Deep exploration as a unifying account of explore-exploit behavior

Robert C. Wilson^{a,b,*}, Siyu Wang^a, Hashem Sadeghiyeh^c,
& Jonathan D. Cohen^{d,e}

^aDepartment of Psychology, University of Arizona

^bCognitive Science Program, University of Arizona

^cDepartment of Psychological Science, Missouri University of Science and Technology

^dDepartment of Psychology, Princeton University

^ePrinceton Neuroscience Institute, Princeton University

*E-mail: bob@arizona.edu

Many decisions involve a choice between exploring unknown opportunities and exploiting well-known options. Work across a variety of domains, from animal foraging to human decision making, has suggested that animals solve such “explore-exploit dilemmas” with a mixture of two strategies: one driven by information seeking (directed exploration) and the other by behavioral variability (random exploration). Here we propose a unifying account in which these two strategies emerge from a kind of stochastic planning, known in the machine learning literature as Deep Exploration. In this model, the explore-exploit decision is made by stochastic simulation of plausible futures that are deep, in that they extend far into the future, and narrow, in that the number of possible futures they consider is small. By applying Deep Exploration to a simple explore-exploit task we show theoretically how directed and random exploration can emerge in these settings. Moreover, we show that Deep Exploration implies a tradeoff between directed and random exploration that is mediated by the number of simulations, or samples — with more samples leading to increased directed exploration and decreased random exploration at the expense of greater time taken to respond. By measuring human behavior on the same simple task, we show that this reaction-time-mediated tradeoff exists in human behavior both between and within participants. We therefore suggest that Deep Exploration is a unifying account of explore-exploit behavior in humans.

When dining at a favorite restaurant, do you exploit the pizza you known and love, or try the ravioli just added to the menu? Exploiting the pizza guarantees a good meal, but is uninformative; you already knew the pizza was good. Alternatively, exploring the ravioli provides information about how tasty the ravioli is at this restaurant, but there’s no guarantee you will enjoy it. Beyond the epicurean, such explore-exploit dilemmas occur at all levels of decision making, from picking a TV show to watch to a finding a person to marry, and there are real advantages to solving it well — exploit too much and you’ll eat pizza for the rest of your life; explore too much and you will never

discover anything new that is good. Unfortunately, however, solving the explore-exploit dilemma optimally is intractable in all but the simplest settings (Bellman, 1956; Gittins, 1979), and so the question arises as to how agents balance exploration and exploitation in practice.

Work in psychology, ecology and neuroscience suggests that humans and non-human animals use two strategies for solving the explore-exploit dilemma (Krebs et al., 1978; Meyer and Shi, 1995; Banks et al., 1997; Ölveczky et al., 2005; Frank et al., 2009; Steyvers et al., 2009; Lee et al., 2011; Payzan-LeNestour and Bossaerts, 2012; Wilson et al., 2014; Ebitz et al., 2018; Costa et al., 2019; Schulz and Gershman, 2019). The first strategy is ‘directed exploration.’ In this strategy, choices are biased towards more informative options by an ‘information bonus’ that increases the relative value of unknown options. In the restaurant example, the known pizza receives no information bonus, because eating it tells us nothing new; while the unknown ravioli receives a large information bonus, because choosing it provides information about how good the ravioli is in this setting. All else being equal, the information bonus tips the balance in favor of exploring the unknown. While in some circumstances it is possible to compute the optimal value for the information bonus, this optimal computation comes at considerable computational cost and, in some cases it is not computable at all (Bellman, 1956; Gittins and Jones, 1979).

The second strategy is ‘random exploration.’ In this strategy, exploratory choices are driven by noise in the decision process that makes the highest value (i.e. exploit) option less likely to be chosen. In the restaurant example, one implementation of random exploration would be to add random decision noise to the values of the pizza and ravioli. In this way, even if the pizza has a higher value, the random noise will sometimes lead to exploration simply by chance. While this is unlikely to generate the optimal strategy for exploration (if and when that can be determined), it is less computationally costly and, in some circumstances, may be the only practical way to favor exploration (Watkins, 1989; Sutton and Barto, 2018).

Considered as above, the mechanisms of directed and random exploration seem distinct. However, recent work in machine learning has shown how directed and random exploration can arise naturally from a decision process based on stochastic simulation known as ‘Deep Exploration’ (Osband et al., 2016), sometimes also referred to as Posterior Sampling for Reinforcement Learning (Strens, 2000). In these algorithms, the explore-exploit choice is made by simulating a small number of (as low as 1) random, but plausible, future experiences in order to approximate the expected value of taking different actions. Once these simulated values are computed, the decision is made by picking the action with the highest simulated value.

Random exploration arises naturally from Deep Exploration because the simulations are stochastic. More subtly, and as outlined in the Supplementary Material, directed exploration arises from this model in the details of how the simulated choices play out. In general, the more simulations that are carried out, the better the estimate of the value of exploration and the closer the decision process approximates optimal directed exploration. Thus, Deep Exploration predicts a tradeoff between directed and random exploration that is mediated by the number of simulations used to compute values, with more simulations corresponding more directed and less random exploration, but at the expense of more computational effort and time taken to make the decision.

Figure 1: The Infinite Bandits Task showing the different trial types. In all cases a trial begins with the presentation of the exploit and explore options. The exploit option is indicated by the amount available for choosing it (in this case 64 points), and the explore option by two question marks, ‘??’. The location (left or right) of the two options is assigned randomly on each trial. Based on this information participants make a choice that takes a reaction time, RT. (A) If the participant chooses to exploit, the exploit value glows white for delay, d , indicating the choice and outcome, after which the boxes are marked with an ‘XX’ to indicate they are closed and cannot be chosen for the duration of the inter-trial interval (ITI). The next trial begins when the next pair of options is revealed. If the game does not end, then the exploit value stays the same. (B, C) If the participant chooses to explore, the value of the explore option is revealed, and then one of two things can happen on the next trial, depending on whether the outcome is worse (B) or better (C) than the exploit value. If the outcome is worse (B) then the exploit value remains the same on the next trial; if the outcome is better (C) then the exploit value for the next trial takes on the value of the explore option from the current trial. In this way, the value of the exploit option reflects the best outcome seen so far, and the choice is always between exploiting this best outcome and exploring a new one.

The Infinite Bandits Task

In this work, we tested the predictions of Deep Exploration in a simple explore-exploit task that we call the ‘Infinite Bandits Task’ (Figure 1). In this task, participants play a series of games in which they make repeated choices between two options — a fully known exploit option with a value between 1 and 100 points, corresponding to the best outcome they have seen so far; and an unknown explore option, corresponding to a random draw from a uniform distribution between 1 and 100. Each game lasts for a random number of trials such that there is a fixed probability — the hazard rate, $h = 0.1$ — that the game will end after every trial.

At the start of each game, the exploit value is set randomly between 1 and 100. After that, the game plays out in the following way: Choosing the exploit option (like eating a known pizza) yields the known outcome that was cued at the start of the trial. Choosing the explore option

(eating unknown ravioli) yields a random outcome, drawn uniformly from 1 to 100. If the explore outcome is greater than the current exploit value (if the ravioli is better than the pizza), then this higher value becomes the exploit option on the next trial (ravioli is the new favorite). If the explore outcome was less than the current exploit value (the ravioli is worse than the pizza), then this value is discarded and the exploit value stays the same for the next trial (pizza remains the favorite). In this way, the exploit option reflects the best outcome seen so far in this game, and the choice is always between exploiting this best outcome and exploring something new. Furthermore, because there is always a new explore option on every trial (a new item is added to the menu), our task is equivalent to choosing between infinitely many options with deterministic payouts, hence the name: the Infinite Bandits Task.

Deep Exploration predicts a tradeoff between directed and random exploration in the Infinite Bandits Task

The simplicity of the Infinite Bandits Task makes it amenable to formal analysis, allowing us to derive expressions for the choice probabilities of Deep Exploration. In particular, we assume that Deep Exploration proceeds by simulating the values of exploiting and exploring out to some time horizon, T , (which is a free parameter of the model, but in principle should be $1/h$ on average) when this particular game is simulated to its end. If the payoff of the known exploit option is c , then the value of choosing the exploit option (assuming we continue to pick it for all trials up to the horizon) is

$$Q_{\text{exploit}} = Tc \quad (1)$$

To simulate the value of the explore option, the algorithm first samples a *plausible* value of the unknown explore option, $\tilde{u}^i$, based on its knowledge of the underlying generative process — in this case a draw from the uniform distribution between 1 and 100. For each sample, i , the agent then computes the expected value of exploring, $\tilde{Q}_{\text{explore}}^i$, as though the simulated value $\tilde{u}^i$ were the actual payoff from the unknown option. If $\tilde{u}^i > c$ the algorithm continues to pick that option for the remainder of the simulated game, otherwise the algorithm switches to the exploit option. Thus, $\tilde{Q}_{\text{explore}}^i$ can be written as

$$\tilde{Q}_{\text{explore}}^i = \begin{cases} T\tilde{u}^i & \text{if } \tilde{u}^i > c \\ \tilde{u}^i + (T-1)c & \text{if } \tilde{u}^i \leq c \end{cases} \quad (2)$$

This mental simulation process is repeated n times, with independent samples for $\tilde{u}^i$ each time giving independent samples of $\tilde{Q}_{\text{explore}}^i$. These samples are then combined to compute the simulated expected value of exploring, Q_{explore} , as

$$Q_{\text{explore}} = \frac{1}{n} \sum_{i=1}^n \tilde{Q}_{\text{explore}}^i \quad (3)$$

The explore-exploit decision is then made by comparing Q_{explore} to the value of exploiting Q_{exploit} , with the algorithm exploring when $Q_{\text{explore}} > Q_{\text{exploit}}$ and exploiting otherwise.

