Uncertainty and Exploration

Samuel J. Gershman

Department of Psychology and Center for Brain Science

Harvard University

Address for correspondence:

Samuel Gershman

Department of Psychology

Harvard University

52 Oxford St., room 295.05

Cambridge, MA 02138

Phone: 773-607-9817

E-mail: gershman@fas.harvard.edu

KEYWORDS: explore-exploit dilemma; reinforcement learning; Bayesian inference

Abstract

In order to discover the most rewarding actions, agents must collect information about their environment, potentially foregoing reward. The optimal solution to this "explore-exploit" dilemma is computationally intractable, but principled algorithmic approximations exist. These approximations utilize uncertainty about action values in different ways. Some *random* exploration algorithms scale the level of choice stochasticity with the level of uncertainty. Other *directed* exploration algorithms add a "bonus" to action values with high uncertainty. Random exploration algorithms are sensitive to *total* uncertainty across actions, whereas directed exploration algorithms are sensitive to *relative* uncertainty. This paper reports a multi-armed bandit experiment in which total and relative uncertainty were orthogonally manipulated. We found that humans employ both exploration strategies, and that these strategies are independently controlled by different uncertainty computations.

Introduction

Uncertainty lies at the heart of decision making in the real world. A bee in search of nectar and a venture capitalist in search of an investment both need to explore their options in order to reduce their uncertainty, at the expense of exploiting the currently best option. In the absence of uncertainty, no explore-exploit dilemma would exist. The question addressed here is how humans use uncertainty to guide exploration.

The optimal solution to the explore-exploit dilemma is computationally intractable, but computer scientists have developed algorithmic approximations that can provably approach optimal behavior (Sutton & Barto, 1998). Psychologists have also studied algorithmic hypotheses about how humans balance exploration and exploitation (Cohen, McClure, & Yu, 2007; Mehlhorn et al., 2015), but only recently have the links between computational and psychological algorithms been systematically investigated (Gershman, 2018). Key to this synthesis is the idea that uncertainty can guide exploration in two qualitatively different ways: by adding *randomness* into choice behavior, or by *directing* choice towards uncertain options.

The pioneering work of Wilson, Geana, White, Ludvig, and Cohen (2014) demonstrated that humans use both random and directed exploration in carefully designed two-armed bandit task. When the subject had more experience with one option (and hence less uncertainty), she favored the more uncertain option, indicating a form of directed exploration. In addition, subjects increased the randomness in their choices when they were more uncertain, spreading their choices across both high- and low-value options. Directed and random exploration strategies are dissociable, developing on different timescales across the lifespan (Somerville et al., 2017) and relying on different neural substrates (Zajkowski, Kossut, & Wilson, 2017).

The directed/random distinction is closely related to the distinction between two families

of exploration algorithms that have been studied extensively in machine learning. Directed exploration can be realized by adding *uncertainty bonuses* to estimated values (Auer, Cesa-Bianchi, & Fischer, 2002; Brafman & Tennenholtz, 2002; Dayan & Sejnowski, 1996; Kolter & Ng, 2009; Srinivas, Krause, Seeger, & Kakade, 2010). One of the most prominent versions of this approach is known as the *upper confidence bound* (UCB) algorithm (Auer et al., 2002), which chooses action a_t on trial t according to:

$$a_t = \operatorname*{argmax}_k Q_t(k) + U_t(k), \tag{1}$$

where k indexes actions and $U_t(k)$ is the upper confidence bound that plays the role of an uncertainty bonus. Under a Bayesian analysis (Srinivas et al., 2010), $Q_t(k)$ corresponds to the posterior mean and the uncertainty bonus is proportional to the posterior standard deviation $\sigma_t(k)$.

