15,000 trial simulation with nonstationary weights. For the Gibbs sampler, we drew 10 burn-in samples, followed by 200 samples in each trial. For the particle filter, we simulated 1,000 particles.

ACKNOWLEDGMENTS. This work was supported by a James S. McDonnell Foundation Scholar Award (220020462) (J.D.) and grants from the National Institute of Mental Health (R01MH115554) (J.D.), the Swiss National Science Foundation (www.snf.ch) (31003A_143707 and 31003A_165831) (A.P.), the Champalimaud Foundation (Z.F.M.), the European Research Council (Advanced Investigator Grants 250334 and 671251) (Z.F.M.), the Human Frontier Science Program (Grant RGP0027/2010) (Z.F.M. and A.P.), the Simons Foundation (Grant 325057) (Z.F.M. and A.P.), and Fundação para a Ciência e a Tecnologia (A.G.M.).

1. K. Doya, S. Ishii, A. Pouget, R. P. N. Rao, *Bayesian Brain: Probabilistic Approaches to Neural Coding* (MIT Press, 2006).
2. R. Ratcliff, A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
3. R. Ratcliff, G. McKoon, The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).

4. R. Bogacz, E. Brown, J. Moehlis, P. Holmes, J. D. Cohen, The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
5. R. Ratcliff, P. L. Smith, A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367 (2004).

6. P. I. Frazier, A. J. Yu, Sequential hypothesis testing under stochastic deadlines. *Adv. Neural Inf. Process. Syst.* **2008**, 1–8 (2008).
7. S. Tajima, J. Drugowitsch, A. Pouget, Optimal policy for value-based decision-making. *Nat. Commun.* **7**, 12400 (2016).
8. J. Drugowitsch, R. Moreno-Bote, A. K. Churchland, M. N. Shadlen, A. Pouget, The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* **32**, 3612–3628 (2012).
9. J. I. Gold, M. N. Shadlen, Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* **36**, 299–308 (2002).
10. J. Drugowitsch, G. C. DeAngelis, E. M. Klier, D. E. Angelaki, A. Pouget, Optimal multisensory decision-making in a reaction-time task. *eLife* **3**, 1–19 (2014).
11. P. Dayan, S. Kakade, P. R. Montague, Learning and selective attention. *Nat. Neurosci.* **3** (suppl), 1218–1223 (2000).
12. P. Dayan, A. J. Yu, Uncertainty and learning. *IETE J. Res.* **49**, 171–181 (2003).
13. K. P. Körding, D. M. Wolpert, Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
14. A. C. Courville, N. D. Daw, D. S. Touretzky, Bayesian theories of conditioning in a changing world. *Trends Cogn. Sci.* **10**, 294–300 (2006).
15. R. Ratcliff, A theory of order relations in perceptual matching. *Psychol. Rev.* **88**, 552–572 (1981).
16. P. Gomez, R. Ratcliff, M. Perea, The overlap model: A model of letter position coding. *Psychol. Rev.* **115**, 577–600 (2008).
17. R. Ratcliff, J. J. Starns, Modeling confidence judgments, response times, and multiple choices in decision making: Recognition memory and motion discrimination. *Psychol. Rev.* **120**, 697–719 (2013).
18. T. D. Hanks, M. E. Mazurek, R. Kiani, E. Hopp, M. N. Shadlen, Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352 (2011).
19. R. Moreno-Bote, Decision confidence and uncertainty in diffusion models with partially correlated neuronal integrators. *Neural Comput.* **22**, 1786–1811 (2010).
20. J. Palmer, A. C. Huk, M. N. Shadlen, The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* **5**, 376–404 (2005).
21. J. O. Berger, *Statistical Decision Theory and Bayesian Analysis* (Springer, ed. 2, 1993).
22. A. Pouget, J. Drugowitsch, A. Kepecs, Confidence and certainty: Distinct probabilistic quantities for different goals. *Nat. Neurosci.* **19**, 366–374 (2016).
23. B. B. Averbeck, P. E. Latham, A. Pouget, Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
24. R. Moreno-Bote et al., *Information-Limiting Correlations* (Nature Publishing Group, 2014).
25. R. Kiani, M. N. Shadlen, Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324**, 759–764 (2009).
26. C. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2006).
27. J. Drugowitsch, G. C. DeAngelis, D. E. Angelaki, A. Pouget, Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making. *Elife* **4**, 1–11 (2015).
28. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (Wiley, ed. 2, 2006).
29. K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
30. T. Graepel, J. Quiñero-Candela, T. Borchert, R. Herbrich, “Web-scale Bayesian click-through rate prediction for sponsored search advertising in Microsoft’s Bing search engine” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, J. Fürnkranz, T. Joachims, Eds. (ACM Press, New York, 2010), pp. 13–20.
31. W. Chu, M. Zinkevich, L. Li, A. Thomas, B. Tseng, “Unbiased online active learning in data streams” in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining—KDD ’11* (ACM Press, New York, 2011), pp. 195–203.
32. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, ed. 2, 2018).
33. A. Zylberberg, D. M. Wolpert, M. N. Shadlen, Counterfactual reasoning underlies the learning of priors in decision making. *Neuron* **99**, 1083–1097.e6 (2018).
34. F. Meyniel, S. Dehaene, Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E3859–E3868 (2017).
35. W. J. Ma, J. M. Beck, P. E. Latham, A. Pouget, Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 (2006).
36. J. M. Beck et al., Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152 (2008).
37. L. Busse et al., The detection of visual contrast in the behaving mouse. *J. Neurosci.* **31**, 11351–11361 (2011).
38. A. J. Yu, J. D. Cohen, Sequential effects: Superstition or rational behavior? *Adv. Neural Inf. Process. Syst.* **21**, 1873–1880 (2008).
39. A. G. Mendonça et al., The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs. bioRxiv:10.1101/501858 (19 December 2018).
40. A. E. Urai, A. Braun, T. H. Donner, Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* **8**, 14637 (2017).
41. V. Rao, G. C. DeAngelis, L. H. Snyder, Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *J. Neurosci.* **32**, 10063–10074 (2012).