When only one sample is used to compute Q_{explore} ($n = 1$), Deep Exploration leads to probability matching behavior. This is easily seen from Equations 1 and 2 where Q_{explore} will be less

than $Q_{exploit}$ when the sampled explore value $\tilde{u}^i$ is less than c . Thus, the model will exploit with probability $p(\tilde{u}^i < c)$, which, if $\tilde{u}^i$ comes from a uniform distribution, is given by

$$p_{exploit} = p(\tilde{u}^i < c) = \frac{c}{100} \quad (4)$$

This probability matching behavior (blue line in Figure 2) is a purely random strategy with no directed component. The indifference point of the choice curve is equal to 50.5, the mean of the uniform distribution over outcomes from the explore option, indicating that there is no information bonus in favor of the explore option. In addition, the $n = 1$ strategy is independent of horizon, T . Such a probability matching strategy for exploration is known in the machine learning literature as Thompson sampling (Thompson, 1933). In psychology, probability matching has been observed in many tasks and across species (Vulkan, 2000), and Thompson sampling models can account for some properties of human random exploration (Speekenbrink and Konstantinidis, 2015; Gershman, 2018).

When infinitely many samples are used ($n = \infty$), Deep Exploration converges to the optimal strategy in the Infinite Bandits Task (see Supplementary Material). This is the yellow line in Figure 2. In this case

$$p_{exploit} = \begin{cases} 1 & \text{if } c > \theta^* \\ 0 & \text{if } c \leq \theta^* \end{cases} \quad (5)$$

where θ^* corresponds to the threshold between exploration and exploitation. This strategy is a purely directed strategy with an indifference point of $\theta^* = 75.97$ points (i.e. much larger than the expected value from a single play of the explore option, which is 50.5), but no random exploration because the behavior is deterministic. Moreover, this strategy is highly dependent on the simulation horizon, T , with the threshold between explore and exploit behavior changing as a function of T as

$$\theta^* = 100 \times \left(\frac{\sqrt{T}}{1 + \sqrt{T}} \right) \quad (6)$$

More generally (for $0 < n < \infty$) it is possible to compute an approximate analytic, albeit somewhat opaque, expression for the probability of exploiting as a function of n and T .

$$p_{exploit}(n, T) \approx \sum_{n_+=0}^n \binom{n}{n_+} \left(\frac{m-c}{m} \right)^{n_+} \left(\frac{c}{m} \right)^{n-n_+} \frac{1}{n_+!} \sum_{k=0}^{\left\lfloor \frac{c(n-n_+)}{2T(m-c)} \right\rfloor} (-1)^k \binom{n_+}{k} \left(\frac{c(n-n_+)}{2T(m-c)} - k \right)^{n_+} \quad (7)$$

As shown in Figure 2A, as the number of samples n increases, behavior becomes more deterministic (less random exploration) and more biased towards exploring (more directed exploration). Thus, assuming that sampling is costly in some form (in computational effort and/or in time), then there is a tradeoff between directed and random exploration that is mediated by the number of samples. In addition, it should be noted that, although more samples leads to better overall performance (Figure 2B), in this particular case the overall improvement in performance is small, approximately 5% between 1 and infinitely many samples. Thus, given that there is a cost to sampling, the agent should favor a small number of samples. More generally, to the extent that the cost of sampling can be quantitatively related to the gains afforded by more sampling, then the formalism can be used to determine an optimal tradeoff between random (minimal sampling) and directed (maximal sampling) exploration.

To illuminate the tradeoff between directed and random exploration more explicitly, we quantified the Deep Exploration choice curves in terms of two parameters: the indifference point, θ , at which the model is equally likely to explore or exploit, and the decision noise, σ , which corresponds to the inverse of the slope of the choice curve at the indifference point (Figure 2C). More directed exploration corresponds to a higher θ , while more random exploration corresponds to a higher σ (Figure 2D).

Figure 2: Behavior of the algorithm as a function of the number of samples, n . (A) Solid lines correspond to the analytic approximation, while the dots correspond to the average behavior of 10,000 simulations. As n increases, the model shows more directed exploration (choice curves shift to the right), and less random exploration (choice curves get steeper). As $n \rightarrow \infty$, the algorithm approaches the behavior of the optimal model. (B) Performance of the model as a function of the number of samples. Increasing the number of samples increases the average number of points earned in the task, but even with one sample, the performance of Deep Exploration is within 6% of optimal. (C) Illustration of indifference point, θ , which captures directed exploration, and slope σ , which captures random exploration. (D) Theoretical tradeoff between directed and random exploration for a given value of hazard rate $h = 0.1$. Note colored stars correspond to different number of samples as given by legend in panel A.

Between-subject tradeoff between directed and random exploration

To test whether the theoretically predicted tradeoff between directed and random exploration is observed in human behavior, 174 participants (ages 18-25, mean 19.1; 80 identified as male, 93 identified as female, 1 identified as transgender) performed the Infinite Bandits Task in the laboratory. In the simplest analysis, we quantified each person’s indifference point, θ , and noise, σ , using a logistic regression model (see Methods). As shown in Figure 3A, we found a strong negative correlation between θ and σ (Pearson’s $r(172) = -0.48$, $p < 0.001$), consistent with a tradeoff between directed and random exploration across the population. Crucially, the negative correlation between θ and σ is robust to a wide variety of modeling assumptions and is not induced as an artifact of the fitting procedure (see Supplementary Figures S5, S6, and S7). This suggests that the tradeoff between directed and random exploration is a real feature of the data.

To further investigate the tradeoff between directed and random exploration, we fit the Deep Exploration model (Equation 7) to behavior. As shown in Figure 3A (insets on the right), the model provides an excellent fit to the behavior of individual participants (see Supplementary Figure S10 for additional fits), and allows us to estimate the average number of simulations each participant used to make their decision (Figure 3B). In particular, we find that most participants used only a small number of simulations (81% between 1 and 5).

Remarkably, the fit number of simulations correlates with overall reaction time across participants ($r(172) = 0.23$, $p = 0.008$, Figure 3C), such that individuals who use more samples are slower overall. Furthermore, by estimating the slope of the line in Figure 3C, we can compute a crude estimate of the rate at which samples are produced, giving us an estimate of 61ms per sample, a number which is broadly consistent with previous work on the speed of working memory search (38ms per symbol (Sternberg, 1966)), preplay in hippocampus (Redish, 2016; Kay et al., 2019), and more generally has been proposed as a fundamental timescale for cognitive processing Newell (1994).

Within-subject tradeoff between directed and random exploration

Explaining the tradeoff between directed and random exploration in terms sampling directly implies that there should be a speed-accuracy tradeoff in explore-exploit decisions: taking fewer samples is faster (as shown in Figure 3D above), but provides a less accurate estimate of the optimal information bonus (i.e., favors random exploration), whereas taking more samples takes more time but improves estimates of the optimal information bonus (i.e., directed exploration) and thereby performance. This suggests that if participants are given more time to respond, they should be more likely to use greater numbers of samples and thereby exhibit greater evidence of directed exploration.

To test this prediction we ran a version of the Infinite Bandits Task in which we manipulated the time *between* trials (inter-trial interval, ITI) from fast (0.5 seconds) to slow (3 seconds) in two blocks of 500 trials each (order counterbalanced across participants). Such block-wise ITI manipulations, while changing nothing about the reward structure of the task, are known to cause changes in speed-accuracy tradeoffs, such that people tend respond more slowly when given more time to respond in a task with longer ITIs (Bogacz et al., 2006). Consistent with previous work, we

Figure 3: Between-subject tradeoff of directed versus random exploration in the Infinite Bandits Task. (A) Explore threshold and decision noise are negatively correlated, consistent with a tradeoff between directed and random exploration. Choice curves for two example participants, one showing more directed than random exploration (subject 146, estimated number of samples $n = 19.0$) and the other showing more random than directed exploration (subject 46, $n = 1.2$). The solid lines correspond to simulated behavior from the Deep Exploration model. (B) Average choice curves (dots) and Deep Exploration behavior (lines) for participants grouped according to the best fitting number of samples. (C) Number of simulations used by each subject showing that most participants use between 1 and 5 simulations to make their decision. (D) participants who use more samples respond more slowly. Each additional sample increases reaction time by approximately 50ms.

Figure 4: Within-subject tradeoff between directed and random exploration. Slowing the inter-trial interval increases directed exploration (A; significantly more points above the dashed identity line designating no change) and decreases random exploration (B; significantly more points below the line). In addition, the change in directed exploration is negatively correlated with the change in random exploration (C) as though the two were trading off. (D) Fitting the Deep Exploration model to this behavior shows that the change in directed and random is explained by an increase in the number of samples used in the slow condition (more points above the identity line).

found that prolonging the ITI lead to an increase in reaction time in 74% of participants (average change = 114ms, $t(75) = 5.08$, $p < 0.001$). Consistent with the predictions of Deep Exploration, this increase in reaction time was associated with an increase in directed exploration ($t(75) = 2.55$, $p = 0.013$) and a decrease in random exploration ($t(75) = -3.01$, $p = 0.0036$) in the slow condition. Moreover, the change in directed exploration was negatively correlated with the change in random exploration ($r(74) = -0.55$, $p < 0.001$), consistent with the idea that the two types of exploration trade off against one another.