Random exploration has an even older pedigree in machine learning, dating back to Thompson's work in the 1930s (Thompson, 1933). In what is now known as *Thompson sampling*, a random value function from the posterior is drawn and then the agent chooses greedily with respect to the random draw. Like UCB, Thompson sampling uses uncertainty to promote exploration, but does so by encouraging stochasticity (i.e., a form of probability matching) rather than via a response bias. In psychophysical terms, Thompson sampling changes the slope of the function relating action values to choice probabilities, whereas UCB changes the intercept (indifference point).

Recognizing the dissociable (and possibly complementary) nature of directed and random exploration, computer scientists have also constructed hybrids of UCB and Thompson sampling (Chapelle & Li, 2011; May, Korda, Lee, & Leslie, 2012). A recent report (Gershman, 2018) provided the most direct evidence for such hybrids in human decision making, demonstrating that uncertainty influences both the intercept and slope of the choice prob-

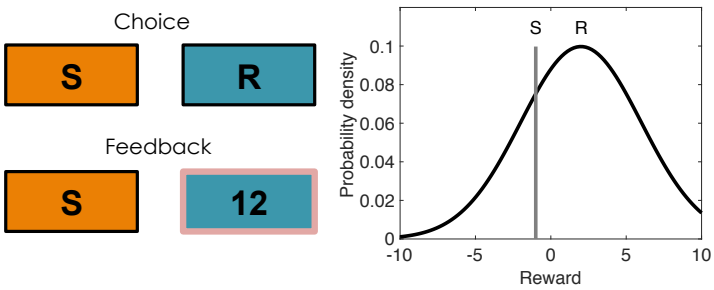

Figure 1: **Task design**. (Left) On each trial, subjects choose between two options and receive reward feedback in the form of points. Safe options are denoted by "S" and risky options are denoted by "R." On each block, one or both of the options may be safe or risky. (Right) The rewards for risky options are drawn from a Gaussian distribution that remains constant during each block. The rewards for safe options are deterministic. Both the mean for the risky option and the reward value of the safe option are drawn from a zero-mean Gaussian distribution that is resampled at each block transition.

ability function. Critically, the intercept and slope effects derive from different uncertainty computations: the intercept is governed by *relative* uncertainty (the difference in posterior uncertainty between the two options, defined formally below), whereas the slope is governed by *total* uncertainty (the sum of posterior uncertainty across the options). This suggests that experimental manipulations of these two factors should produce dissociable effects.

Here we pursue this line of reasoning using a two-armed bandit task, with the additional twist that we inform subjects about the riskiness of each option (Figure 1). On each trial, participants were given a choice between two options, labeled either as "safe" (S) or "risky" (R). The safe options always delivered the same reward, whereas the risky options delivered stochastic rewards (with Gaussian noise). Denoting trial types by compound labels (e.g., "SR" denote trials in which the left option is safe and the right option is risky), we used the comparison between SR and RS trials to isolate the effects of relative uncertainty, and the comparison between SS and RR trials to isolate the effects of total uncertainty (Figure 2).

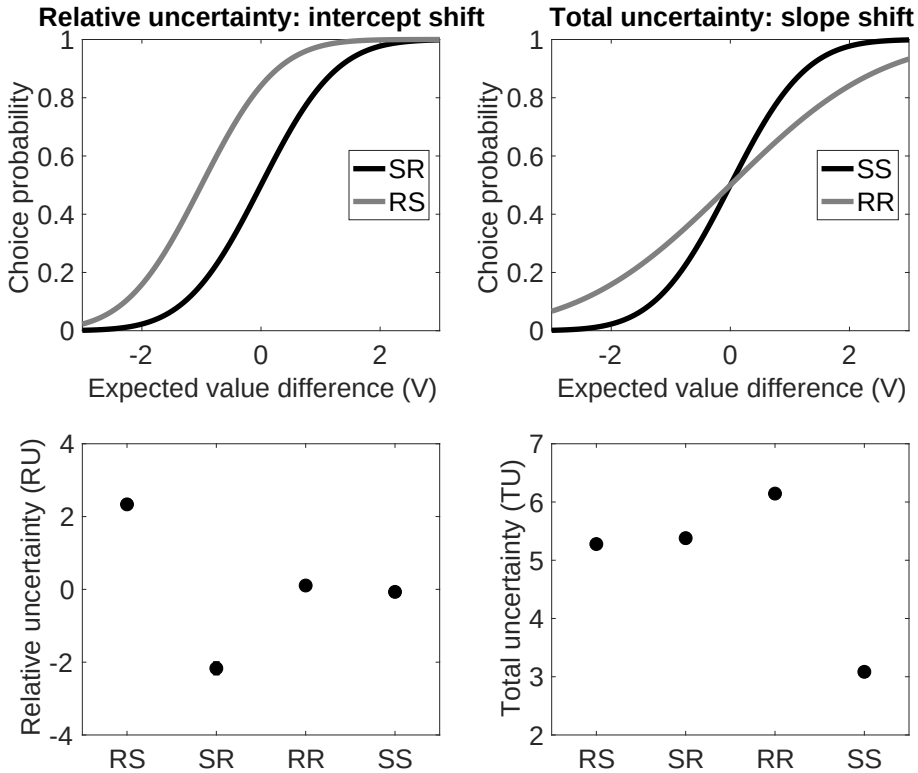

Figure 2: **Relative and total uncertainty**. (Top) Illustration of how choice probabilities change as a function of the experimental condition and form of uncertainty. (Bottom) Average relative and total uncertainty in each condition. Conditions are denoted by safe/risky compounds; for example, "SR" denotes a trial in which the left arm is safe and the right arm is risky.

Materials and Methods

Code and data for reproducing all analyses reported in this paper are available at

`https://github.com/sjgershm/exploration_uncertainty`.

Subjects

Forty-six subjects were recruited via the Amazon Mechanical Turk web service and paid $2.00. The experiments were approved by the Harvard Institutional Review Board.

Stimuli and Procedure

Participants played 30 two-armed bandits in blocks of 10 trials. On each trial, participants chose one of the arms and received reward feedback (points). They were instructed to choose the arm that maximizes their total points. On each block, the mean reward $\mu(k)$ for each arm was drawn from a Gaussian with mean 0 and variance $\tau_0^2(k) = 100$. The arms were randomly designated "safe" or "risky," indicated by an S or R (respectively), and these designations were resampled after a block transition. When participants chose the risky arm, they received stochastic rewards drawn from a Gaussian with mean $\mu(R)$ and variance $\tau^2(R) = 16$. When they chose arm 2, they always received a reward of 0. When participants chose the safe arm, they received a reward of $\mu(S)$.

The exact instructions for participants were as follows:

> *In this task, you have a choice between two slot machines, represented by colored buttons. When you click one of the buttons, you will win or lose points. Choosing the same slot machine will not always give you the same points, but one slot machine is always better than the other. Your goal is to choose the slot machine that will give you the most points. After making your choice, you will receive feedback about the outcome. Sometimes the machines are "safe" (always delivering the same feedback), and sometimes the machines are "risky" (delivering variable feedback). Before you make a choice, you will get information about each machine: "S" indicates SAFE, "R" indicates RISKY. Note that safe/risky is independent of how rewarding a machine is: a risky machine may deliver more points on average than a safe machine, and vice versa. You cannot predict how good a machine is based on whether it is safe or risky. You will play 30 games, each with a different pair of slot machines. Each game will consist of 10 trials.*