Finally, by fitting the Deep Exploration model to behavior, we investigated whether people actually used more samples in the slow ITI condition. As shown in Figure 4D, this is indeed what we find with the average number of samples increasing by 1.4 (t-test on $\log n$, $t(75) = 3.41$, $p = 0.0011$). This suggests that the tradeoff between the two types of exploration can be explained in terms of a single underlying variable: number of samples.

Discussion

In this article we have proposed a unifying account of two different types of exploration — directed and random — that have previously been described for explore-exploit decision making behavior. In our Deep Exploration account, we propose that explore-exploit decisions are made by simulating a small number of plausible futures and choosing the action that maximizes their simulated future value. We showed that directed and random exploration emerge naturally from this account and, further, that Deep Exploration implies a tradeoff between the two types of exploration. Using a simple behavioral task, we found evidence for this tradeoff both between and within participants, such that individuals who take more time to make their decisions exhibit more directed and less random exploration. Furthermore, greater time taken is associated with a greater number of samples taken, as estimated by a fit of the Deep Exploration model, with a time cost of approximately 61ms per sample.

Beyond the Infinite Bandits Task, we show in the Supplementary Material that Deep Exploration can account for a range of previously reported behavioral findings in the explore-exploit literature; that directed and random exploration both: increase with horizon (Wilson et al., 2014); depend on uncertainty (Speekenbrink and Konstantinidis, 2015; Gershman, 2018); and are reduced when future information is not contingent on participants’ present choice (Findling et al., 2019). Moreover, because of the simplicity of the algorithm — all that is needed to implement Deep Exploration is the ability to simulate a small number of plausible futures — Deep Exploration is computationally tractable even in complex cases such as those found in the real world. In this regard, Deep Exploration stands in contrast to optimal (and even many heuristic) exploration policies that quickly become intractable for problems of even modest complexity Bellman (1956). This makes it appealing not only as an account of how people manage the explore-exploit tradeoff, but also as an efficient algorithm in artificial systems.

More generally, our model accords with work showing that mental simulation (including sampling) may be involved in a variety of judgments and decisions (Stewart et al., 2006; Vul et al., 2014). Perhaps most prominent in this literature is the work on intuitive physics, where judgments about the stability of a tower of blocks, or the path of a ball on a pool table, are well described by mental simulation models (Kubricht et al., 2017). Moreover, as we found above, human behavior in these tasks is best described by models using only a small number of simulations (between 1 and 10, (Battaglia et al., 2013; Hamrick et al., 2016)), with the exact number depending in part on the individual but also on the demands of the task (Hamrick et al., 2015). Thus mental simulation may be a general strategy for judgment and decision making, as well as an important factor underlying individual differences in these behaviors.

References

- Banks, J., Olson, M., and Porter, D. (1997). An experimental analysis of the bandit problem. *Economic Theory*, 10(1):55–77.
- Battaglia, P. W., Hamrick, J. B., and Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45):18327–18332.

- Bellman, R. (1956). A problem in the sequential design of experiments. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 16(3/4):221–229.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4):700.
- Costa, V. D., Mitz, A. R., and Averbeck, B. B. (2019). Subcortical substrates of explore-exploit decisions in primates. *Neuron*.
- Ebitz, R. B., Albarran, E., and Moore, T. (2018). Exploration disrupts choice-predictive signals and alters dynamics in prefrontal cortex. *Neuron*, 97(2):450–461.
- Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S., and Wyart, V. (2019). Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature neuroscience*, pages 1–12.
- Frank, M. J., Doll, B. B., Oas-Terpstra, J., and Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature neuroscience*, 12(8):1062.
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173:34–42.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164.
- Gittins, J. C. and Jones, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3):561–565.
- Hamrick, J. B., Battaglia, P. W., Griffiths, T. L., and Tenenbaum, J. B. (2016). Inferring mass in complex scenes by mental simulation. *Cognition*, 157:61–76.
- Hamrick, J. B., Smith, K. A., Griffiths, T. L., and Vul, E. (2015). Think again? the amount of mental simulation tracks uncertainty in the outcome. In *CogSci*. Citeseer.
- Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., Liu, D. F., and Frank, L. M. (2019). Constant sub-second cycling between representations of possible futures in the hippocampus. *bioRxiv*, page 528976.
- Krebs, J. R., Kacelnik, A., and Taylor, P. (1978). Test of optimal sampling by foraging great tits. *Nature*, 275(5675):27.
- Kubricht, J. R., Holyoak, K. J., and Lu, H. (2017). Intuitive physics: Current research and controversies. *Trends in cognitive sciences*, 21(10):749–759.
- Lee, M. D., Zhang, S., Munro, M., and Steyvers, M. (2011). Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*, 12(2):164–174.

- Meyer, R. J. and Shi, Y. (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Science*, 41(5):817–834.
- Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.
- Ölveczky, B. P., Andalman, A. S., and Fee, M. S. (2005). Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS biology*, 3(5):e153.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. (2016). Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, pages 4026–4034.
- Payzan-LeNestour, É. and 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*, 6:150.
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3):147.
- Schulz, E. and Gershman, S. J. (2019). The algorithmic architecture of exploration in the human brain. *Current opinion in neurobiology*, 55:7–14.
- Speekenbrink, M. and Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2):351–367.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153(3736):652–654.
- Stewart, N., Chater, N., and Brown, G. D. (2006). Decision by sampling. *Cognitive psychology*, 53(1):1–26.
- Steyvers, M., Lee, M. D., and Wagenmakers, E.-J. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179.
- Strens, M. (2000). A bayesian framework for reinforcement learning. In *ICML*, volume 2000, pages 943–950.
- Sutton, R. S. and Barto, A. G. (2018). *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(3/4):285–294.
- Vul, E., Goodman, N., Griffiths, T. L., and Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive science*, 38(4):599–637.
- Vulkan, N. (2000). An economists perspective on probability matching. *Journal of economic surveys*, 14(1):101–118.
- Watkins, C. J. C. H. (1989). Learning from delayed rewards.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., and Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074.

Methods

Participants

All participants were recruited from the University of Arizona Psychology subject pool (217 for Experiment 1 and 85 for Experiment 2) and the Princeton student population (27 participants for Experiment 2). Participants were unpaid but received course credit for taking part in the experiment. All participants gave informed consent and the study was approved by the Institutional Review Boards at the University of Arizona and Princeton.

Experiment 1: Baseline Infinite Bandits Task

In Experiment 1, we aimed to measure the tradeoff between directed and random exploration across the population. To this end 217 participants (100 male, 116 female, 1 transgender, ages 18-25, mean 19.1) performed the Infinite Bandits Task. In this version of the task, participants who completed the task performed 500 trials with a fixed hazard rate of $h = 0.1$. Participants who completed fewer than 400 trials were excluded from further analysis.

Experiment 2: Timing manipulation

In Experiment 2, we aimed to manipulate the directed-random tradeoff by changing the timing in the task to favor fast or slow responding. The motivation for this manipulation was that slower responding should lead to people using more mental simulations to make their decision. This, in turn, should increase directed exploration and decrease random exploration. To bias faster or slower responding we changed the inter-trial interval (ITI) in the experiment from 0.5 to 3 seconds while keeping the hazard rate fixed. Such changes in ITI have been previously shown to induce faster and slower responding without explicitly instructing participants to speed up or slow down (Bogacz et al., 2006). 112 participants (35 male, 74 female, 3 declined to answer, ages 18-25, mean age = 19.1) played two blocks of the Infinite Bandits Task. Across blocks we varied the inter-trial interval (ITI) from 0.5 to 3 seconds. The hazard rate was kept constant across the two blocks (at $h = 0.1$) and the order of the ITI conditions was counterbalanced across participants. Participants who completed the task (see below) were presented with either 500 or 600 trials in each block.

Basic performance and exclusion criteria

Participants were excluded from further analysis if they did not complete the required number of trials or if they did not perform the task better than chance. For trials, in Experiment 1, participants were excluded if they performed fewer than 400 trials. In Experiment 2, participants were excluded if they performed fewer than 400 trials in either of the two blocks.

For performance, in both experiments we excluded participants on the basis of a normalized reward score. In particular, we took the mean reward per trial, RPT , for each subject in each block of the experiment and normalized this relative to optimal RPT_{opt} and random performance RPT_{rand} , which were computed according to Equations S28 and S29 in the Supplementary material. Thus

we computed the normalized reward as

$$NR = \frac{RPT - RPT_{rand}}{RPT_{opt} - RPT_{rand}} \quad (8)$$

This normalized reward has value 1 for optimal performance and 0 for random performance. As a conservative threshold, we excluded participants whose normalized reward was less than 0. Thus, in Experiment 1, participants were excluded if their normalized reward was less than 0 overall in Experiment 1, or less than 0 in either block in Experiment 2. With these exclusion criteria, as shown in Figure 5, we excluded 43 participants from Experiment 1 and 36 participants from Experiment 2. This left 174 participants in Experiment 1 (80 male, 93 female, 1 transgender, ages 18-25, mean age = 19.1) and 76 participants in Experiment 2 (24 male, 51 female, 1 declined to answer, ages 18-22, mean age = 19.2).

Figure 5: Normalized performance and exclusion criteria. The normalized reward score is 1 for performance on a par with the optimal model and 0 for purely random performance. (A) In Experiment 1, participants were excluded if they achieved a normalized reward of less than 0, or if they performed fewer than 400 trials. (B) In Experiment 2, participants were excluded if they achieved a normalized reward of less than 0 in either the fast ITI or slow ITI block, or if they performed fewer than 400 trials in either block.