Belief updating model

To derive estimates of expected value and uncertainty, we assume that subjects approximate an ideal Bayesian learner. Given the Gaussian distributional structure underlying our task, the posterior over the value of arm k is Gaussian with mean $Q_t(k)$ and variance $\sigma_t^2(k)$. These sufficient statistics can be recursively updated using the Kalman filtering equations:

$$Q_{t+1}(a_t) = Q_t(a_t) + \alpha_t[r_t - Q_t(a_t)] \tag{2}$$

$$\sigma_{t+1}^2(a_t) = \sigma_t^2(a_t) - \alpha_t \sigma_t^2(a_t), \tag{3}$$

where the learning rate α_t is given by:

$$\alpha_t = \frac{\sigma_t^2(a_t)}{\sigma_t^2(a_t) + \tau^2(a_t)}. \tag{4}$$

The initial values were set to the prior means, $Q_1(k) = 0$ for all k, and the initial variances were set to the prior variance, $\sigma_1^2(k) = \tau_0^2(k)$. Although the Kalman filter is an idealization of human learning, it has been shown to account well for human behavior in bandit tasks (Daw, O'doherty, Dayan, Seymour, & Dolan, 2006; Gershman, 2018; Schulz, Konstantinidis, & Speekenbrink, 2015; Speekenbrink & Konstantinidis, 2015).

Choice probability analysis

Gershman (2018) showed that Thompson sampling, UCB, and a particular hybrid of the two imply a probit regression model of choice:

$$P(a_t = 1|\mathbf{w}) = \Phi(w_1 \mathrm{V}_t + w_2 \mathrm{RU}_t + w_3 \mathrm{V}_t/\mathrm{TU}_t), \tag{5}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard Gaussian distribution (mean 0 and variance 1), and the regressors are defined as follows:

- Estimated value difference, $\mathrm{V}_t = Q_t(1) - Q_t(2)$.
- Relative uncertainty, $\mathrm{RU}_t = \sigma_t(1) - \sigma_t(2)$.
- Total uncertainty, $\mathrm{TU}_t = \sqrt{\sigma_t^2(1) + \sigma_t^2(2)}$.

Thompson sampling predicts a significant positive effect of V/TU on choice probability, but not of RU or V, whereas UCB predicts a significant positive effect of both V and RU, but not of V/TU (Figure 2, top left). We used maximum likelihood estimation to fit the coefficients ($\mathbf{w}$) in the probit regression model.

In addition to this model-based analysis, we analyzed choices as a function of experimental condition.

$$P(a_t = 1|\mathbf{w}) = \Phi\left(\sum_j w_1^j \pi_{tj} + w_2^j \pi_{tj} \mathrm{V}_t\right), \tag{6}$$

where $\pi_{tj} = 1$ if trial t is assigned to condition j, and 0 otherwise. We refer to the w_1 terms as intercepts and the w_2 terms as slopes.[1]

Results

The hybrid random/directed exploration model hypothesizes that all three computational regressors (V, RU, V/TU) should be predictors of choice. We therefore confirmed that maximum likelihood estimates of the corresponding coefficients were significantly greater than 0 [Figure 1; V: $t(45) = 13.94, p < 0.0001$; RU: $t(45) = 4.77, p < 0.0001$; V/TU:

[1]Note that although this analysis is "model-free" in the sense that it does not use the computationally derived uncertainty regressors, it is still dependent on the model-based value estimates.

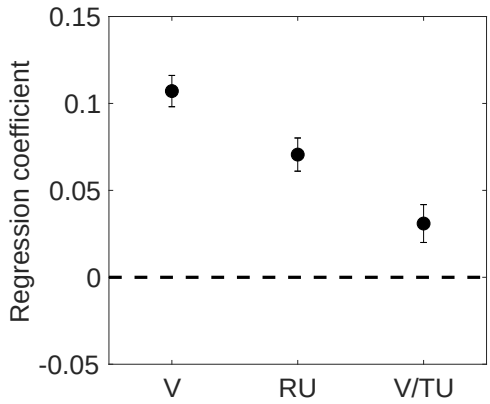

Figure 3: **Probit regression results: computational variables**. Average maximum likelihood estimates and within-subject standard errors for each regressor. V = estimated value difference between arms, RU = relative uncertainty, TU = total uncertainty (see Materials and Methods for details).

$t(45) = 2.29, p < 0.05$]. Bayesian random-effects model comparison (Rigoux, Stephan, Friston, & Daunizeau, 2014) strongly favored this three-parameter model over a one-parameter model with only the V regressor (protected exceedance probability: 0.96). This supports previous results showing that humans use both random and directed exploration strategies (Gershman, 2018; Somerville et al., 2017; Wilson et al., 2014).

We next addressed the central question of the paper: can random and directed exploration be independently manipulated? As predicted, changing relative uncertainty (RS vs. SR) alter the intercept of the choice probability function (Figure 3): there was a significant difference between the intercepts for the RS and SR conditions [$t(45) = 3.58, p < 0.001$]. Furthermore, the RS intercept was significantly greater than 0 [$t(45) = 2.64, p < 0.05$] and the SR intercept was significantly less than 0 [$t(45) = 3.66, p < 0.001$], indicating an uncertainty-directed choice bias, as predicted by the theory. Critically, there was no significant effect of total uncertainty [RR vs. SS, $p = 0.48$], consistent with the hypothesis that uncertainty-directed biases only emerge when there is a difference in relative uncertainty.

The pattern is flipped when we inspect the slope parameter estimates (Figure 4): increas-

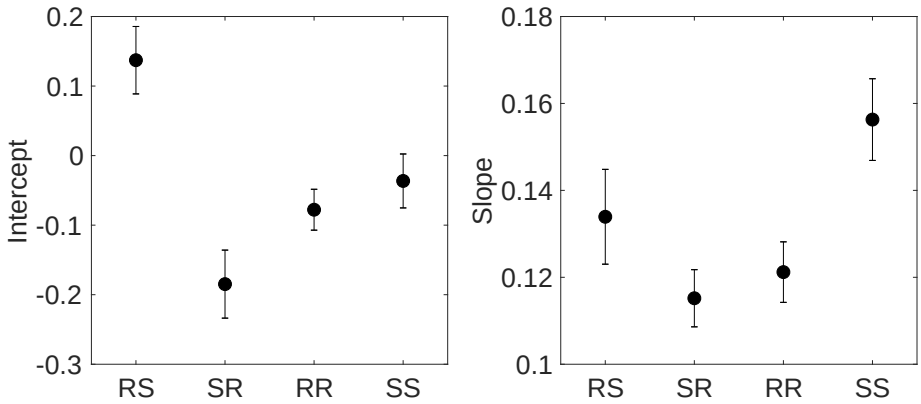

Figure 4: **Probit regression results: experimental variables**. Average maximum likelihood estimates and within-subject standard errors for each regressor. (Left) Bias coefficients. (Right) Slope coefficients.

ing total uncertainty (RR vs. SS) reduced the slope [$t(45) = 3.20, p < 0.005$], but relative uncertainty (RS vs. SR) did not have a significant effect ($p = 0.16$). This finding is consistent with our hypothesis that the random component of exploration would be specifically sensitive to changes in total uncertainty.

Finally, we examined response times as an additional source of evidence about exploration strategies. One of the most successful theories of value-based choice is the drift-diffusion model (DDM), which assumes that the value difference between two options drives a stochastic accumulator until it reaches a decision threshold (Krajbich, Armel, & Rangel, 2010; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010; Ratcliff & Frank, 2012; Summerfield & Tsetsos, 2012). The DDM has also been shown to predict choice in reinforcement learning tasks, where the value is dynamically updated based on experience (Frank et al., 2015; Millner, Gershman, Nock, & den Ouden, 2018; Pedersen, Frank, & Biele, 2017). Within the DDM framework, a directed exploration strategy predicts that response times should be faster for risky choices than for safe choices on RS and SR trials, because uncertainty acts as a bonus added to the values. This prediction was confirmed in our data [Figure 5; $t(45) = 2.01, p < 0.05$].