Descriptive model

We quantified behavior on the Infinite Bandits Task using a series of descriptive models that quantified the overall level of directed and random exploration as well as accounting for possible confounding factors such as lapsing and sequential effects.

Basic form of the θ - σ model

In its most basic form, the descriptive model assumes that participants make choices by assigning different values to each of possible options (either explore and exploit). We assume that the exploit option has a value equal to its cued value, c , and the explore option has a value given by the ‘explore threshold’, θ , that corresponds to the expected value of the explore option plus the information bonus associated with choosing it. Next we assume that participants make a decision based on a noisy difference between these two values,

$$\Delta = \theta - c + \sigma n \quad (9)$$

where n is zero mean logistic noise with variance 1, and σ describes the magnitude of the decision noise. Finally we assume that participants choose to explore if $\Delta > 0$, exploit if $\Delta < 0$ and choose randomly if $\Delta = 0$. Thus on average (over the decision noise) the model assumes that participants choose to exploit with probability

$$p_{\text{exploit}} = \frac{1}{1 + \exp\left(\frac{\theta - c}{\sigma}\right)} \quad (10)$$

The choice probability for the ‘ θ - σ model’ is plotted in Figure 6A as function of exploit value c and shows how the parameters θ and σ relate to the choice curve.

Figure 6: Relationship between model parameters and choice curves. As the exploit value increases the model is more likely to choose the exploit option. (A) In the θ - σ model, the explore threshold, θ , (in this case $\theta = 70$) determines the exploit value at which the model is indifferent between the explore and exploit options. The decision noise, σ , (in this case $\sigma = 10$) determines the slope of the choice curve at the indifference point. (B) In the θ - σ - ϵ model, $\epsilon_1 = 0.3$, $\epsilon_2 = 0.15$

Variants of the θ - σ model to include lapsing and sequential effects

To account for the possibility of factors other than directed and random exploration driving behavior, we extended the θ - σ model to incorporate lapsing and sequential effects.

Lapsing occurs when participants choose randomly, with no regard to the exploit value c . Lapses are most apparent at the far asymptotes of the choice curve and is captured by two parameters (Figure 6B): ϵ_1 , for the lapse rate for exploiting (left asymptote of curve for small c), and ϵ_2 the lapse rate for exploring (right asymptote of curve for large c). Such lapsing is a common feature of many decisions (e.g. Busse et al., 2011; Gold and Ding, 2013; Carandini and Churchland, 2013) and models that include two lapse rates have been shown to fit data better and provide better estimates of other parameters such as θ and σ (Wichmann and Hill, 2001; Prins and Kingdom, 2018).

Sequential effects occur when information from past trials affects the decision on the present trial. Here we consider three such factors: the effect of past exploration (e_{past} , whether the last choice was to explore or exploit), the effect of past physical actions (a_{past} , whether the last choice left or right), and the effect of past change points (cp_{past} , whether the last trial was the start of a new game). We assume that these sequential effects combine linearly with weights β_e for past exploration, β_a for past physical action, and β_{cp} for past change points.

Combining lapses and sequential effects gives us the following general form for choice probabilities

$$p_{exploit} = (1 - \epsilon_1 - \epsilon_2) \frac{1}{1 + \exp\left(\frac{\theta - c + \beta_e e_{past} + \beta_a a_{past} + \beta_{cp} cp_{past}}{\sigma}\right)} + \epsilon_1 \quad (11)$$

The free parameters of this extended model are described in Table 1.

Parameter	Description	Range
θ	Explore threshold	0 to 100
σ	Decision noise	0 to ∞
ϵ_1	Lapse rate (exploit)	0 to 1
ϵ_2	Lapse rate (explore)	0 to 1
β_e	Previous action (explore or exploit) weight	$-\infty$ to ∞
β_a	Previous action (left or right) weight	$-\infty$ to ∞
β_{cp}	Previous change-points	$-\infty$ to ∞

Table 1: Parameters in the θ - σ variant models

By including and excluding the factors ϵ_1 , ϵ_2 , β_e , β_a , and β_{cp} in Equation 11, we created a suite of different models. For the effects of past side and past change points we only included the effect of the last trial, while for the effect of past explore/exploit choice we included the effects of up to two trials into the past. This lead to 48 different forms of the model generated from the presence or absence of ϵ_1 , ϵ_2 , β_e , and β_{cp} ($2^4 = 16$ combinations), and the absence or presence of 1 or 2 trials back from β_e (giving $16 \times 3 = 48$ combinations).

Deep Exploration model

As with the Descriptive model above, to fit the Deep Exploration model to behavior, we augmented the basic choice probability (Equation 7) to include the possibility of lapsing and sequential effects. That is, we assumed the probability of exploiting was given by

$$\tilde{p}_{exploit} = (1 - \epsilon_1 - \epsilon_2) \frac{1}{1 + \exp(\text{logit}(p_{exploit}(n, T)) + \beta_e e_{past} + \beta_a a_{past} + \beta_{cp} cp_{past})} + \epsilon_1 \quad (12)$$

where $p_{exploit}(n, T)$ was given by equation 7. Note that for the purposes of fitting, instead of parameterizing the model through the simulation horizon T , we used the implied hazard rate, $h = 1/T$. While this substitution does not change the properties of the model, it has the advantage (especially for visualization purposes) that h is between 0 and 1 while T is unconstrained.

Model fitting

All models were fit using a maximum likelihood approach. In particular, each model gives us a closed form expression for the probability of choosing the exploit option, $p_{exploit}$, as a function of parameters, ϕ , which we then use to compute the likelihood of the observed actions, $a_{1:t}$, as

$$p(a_{1:t}|\phi) = \prod_{i=1}^t p(a_i|\phi, a_{1:i-1}) \quad (13)$$

where

$$p(a_i|\phi, a_{1:i-1}) = \begin{cases} p_{exploit} & a_i = \text{exploit} \\ 1 - p_{exploit} & a_i = \text{explore} \end{cases} \quad (14)$$

Maximum likelihood parameter fits, $\hat{\phi}$, were determined by maximizing the log likelihood

$$\hat{\phi} = \operatorname{argmax}_{\phi} \log p(a_{1:t}|\phi) \quad (15)$$

Model comparison

Model fits were compared using Bayes Information Criterion scores (BIC, Schwarz et al., 1978) computed as

$$BIC = k \log T - 2 \log \hat{LL} \quad (16)$$

where k is the number of free parameters in the model, T is the number of trials and $\hat{LL}$ is the maximum log likelihood.

For each model, we computed the number of participants best fit by that model as well as the exceedance probability (Rigoux et al., 2014). This latter measure estimates the probability that each model generated the data from all participants under the assumption that all participants used the same model.

Methods References

- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4):700.
- Busse, L., Ayaz, A., Dhruv, N. T., Katzner, S., Saleem, A. B., Schölvinck, M. L., Zaharia, A. D., and Carandini, M. (2011). The detection of visual contrast in the behaving mouse. *Journal of Neuroscience*, 31(31):11351–11361.
- Carandini, M. and Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature neuroscience*, 16(7):824.

- Gold, J. I. and Ding, L. (2013). How mechanisms of perceptual decision-making affect the psychometric function. *Progress in neurobiology*, 103:98–114.
- Prins, N. and Kingdom, F. A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the palamedes toolbox. *Frontiers in psychology*, 9.
- Rigoux, L., Stephan, K. E., Friston, K. J., and Daunizeau, J. (2014). Bayesian model selection for group studies revisited. *Neuroimage*, 84:971–985.
- Schwarz, G. et al. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.
- Wichmann, F. A. and Hill, N. J. (2001). The psychometric function: I. fitting, sampling, and goodness of fit. *Perception & psychophysics*, 63(8):1293–1313.

Acknowledgements

This project was made possible through the support of National Institute on Aging grants R56AG061888 and R01AG061888 to RCW and a grant from the John Templeton Foundation to JDC. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the funders.

Author contributions

RCW and JDC conceived the experiment. HS and SW acquired the data. RCW derived the theory and performed the analyses with help from SW and HS. RCW and JDC wrote the paper with input from SW and HS.

Supplementary Material: Deep exploration as a unifying account of explore-exploit behavior

Robert C. Wilson^{a,b,*}, Siyu Wang^a, Hashem Sadeghiyeh^c,
& Jonathan D. Cohen^{d,e}

Contents

1	Optimal strategy for the Infinite Bandits Task	2
2	Deep Exploration algorithm for the Infinite Bandits Task	3
2.1	Choice curves for $n = 1$	3
2.2	Choice curves for $n = 2$	4
2.3	Approximate choice curves for general n	5
3	Evaluating performance of different strategies on the Infinite Bandits Task	7
4	Descriptive model	8
4.1	Model comparison for descriptive model variants	8
4.2	Correlation between directed and random exploration is conserved across models .	8
4.3	The five parameter descriptive model	10
4.4	Parameter recovery for the five-parameter descriptive model	10
5	Fitting the Deep Exploration model to behavior	11
5.1	Model comparison	12
5.2	Parameter recovery	12
5.3	Comparison between simulated and actual choice curves	12
6	Deep exploration can account for horizon, uncertainty, and feedback dependence of directed and random exploration	16
6.1	The Horizon Task and its variants	16
6.1.1	Horizon Task	16
6.1.2	Feedback Task	17
6.1.3	Parametric Uncertainty Task	19
6.2	Optimal model for Horizon Task variants	20
6.3	Deep Exploration algorithm in the Horizon Task variants	22
6.3.1	Deep exploration can account for behavior in Horizon Task variants	23

1 Optimal strategy for the Infinite Bandits Task

The optimal strategy in the Infinite Bandits Task is a threshold policy whereby one should always exploit when the known option is greater than some threshold θ^* , otherwise one should explore. To compute the value of the optimal threshold we imagine a situation where the exploit value c exactly equals the threshold, θ^* . Also, to begin, we assume that the game length, T , is fixed. In this case, the value of choosing the exploit option is simply

$$Q_c(T) = Tc = T\theta^* \quad (\text{S1})$$

This is because if one exploits on the first choice one should exploit to the end of the game because, by definition the value of exploration and exploitation are equal when $c = \theta^*$.