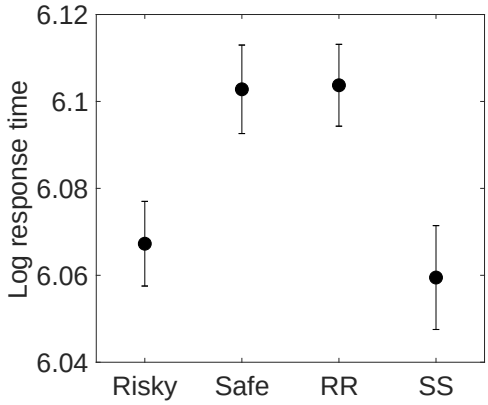

Figure 5: **Response time analysis**. Log response times (mean ± within-subject standard error). "Risky" denotes risky choices on SR and RS trials; "Safe" denotes safe choices on SR and RS trials. SS and RR response times are collapsed across arms.

When uncertainty is equated across the arms, directed exploration does not predict any difference in response time. Random exploration strategies, in contrast, correctly predict that total uncertainty will act divisively on values, thereby slowing response times when both arms are risky compared to when both arms are safe [Figure 5; $t(45) = 2.28, p < 0.05$].

Discussion

By separately manipulating relative and total uncertainty, we were able to independently influence directed and random exploration, lending support to the contention that these strategies coexist and jointly determine exploratory behavior (Gershman, 2018; Wilson et al., 2014). In particular, increasing relative uncertainty by making one option riskier than the other caused participants to shift their preference towards the risky option in a value-independent manner, consistent with a change in the intercept (indifference point) of the choice probability function. This manipulation had no effect on the slope of the choice probability function. In contrast, increasing total uncertainty by making both options risky decreased the slope relative to when both options were safe, with no effect on the intercept.

In other words, increasing total uncertainty caused choices to become more stochastic.

This dissociation between strategies guiding choice behavior was mirrored by a dissociation in response times. When one option was riskier than the other, subjects were faster in choosing the risky option. However, the same subjects were *slower* when both options were risky compared to when both options were safe. Taken together, these findings demonstrate that risk can have qualitatively different effects on both choice and response time, depending on the underlying uncertainty computation (i.e., total vs. relative).

We note here that preference for the risky option appears to be in opposition to what would be predicted by theories of risk aversion, as pointed out by previous authors (Gershman, 2018; Payzan-LeNestour & Bossaerts, 2012; Wilson et al., 2014). This does not mean that our subjects were not risk averse; rather, we suspect that the imperative to explore in a multi-armed bandit task overwhelmed the tendency to avoid risks seen in typical gambling paradigms.

The role of uncertainty-guided exploration has come to occupy an increasingly important place in theories of reinforcement learning (Gershman & Niv, 2015; Knox, Otto, Stone, & Love, 2011; Navarro, Newell, & Schulze, 2016; Payzan-LeNestour & Bossaerts, 2011; Pearson, Hayden, Raghavachari, & Platt, 2009; Schulz et al., 2015; Speekenbrink & Konstantinidis, 2015; Zhang & Yu, 2013), superseding earlier models of exploratory choice based on a fixed source of decision noise, as in ϵ-greedy and softmax policies (e.g., Daw et al., 2006). This shift has been accompanied by a deeper understanding of how reinforcement learning circuits in the basal ganglia compute, represent, and transmit uncertainty to downstream decision making circuits (Gershman, 2017; Lak, Nomoto, Keramati, Sakagami, & Kepecs, 2017; Starkweather, Babayan, Uchida, & Gershman, 2017). Despite this progress, we are only beginning to understand how these circuits implement dissociable channels regulating directed and random exploration (Warren et al., 2017; Zajkowski et al., 2017).

Acknowledgments

I am grateful to Eric Schulz for helpful discussions. This research was supported by the NSF Collaborative Research in Computational Neuroscience (CRCNS) Program Grant R01-1207833.

References

Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, *47*, 235–256.