We can also compute the value of exploration in the case where $c = \theta^*$ as

$$\begin{aligned} Q_u(T) &= \overbrace{Tp(u > \theta^*)\mathbb{E}(u|u > \theta^*)}^{\text{good outcome } u > \theta^*} + \overbrace{p(u \leq \theta^*)(\mathbb{E}(u|u \leq \theta^*) + (T-1)\theta^*)}^{\text{bad outcome } u \leq \theta^*} \\ &= T \left(\frac{m - \theta^*}{m} \right) \left(\frac{m + \theta^*}{2} \right) + \frac{\theta^*}{m} \left(\frac{\theta^*}{2} + (T-1)\theta^* \right) \\ &= \frac{T(m^2 - \theta^{*2})}{2m} + \frac{\theta^{*2}}{2m} + \frac{(T-1)\theta^{*2}}{m} \\ &= T \frac{m}{2} - (T-1) \frac{\theta^{*2}}{2m} + (T-1) \frac{\theta^{*2}}{m} \\ &= T \frac{m}{2} + (T-1) \frac{\theta^{*2}}{2m} \\ &= T\bar{u} + (T-1) \left(\frac{\theta^{*2}}{2m} \right) \end{aligned} \quad (\text{S2})$$

where $\bar{u}$ is the average reward for choosing the explore option.

Taking the average over T we get the average value associated with this policy when $c = \theta^*$ as

$$\begin{aligned} Q_u &= \sum_{T=1}^{\infty} p(T) \left(T\bar{u} + (T-1) \left(\frac{\theta^{*2}}{2m} \right) \right) \\ &= \bar{T}\bar{u} + (\bar{T}-1) \left(\frac{\theta^{*2}}{2m} \right) \end{aligned} \quad (\text{S3})$$

where $\bar{T} = 1/h$. Likewise, the expected value of exploiting is

$$Q_c = \bar{T}\theta^* \quad (\text{S4})$$

By definition $Q_c = Q_u$ when $c = \theta^*$ so we have

$$\begin{aligned} \bar{T}\bar{u} + (\bar{T}-1) \left(\frac{\theta^{*2}}{2m} \right) &= \bar{T}\theta^* \\ \implies \left(\frac{\bar{T}-1}{2m} \right) \theta^{*2} - \bar{T}\theta^* + \bar{T}\bar{u} &= 0 \end{aligned} \quad (\text{S5})$$

Solving for θ^* we get

$$\theta^* = m \left(\frac{\sqrt{T}}{1 + \sqrt{T}} \right) = \frac{m}{1 + \sqrt{h}} \quad (\text{S6})$$

which for the experiment, where $h = 0.1$ and $m = 100$, gives $\theta^* = 75.97$ points.

2 Deep Exploration algorithm for the Infinite Bandits Task

A simple deep exploration algorithm for the Infinite Bandits Task proceeds in the following way. First sample n plausible values for the unknown bandit, $\tilde{\mathbf{u}} = (\tilde{u}^1, \tilde{u}^2, \dots, \tilde{u}^n)$. If the agent knows the prior is a uniform between 0 and m then this is done by sampling each $\tilde{u}^i$ as $\tilde{u}^i \sim U(0, m)$. For each sample, i , the agent then computes the expected value of exploring, $\tilde{Q}_u^i$, as though $\tilde{u}^i$ were the actual payoff from the unknown option. This expected value can be written as

$$\tilde{Q}_u^i = \begin{cases} T\tilde{u}^i & \text{if } \tilde{u}^i > c \\ \tilde{u}^i + (T-1)c & \text{if } \tilde{u}^i \leq c \end{cases} \quad (\text{S7})$$

Dividing by T we get the expected value per trial for this sample

$$\tilde{q}_u^i = \begin{cases} \tilde{u}^i & \text{if } \tilde{u}^i > c \\ \frac{1}{T}(\tilde{u}^i + (T-1)c) & \text{if } \tilde{u}^i \leq c \end{cases} \quad (\text{S8})$$

Finally, we compute the approximate expected value, $\tilde{q}_u$, by averaging over the n samples,

$$\tilde{q}_u = \frac{1}{n} \sum_{i=1}^n \tilde{q}_u^i \quad (\text{S9})$$

The decision is made by comparing $\tilde{q}_u$ to the value of exploiting c , with the algorithm exploring when $\tilde{q}_u > c$ and exploiting otherwise.

Behavior of the algorithm for different values of n for fixed horizon, $T = 16$, is shown in Figure S3. As the number of samples increases the algorithm gradually converges to the optimal choice curve. More insight into this behavior can be obtained by computing analytic forms for the choice curves.

2.1 Choice curves for $n = 1$

When $n = 1$ the $\tilde{q}_u$ is based on just one sample. From equation S8, we can easily see that the probability that the $\tilde{q}_u$ is less than c is equal to the probability that the single sample u^1 is less than c ; i.e.

$$p_{\text{exploit}} = p(\tilde{q}_u < c) = p(u < c) = \frac{c}{m} \quad (\text{S10})$$

Thus the $n = 1$ case is a simple *probability matching* strategy whereby the algorithm choose each option with the probability that it will be the best. Note that this strategy has only random exploration with the indifference point being equal to the mean of the prior, $m/2$. In addition, the $n = 1$ strategy is independent of horizon, T . In the bandits literature such a probability matching strategy is known as Thompson sampling (Thompson, 1933) and it is a general result that the $n = 1$ case of Deep Exploration is equivalent to Thompson sampling.

Figure S1: Probability distribution over the expected value of one sample, $p(\tilde{q}_u^i)$. From equation S8, this probability is T/m for $(T-1)c/T \leq \tilde{q}_u^i \leq c$, $1/m$ for $c < \tilde{q}_u^i \leq m$ and zero otherwise. From here it is easy to see that the probability of exploiting in the one-sample case is simply c/m .

2.2 Choice curves for $n = 2$

In this case we can compute an exact expression for the choice curves. To do this we need to determine the probability that the mean of two samples, $\frac{1}{2}(\tilde{q}_u^1 + \tilde{q}_u^2)$, is less than c . This can occur in two different ways. First, both samples could be less than c , that is, they could lie in the region 1 in figure S2. When this happens, the average of the two samples is guaranteed to be less than c and so the algorithm will always exploit when this happens. From equation S8 it is easy to see that both samples will be less than c with probability $(c/m)^2$.

The second way the mean of the samples could be less than c is if one sample was greater than c , the other was less than c and the sum of the two was less than $2c$, that is if the samples lie in one of the red triangle regions in figure S2. The probability of one sample being larger than c and on smaller than c is

$$p(u^1 > c)p(u^2 \leq c) = \left(\frac{m-c}{m}\right) \left(\frac{c}{m}\right) \quad (\text{S11})$$

The probability that the sum of the two samples is less than $2c$ depends on the size of c . If $c \leq mT/(T+1)$ then the probability is the relative size of red triangle to the rectangle 2 in panel A of figure S2. This is simply $c/2T(m-c)$. If $c > mT/(T+1)$ then the probability is the relative size of the red area to the rectangle in panel B of figure S2. This is $1 - c/2T(m-c)$.

Putting it all together this gives the probability of exploiting in the $n = 2$ case as

$$p_{\text{exploit}} = \begin{cases} \frac{c^2}{m^2} \left(\frac{T+1}{T}\right) & c \leq Tm/(T+1) \\ \frac{c^2}{m^2} \left(1 + \left(\frac{m-c}{c}\right) \left(2 - T \left(\frac{m-c}{c}\right)\right)\right) & c > Tm/(T+1) \end{cases} \quad (\text{S12})$$

From equation S12, we can compute analytic expressions for the indifference point as the value of c at which $p_{\text{exploit}} = 1/2$

$$\frac{c^*}{m} = \sqrt{\frac{T}{2(T+1)}} \quad (\text{S13})$$

Note that $c^* < Tm/(T+1)$ and so this expression holds regardless of T . In addition we can compute the inverse of the gradient of the choice curve at the indifference point, a proxy for the decision noise, as

$$\frac{\sigma^*}{m} = \frac{1}{m \left(\frac{dp_{\text{exploit}}}{dc}\right)_{c^*}} = \sqrt{\frac{T}{2(T+1)}} \quad (\text{S14})$$

Figure S2: Intuition for the choice-curve derivation for $n = 2$.

Note that c^* and σ^* have identical dependence on T . This means that for the two-sample case both directed and random exploration have the same horizon dependence. Note also that the T -dependence for c^* is different to the optimal model.

2.3 Approximate choice curves for general n

To compute choice curves for general n we will first compute an approximation for the choice curves when exactly n_+ of the samples have $u^i > c$ and $n_- (= n - n_+)$ of the samples have $u^i \leq c$. That is, we write

$$\begin{aligned}
 p_{\text{exploit}} &= p(\tilde{q}_u \leq c) \\
 &= \sum_{n_+=0}^n p(\tilde{q}_u(n_+) \leq c) p(n_+|c) \\
 &= \sum_{n_+=0}^n p_{\text{exploit}}^{n_+} p(n_+|c)
 \end{aligned} \tag{S15}$$

where $p_{\text{exploit}}^{n_+}$ is the probability of exploiting given that n_+ of the samples have $u^i > c$ and $p(n_+|c)$ is the probability of seeing n_+ samples with $u^i > c$, which is simply the binomial probability

$$p(n_+|c) = \binom{n}{n_+} \left(\frac{m-c}{m} \right)^{n_+} \left(\frac{c}{m} \right)^{n-n_+} \tag{S16}$$

To compute $p_{\text{exploit}}^{n_+}$, we first write

$$\tilde{q}_u(n_+) = \frac{1}{n} \sum_{i \in u^i > c} \tilde{q}_u^{i+} + \frac{1}{n} \sum_{i \in u^i \leq c} \tilde{q}_u^{i-} \tag{S17}$$

where $\tilde{q}_u(n_+)$ is the expected value given that n_+ of the samples have $u^i > c$, $\tilde{q}_u^{i+}$ is the expected value of a single sample given that $u^i > c$ and $\tilde{q}_u^{i-}$ is the expected value of a single sample given that $u^i \leq c$.

The point of the decomposition in equation S17 is that, unlike $\tilde{q}_u^i$, $\tilde{q}_u^{i+}$ and $\tilde{q}_u^{i-}$ are distributed *uniformly*. Thus, the sums in equation S17 are sums of uniformly sampled random variables and will follow an Irwin-Hall distribution. In addition, we note that $\tilde{q}_u^{i-}$ ranges between $c(T-1)/T$ and c , which for large T , becomes small. Thus we can approximate the second sum as

$$\sum_{i \in u^i \leq c} \tilde{q}_u^{i-} \approx \frac{(2T-1)(n-n_+)c}{2T} \quad (\text{S18})$$

Thus $p_{exploit}^{n_+}$ is approximately given by

$$p_{exploit}^{n_+} \approx p \left(\sum_{i \in u^i > c} \tilde{q}_u^{i+} > c \left(n - \frac{(2T-1)(n-n_+)}{2T} \right) \right) \quad (\text{S19})$$

Using the standard form of the cumulative distribution function of the Irwin-Hall distribution, after some algebra, this can be written as

$$p_{exploit}^{n_+} \approx \frac{1}{n_+!} \sum_{k=0}^{\lfloor \frac{c(n-n_+)}{2T(m-c)} \rfloor} (-1)^k \binom{n_+}{k} \left(\frac{c(n-n_+)}{2T(m-c)} - k \right)^{n_+} \quad (\text{S20})$$

Putting this all together we can write the probability of exploiting as

$$p_{exploit} \approx \sum_{n_+=0}^n \binom{n}{n_+} \left(\frac{m-c}{m} \right)^{n_+} \left(\frac{c}{m} \right)^{n-n_+} \frac{1}{n_+!} \sum_{k=0}^{\lfloor \frac{c(n-n_+)}{2T(m-c)} \rfloor} (-1)^k \binom{n_+}{k} \left(\frac{c(n-n_+)}{2T(m-c)} - k \right)^{n_+} \quad (\text{S21})$$

As shown in Figure S3, this expression provides a good approximation of the choice curves.

Figure S3: Behavior of the algorithm as a function of the number of samples, n . Solid lines correspond to the analytic approximation, while the dots correspond to the average behavior of 10,000 simulations. As n increases, the model shows more directed exploration, choice curves shift to the right, and less random exploration, choice curves get steeper. As $n \rightarrow \infty$, the algorithm approaches the behavior of the optimal model.

3 Evaluating performance of different strategies on the Infinite Bandits Task

In this section we outline how to compute the performance of a given strategy on the Infinite Bandits Task as used in Figure 2B in the main text. In particular, we show how to compute the expected number of points per trial, $V(\pi)$, for playing the task with strategy (or policy) π . Our approach for computing this value involves three steps. First we compute the average earnings per trial when the exploit value is c , $R(c|\pi)$. Next we compute the frequency with which trials with exploit value c occur, $F(c|\pi)$. Finally we compute $V(\pi)$ by averaging over c .

Step 1 - Computing the reward per trial, $R(c|\pi)$ Because an exploit choice always yields reward c and an explore choice yields an average reward of $\bar{u} = 50.5$ (the average of all possible reward values from 1 to 100), $R(c, \pi)$, is simply the average of c and $\bar{u}$ over the choice probability, $p(a|c, \pi)$; i.e.,

$$R(c|\pi) = cp(a = \text{exploit}|c, \pi) + \bar{u}\pi(a = \text{explore}|c, \pi) \quad (\text{S22})$$

Step 2 - Computing the frequency of trials with exploit value c , $F(c, \pi)$ This step is slightly more involved. First we write down the transition matrix, T , that is made up of elements $T_{c',c}$ denoting the probability of moving to an exploit value of c' on the next trial given that we have an exploit value c on the current trial. In particular, $T_{c',c}$ can be written in terms of the choice probability in Equation S21, $p(a|c, \pi)$, and the transition probability given the choice, $p(c'|c, a)$, as

$$T_{c',c} = \sum_a p(c'|c, a)\pi(a|c, \pi) \quad (\text{S23})$$

It is then straightforward to write down $p(c'|c, a)$ for the two different actions, exploit and explore.

For exploit choices, when there is no new game (which occurs with probability $1 - h$), the exploit value does not change on the next trial, i.e. $c' = c$. When there is a new game (which occurs with probability h) c' is equally likely to take on any value from 1 to 100 (including c). Thus we can write

$$p(c'|c, a = \text{exploit}) = \begin{cases} 1 - h + h/100 & c' = c \\ h/100 & c' \neq c \end{cases} \quad (\text{S24})$$

For explore choices, when there is no new game, the exploit value only stays the same ($c' = c$) if the explore value was less than or equal to c , which occurs with probability $c/100$. Otherwise, a given value of c' above c occurs with probability of $1/100$. Finally, when there is a new game, as for exploit choices, c' is equally likely to take on any value. Thus,

$$p(c'|c, a = \text{explore}) = \begin{cases} (1 - h)/100 + h/100 & c' > c \\ (1 - h)s/100 + h/100 & c' = c \\ h/100 & c' < c \end{cases} \quad (\text{S25})$$

Equations S24 and S25 allow us to compute the transition probability $T_{c',c}$, which in turn allows us to compute $F(c|\pi)$. In particular, since $F(c|\pi)$ is the stationary distribution over exploit values given the choice probabilities in Equation S21, it will be invariant under transformation by $\mathbf{T}$; i.e.

$$F(c'|\pi) = \sum_{c=1}^{100} T_{c',c}F(c|\pi) \quad (\text{S26})$$

This equation implies that $F(c|\pi)$ is the eigenvector of the matrix $\mathbf{T}$ with eigenvalue 1, something that is readily obtained by standard numerical methods (in our case using the `eigs` function in Matlab).

Step 3 - Computing the expected number of points per trial, $V(\pi)$ Now that we have $R(c|\pi)$, the expected reward given the exploit value is c , and $F(c|\pi)$, the frequency with which trials with exploit value c occur, it is straightforward to write down the expected performance, $V(\pi)$, as the average over s ; i.e.,

$$V(\pi) = \sum_{s=1}^{100} F(s|\pi) R(s|\pi) \quad (\text{S27})$$

In this way we can compute the reward per trial for optimal and random performance used in the normalized reward (Equation 8 in the main paper). In particular, for optimal performance we have

$$RPT_{opt} = V(\pi = \text{optimal}) = 79.35 \text{ points} \quad (\text{S28})$$

where the optimal policy is a deterministic threshold policy with threshold equal to 75.97. Similarly, for random performance

$$RPT_{rand} = V(\pi = \text{random}) = 67.34 \text{ points} \quad (\text{S29})$$

where the random policy is to choose each option with 50% probability.

4 Descriptive model

4.1 Model comparison for descriptive model variants

We fit all 48 combinations of the descriptive model along with a purely random model that assumed a fixed probability b for exploring or exploiting. Model performance was compared using BIC as described above. For each model we computed the number of participants best fit and the exceedance probability - the probability that a model can best fit data from all participants under the assumption that all participants used the same model. Results are summarized in Figure S4. The best two models — that were almost tied for both number of participants best fit and exceedance probability — were the pure θ - σ model, with no additional factors, and the θ - σ model including the effect exploring or exploiting on the last trial. These findings accord well with recent work by (Song et al., 2019), who showed a similar dependence on previous exploration in a related task.

4.2 Correlation between directed and random exploration is conserved across models

Is the negative correlation between directed and random exploration robust to the inclusion of different factors in the model? To answer this question, we computed the correlation between θ and σ parameters for all 48 versions of the descriptive model. Regardless of the form of the model, we found that the strong negative correlation between directed and random exploration was maintained (Figure S5). This suggests that the trade off is not an artifact of the parameterization of the model.

Figure S4: Model comparison for different versions of the descriptive model. *Left* Parameters in each model ranked by number of participants best fit. Each row corresponds to a different model and each column to a different parameter. White cells correspond to the absence of the parameter in that model, grey cells to the presence of a parameter and (for β_e) black cells correspond to case where the sequential effect is considered up to two trials in the past. The purely random model (the worst fit of all at rank 49) is represented by the absence of all parameters, i.e. all entries are white. *Middle* The number of participants best fit by each model. *Right* The exceedance probability for each model.

Figure S5: Correlation between theta and sigma is robust to modeling assumptions.