Brafman, R. I., & Tennenholtz, M. (2002). R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, *3*, 213–231.

Chapelle, O., & Li, L. (2011). An empirical evaluation of Thompson sampling. In *Advances in neural information processing systems* (pp. 2249–2257).

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should i stay or should i go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *362*, 933–942.

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*, 876–879.

Dayan, P., & Sejnowski, T. J. (1996). Exploration bonuses and dual control. *Machine Learning*, *25*, 5–22.

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fmri and eeg predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, *35*, 485–494.

Gershman, S. J. (2017). Dopamine, inference, and uncertainty. *Neural Computation*, *29*, 3311–3326.

Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, *173*, 34–42.

Gershman, S. J., & Niv, Y. (2015). Novelty and inductive generalization in human reinforcement learning. *Topics in Cognitive Science*, *7*, 391–415.

Knox, W. B., Otto, A. R., Stone, P., & Love, B. C. (2011). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*, *2*.

Kolter, J. Z., & Ng, A. Y. (2009). Near-Bayesian exploration in polynomial time. In *Proceedings of the 26th annual international conference on machine learning* (pp. 513–520).

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*, 1292.

Lak, A., Nomoto, K., Keramati, M., Sakagami, M., & Kepecs, A. (2017). Midbrain dopamine neurons signal belief in choice accuracy during a perceptual decision. *Current Biology*, *27*, 821–832.

May, B. C., Korda, N., Lee, A., & Leslie, D. S. (2012). Optimistic Bayesian sampling in contextual-bandit problems. *Journal of Machine Learning Research*, *13*, 2069–2106.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*, 191–215.

Millner, A. J., Gershman, S. J., Nock, M. K., & den Ouden, H. E. (2018). Pavlovian control of escape and avoidance. *Journal of Cognitive Neuroscience*, 1–12.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, *5*, 437–449.

Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore-exploit dilemma in static and dynamic

environments. *Cognitive Psychology*, *85*, 43–77.

Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, *7*, e1001048.

Payzan-LeNestour, E., & Bossaerts, P. (2012). Do not bet on the unknown versus try to find out more: estimation uncertainty and "unexpected uncertainty" both modulate exploration. *Frontiers in Neuroscience*, *6*.

Pearson, J. M., Hayden, B. Y., Raghavachari, S., & Platt, M. L. (2009). Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Current Biology*, *19*, 1532–1537.

Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*, 1234–1251.

Ratcliff, R., & Frank, M. J. (2012). Reinforcement-based decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. *Neural Computation*, *24*, 1186–1229.

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studiesrevisited. *NeuroImage*, *84*, 971–985.

Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2015). Learning and decisions in contextual multi-armed bandit tasks. In *Proceedings of the 37th annual conference of the cognitive science society* (pp. 2122–2127).

Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology: General*, *146*, 155–164.

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, *7*, 351–367.

Srinivas, N., Krause, A., Seeger, M., & Kakade, S. M. (2010). Gaussian process optimization

in the bandit setting: No regret and experimental design. In *Proceedings of the 27th international conference on machine learning* (pp. 1015–1022).

Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, *20*, 581.

Summerfield, C., & Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Frontiers in Neuroscience*, *6*.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.

Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, *25*, 285–294.

Warren, C. M., Wilson, R. C., van der Wee, N. J., Giltay, E. J., van Noorden, M. S., Cohen, J. D., & Nieuwenhuis, S. (2017). The effect of atomoxetine on random and directed exploration in humans. *PloS One*, *12*, e0176034.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology: General*, *143*, 2074–2081.

Zajkowski, W. K., Kossut, M., & Wilson, R. C. (2017). A causal role for right frontopolar cortex in directed, but not random, exploration. *eLife*, *6*, e27430.

Zhang, S., & Yu, A. J. (2013). Forgetful Bayes and myopic planning: Human learning and decision-making in a bandit setting. In *Advances in neural information processing systems* (pp. 2607–2615).