4.3 The five parameter descriptive model

The remaining descriptive model analyses (including those in Figures 3 and 4 from the main paper) used a five parameter version of the descriptive model that included terms for directed and random exploration (θ and σ), lapsing (ϵ_1 and ϵ_2), and past effects of the explore-exploit choice (β_e). While this model was only ranked fifth in terms of BIC (Figure S4), it captures the main suboptimalities exhibited by people (i.e. lapsing and sequential effects of last choice), and includes most of the better models as special cases (e.g. $\epsilon_1 = \epsilon_2 = 0$). Thus it represents a flexible choice of model that accounts for most of the observed behavior without introducing too many free parameters.

4.4 Parameter recovery for the five-parameter descriptive model

While the analysis in Figure S5 shows that all versions of the descriptive model exhibit a negative correlation between directed and random exploration, it is still possible that this correlation could result from some bias in the fitting procedure or even a bug in the code. To determine whether the negative correlations were introduced artificially we ran two parameter recovery analyses for the five-parameter descriptive model (Wilson and Collins, 2019).

In these analyses we simulated data according to the model with known parameter values. We then refit this simulated data to ask whether the parameters fit to the simulated data were similar to the actual parameters used in the simulation. In the first analysis, we simulated data for each subject using the parameter values fit for that subject. As shown in Figure S6 this leads to a close correspondence between simulated and fit parameters; i.e. parameter recovery is excellent for this model in this task.

Figure S6: Parameter recovery for five-parameter descriptive model.

Next we asked whether the tradeoff between directed and random exploration could be induced artificially by the fitting procedure. To this end we simulated data from 174 participants using *scrambled* versions of the fit parameters that removed the negative correlation between θ and σ across the population (Figure S7 left). Because there was no correlation between θ and σ in the simulations a negative correlation in the recovered parameters would show that the tradeoff was an artifact. As shown in Figure S7 our fitting procedure did not artificially induce such a correlation suggesting that the tradeoff between directed and random exploration is a real feature of human behavior.

5 Fitting the Deep Exploration model to behavior

As with the Descriptive model above, to fit the Deep Exploration model to behavior, we augmented the basic choice probability (Equation 7 in the main paper) to include the possibility of lapsing and sequential effects. That is, we assumed the probability of exploiting was given by

$$\tilde{p}_{exploit} = (1 - \epsilon_1 - \epsilon_2) \frac{1}{1 + \exp(\text{logit}(p_{exploit}(n, T)) + \beta_e e_{past} + \beta_a a_{past} + \beta_{cp} c p_{past})} + \epsilon_1 \quad (\text{S30})$$

Figure S7: Parameter recovery for theta-sigma model shows that correlation between theta and sigma is not induced by model fitting procedure

where $p_{exploit}(n, T)$ was given by equation S21. Note that for the purposes of fitting, instead of parameterizing the model through the simulation horizon T , we used the implied hazard rate, $h = 1/T$. While this substitution does not change the properties of the model, it has the advantage (especially for visualization purposes) that h is between 0 and 1 while T is unconstrained.

5.1 Model comparison

As with the descriptive models, we fit multiple versions of the Deep Exploration that included/excluded different factors. As shown in Figure S8 the best fitting models included lapse rates and the effect of the past explore/exploit choice.

5.2 Parameter recovery

As with the descriptive model we performed a parameter recovery analysis to confirm that parameters could be recovered in the best case scenario in which the behavior is actually generated by the model. In particular, we focused on the version of the Deep Exploration model with yoked lapse rates (i.e. $\epsilon_1 = \epsilon_2 = \epsilon$) that included an exploration kernel (i.e. $\beta_e \neq 0$). Simulated data were generated using the fit parameter values. As shown in Figure S9 simulation and fit parameter value are highly correlated for all four parameters demonstrating good parameter recovery for this model.

5.3 Comparison between simulated and actual choice curves

To confirm that the Deep Exploration model provides a good fit for all participants, in Figure S10 we plot choice curves from all 174 participants along with the simulated choice curves from the Deep Exploration model. This shows excellent correspondence between model and experiment for all participants in our data set.

Figure S8: Model comparison for different versions of Deep Exploration model. *Left* Each row corresponds to a different model, each column to a different parameter. White indicates the absence of a parameter, grey the presence of a parameter, and black indicates cases in which the two lapse rates are constrained to be equal, i.e. $\epsilon_1 = \epsilon_2 = \epsilon$. *Middle* Number of participants best fit by each model. *Right* Exceedance probability for each model.

Figure S9: Parameter recovery for the Deep Exploration model with lapse rate and explore kernel parameters.

Figure S10: Simulated (blue) vs actual (red) choice curves for all participants in Experiment 1.

6 Deep exploration can account for horizon, uncertainty, and feedback dependence of directed and random exploration

Previous work has revealed three key features of directed and random exploration: (1) that both directed and random exploration depend on the time horizon (Wilson et al., 2014); (2) that directed and random exploration are both diminished when feedback is not contingent on choice (Findling et al., 2019); and (3) that directed and random exploration depend on the uncertainty in the options (Speekenbrink and Konstantinidis, 2015; Gershman, 2018). In this section we show that all of these results can be accounted for by Deep Exploration.

6.1 The Horizon Task and its variants

To ground our analysis, we focus on three variants of a gambling task known as the ‘Horizon Task’ (Wilson et al., 2014). In these tasks participants make a series of choices between two one-armed bandits that pay out probabilistic rewards. To maximize their rewards, participants need to exploit the option with the highest average payout, but because the reward probabilities are initially unknown, participants must explore both options in order to find out which option is best. Three variants of this task manipulate the horizon, feedback structure, and uncertainty, allowing us to investigate the three key features of directed and random exploration.

6.1.1 Horizon Task

In the Horizon Task (Wilson et al., 2014), participants play a series of ‘games,’ lasting either 5 or 10 trials each. On each trial they choose between two one-armed bandits that pay out rewards from Gaussian distributions. The means of the Gaussians are different for each bandit and change from game to game. In each game, one option is randomly assigned a mean of either 40 or 60 points, and the other a mean that is either higher or lower by 4, 8, 12, 16, or 20 points (thus one option is always better on average). The standard deviation of the Gaussians is identical for both bandits and fixed throughout the experiment at 8 points.

The first 4 trials of each game are ‘forced’ trials, in which participants are instructed as to which options they must play. By controlling which options participants play, these forced trials allow us to control the information people have about the options before their first ‘free’ choice. In particular, we use the forced trials to setup one of two information conditions: an unequal (or [1 3]) condition, in which participants play one option once and the other three times, and an equal (or [2 2] condition) in which participants play both options twice.

In addition to the information manipulation, we also manipulate the horizon by setting the number of trials in each game to be either 5 or 10. Because the first four trials are forced trials, this gives participants a horizon of either 1 or 6 free choices per game. The horizon dependence of directed and random exploration can then be measured by comparing behavior on the first free choice of each game, where the only difference between conditions is the number of decisions participants *will* make in the future.

To quantify directed and random exploration we fit a logistic model to the behavior. In particular, we assume that choice is a function of the difference in the observed mean rewards for each option, $\Delta R = R_{left} - R_{right}$, and the observed difference in information between the two options,

ΔI ($\Delta I = 0$ in the [2 2] condition, $\Delta I = +1$ when the left option was played once in the [1 3] condition, $\Delta I = -1$ when the right option was played once in the [1 3] condition). That is, we write the probability of choosing the left option as

$$p(\text{choose left}) = \frac{1}{1 + \exp\left(-\frac{\Delta R + A\Delta I + B}{\sqrt{2}\sigma}\right)} \quad (\text{S31})$$

where the free parameters A , B , and σ are the information bonus, side bias, and decision noise respectively. These free parameters are fit to each subject in each of the conditions: the two horizon conditions for A , and the four horizon $\times$ uncertainty conditions for B and σ . From there, directed exploration is quantified using the information bonus and random exploration by the decision noise.

Data from the original Horizon Task paper (Wilson et al., 2014) (30 participants after exclusions, 20 women, 10 men, ages 18-24) are shown in the top row of Figure S11. For simplicity we focus on decision noise in the [2 2] condition only, although, as shown in the original paper, similar results hold for decision noise in the [1 3] condition. The key finding is that both information bonus and decision noise increase with horizon (paired t-test comparing information bonus between horizon 1 and 6: $t(29) = 5.06$, $p < 0.001$; paired t-test comparing decision noise in [2 2] condition between horizon 1 and 6: $t(29) = 5.61$, $p < 0.001$). This indicates that both directed and random exploration are horizon dependent.

6.1.2 Feedback Task

In the original Horizon Task, participants only receive information about the options they choose. Such ‘partial feedback’ is critical for the explore-exploit dilemma to be present because it makes information contingent on choice. When information is not contingent on choice, for example in a ‘full feedback’ condition where participants gain information about both options regardless of choice, then exploration has no value because it yields the same information as exploiting. In the Feedback Task we explicitly manipulate the feedback participants receive between partial and full feedback.

The Feedback Task replicates the main features of the Horizon Task, with the same two horizon and two uncertainty conditions controlled by four forced trials (note that the forced trials always have partial feedback). On top of this, however, the Feedback Task also manipulates feedback. In the partial feedback condition, participants play 160 games of the original Horizon Task where information is dependent on choice. In the full feedback condition, participants play 160 games in which they see the outcome of both options regardless of choice on the free-choice trials. Critically, participants only receive reward for the option they choose. Thus, in the full feedback condition, participants are incentivized to exploit, but not to explore.

Behavior is fit using the same logistic model as used in the Horizon Task, and the key question is whether the information bonus and decision noise differ between feedback conditions. Data from 31 subjects (23 women, 8 men, ages 18-42) performing the Feedback Task is shown in the middle row of Figure S11. We see an interaction between horizon and feedback condition for both information bonus and decision noise (see Tables S1) that is driven by a reduced information bonus ($t(27) = -3.5$, $p = 0.0018$) and reduced decision noise ($t(27) = -3.6$, $p = 0.0014$) in horizon 6 in the full feedback condition. Thus, both directed and random exploration are reduced when feedback is not contingent on choice. Note that the result for noise replicates the earlier finding of

Figure S11: Behavior in the Horizon Task variants showing that directed and random exploration depend on (Top Row) horizon, (Middle Row) feedback and (Bottom Row) uncertainty.

Findling et al. (2019) who showed that exploratory noise is reduced in a full feedback task. Also, of note — and as we shall see, something that is predicted by Deep Exploration — is that decision noise is horizon dependent in the full feedback condition ($t(27) = 2.7, p = 0.011$).

Information bonus					
Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
horizon	697.96	1	697.96	18.67	0.0002
feedback	460.85	1	460.85	18.49	0.00021
order	187.10	1	187.10	3.73	0.064
subject(order)	1302.71	26	50.10	1.98	0.23
horizon*feedback	209.74	1	209.74	5.67	0.025
horizon*order	74.51	1	74.51	1.99	0.17
horizon*subject(order)	972.07	26	37.39	1.01	0.49
feedback*order	1.35	1	1.35	0.05	0.82
feedback*subject(order)	647.97	26	24.92	0.67	0.84
Error	999.21	27	37.01		
Total	5530.67	111			

Decision noise					
Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
horizon	514.77	1	514.77	9.51	0.0048
feedback	36.33	1	36.33	4.25	0.049
order	137.25	1	137.25	1.50	0.23
subject(order)	2372.57	26	91.25	1.56	0.12
horizon*feedback	20.82	1	20.82	4.96	0.034
horizon*order	69.99	1	69.99	1.29	0.27
horizon*subject(order)	1406.77	26	54.11	12.89	1.8e-09
feedback*order	21.99	1	21.99	2.57	0.12
feedback*subject(order)	222.35	26	8.55	2.04	0.036
Error	113.33	27	4.20		
Total	4942.66	111			

Table S1: ANOVA tables for the information bonus and decision noise in the Feedback Task

6.1.3 Parametric Uncertainty Task

(Gershman, 2018) investigated the uncertainty dependence of directed and random exploration by having a wider variety of forced trials. Here we focus on a conceptual replication of that study. In this Parametric Uncertainty Task, we use up to six forced trials to create six uncertainty conditions, three unequal uncertainty conditions ([1 2], [1 3], and [2 3]) and three equal uncertainty conditions ([1 1], [2 2], and [3 3]). In this task the horizon is fixed at 5 across all games.

29 participants (21 female, 8 male, age 18-42, mean 22.2) performed 240 games of the Parametric Uncertainty Task. Behavior was fit using the same logistic model, with the proviso that the information variable ΔI indicates which option is more informative to choose. That is $\Delta I = 0$ for

the equal uncertainty conditions $\Delta I = \pm 1$ for the unequal uncertainty conditions. This allows us to estimate an information bonus and decision noise for each uncertainty condition.

Analyses of variance revealed a trend towards a main effect of uncertainty condition on information bonus and a strong main effect of uncertainty condition on decision noise (see table S2). This trend for directed exploration was driven by an increased information bonus between the [2 3] and [1 2] conditions (paired t-test, $t(28) = 2.6, p = 0.015$) and the [2 3] and [1 3] conditions (paired t-test $t(28) = 2, p = 0.059$).

Information bonus					
Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
uncertainty	125.77	2	62.89	2.88	0.064
subject	2465.68	28	88.06	4.04	4.4e-06
Error	1221.50	56	21.81		
Total	3812.95	86			

Decision noise					
Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
uncertainty	923.54	5	184.71	21.71	4.7e-16
subject	2309.77	28	82.49	9.70	4.4e-21
Error	1190.96	140	8.51		
Total	4424.26	173			

Table S2: ANOVA table for the information bonus and decision noise in the Parametric Uncertainty Task

Further insight into these results can be obtained by plotting the information bonus and decision noise as a function of the uncertainty that a Bayesian ideal observer would have about the options after the forced choice trials. In particular, if we assume that the model knows the true variance of the reward distribution ($\sigma_n = 8$ points), then the uncertainty in the estimate of the mean of the rewards after the forced choice trials will be

$$\sigma^i = \frac{\sigma_n}{\sqrt{n^i}} \quad (\text{S32})$$

where n is the number of times option i has been played during the forced choice trials.

Plotting the information bonus as a function of difference in uncertainty between the two options and the decision noise as a function of total uncertainty (S11, bottom row) reveals a near linear relationship between the variables. This result is consistent with the findings of (Gershman, 2018) and represents a conceptual replication of his results.

6.2 Optimal model for Horizon Task variants

Subject to some assumptions, the optimal policy in the Horizon Task (and its variants) is computationally tractable. This analysis was first described in the Supplementary Material of (Wilson et al., 2014), but is summarized here as it helps frame the Deep Exploration model.

In particular, the optimal model assumes that participants know that the rewards on trial t , r_t , are generated from Gaussian distributions and, further, that participants know the standard

deviation of the Gaussian distributions, $\sigma_n = 8$. The model assumes that participants do not know the mean of the Gaussian distributions for option, i , μ^i , but that they can infer it from the previously observed rewards, the ‘data,’ D_t . If participants are good Bayesians then the inferred distribution over μ_i is

$$p(\mu_i|D_t) \propto \sqrt{\frac{n_t^i}{2\pi}} \frac{1}{\sigma_n} \exp\left(-\frac{n_t^i(\mu^i - R_t^i/n_t^i)^2}{2\sigma_n^2}\right) p(\mu_i) \quad (\text{S33})$$

where n_t^i is the number of times option i has been played, R_t^i is the cumulative sum of the rewards obtained from playing option i and $p(\mu_i)$ is the prior on the mean. For simplicity we assume a uniform prior for $p(\mu_i)$, although other forms, such as a Gaussian prior, are relatively straightforward to implement.

Equation S33 shows that the optimal model’s state of knowledge about option i can be summarized by two sufficient statistics: the number of plays, n_t^i , and the cumulative rewards, R_t^i . This allows us to define the state of information of the optimal model as the number of plays and cumulative rewards from both options; i.e.,

$$S_t = (n_t^1, R_t^1, n_t^2, R_t^2) \quad (\text{S34})$$

This in turn defines a state space from which the optimal policy can be computed by solving a dynamic programming problem (Bellman, 1956; Duff and Barto, 2002). In particular, the optimal model takes the action that maximizes the expected future reward, $Q(a|S_t)$; i.e.,

$$a_t^{optimal} = \arg \max_a Q(a|S_t) \quad (\text{S35})$$

Following (Bellman, 1956), $Q(a|S_t)$ satisfies

$$Q(a_t|S_t) = \sum_{S_{t+1}} T(S_{t+1}|S_t, a_t)(r_t(S_{t+1}|S_t) + V(S_{t+1})) \quad (\text{S36})$$

where, $T(S_{t+1}|S_t, a_t)$ is the transition matrix, denoting the probability of moving to state S_{t+1} given that action a_t is taken in the current state S_t , $r_t(S_{t+1}|S_t)$ denotes the reward obtained going into state S_{t+1} from S_t , and

$$V(S_{t+1}) = \max_a Q(a|S_{t+1}) \quad (\text{S37})$$

denotes the value of being in state S_{t+1} under the optimal policy.

The transition matrix, $T(S_{t+1}|S_t, a_t)$, is computed by noting that if option a_t is played and reward r_t is observed, then the sufficient statistics for option a_t update as

$$\begin{aligned} n_t^{a_t} &\leftarrow n_t^{a_t} + 1 \\ R_t^{a_t} &\leftarrow R_t^{a_t} + r_t \end{aligned} \quad (\text{S38})$$

updating the state as

$$S_{t+1} = \begin{cases} (n_t^1 + 1, R_t^1 + r_t, n_t^2, R_t^2) & a_t = 1 \\ (n_t^1, R_t^1, n_t^2 + 1, R_t^2 + r_t) & a_t = 2 \end{cases} \quad (\text{S39})$$

and that this state transition occurs with probability

$$\begin{aligned}
T(S_{t+1}|S_t, a) &= p(r_t|S_t, a) \\
&= \int d\mu^a p(r_t|\mu^a) p(\mu^a|S_t) \\
&= \sqrt{\frac{n_t^a}{2\pi(1+n_t^a)}} \frac{1}{\sigma_n} \exp\left(-\frac{(1+n_t^a)(r_t - R_t^a/n_t^a)^2}{2n_t^a\sigma_n^2}\right)
\end{aligned} \tag{S40}$$

Finally, to compute the values state and action values, $V(S_t)$ and $Q(a|S_t)$, we start at the horizon $t = H$ where the action values are simply the mean of the posterior

$$Q(a|S_H) = \frac{R_H^a}{n_H^{aH}} \tag{S41}$$

and then iterate equations S36 and S37 backwards in time to compute the action values for the initial state $Q(a|S_1)$.

6.3 Deep Exploration algorithm in the Horizon Task variants

The Deep Exploration model uses the same state space as the optimal model, but instead of considering all possible states, it only considers a small number visited during mental simulation.

In order to approximate the action values for the two options in the initial state, for each initial action ($a_t = 1$ or 2) the model simulates a sequence of states, actions, and rewards up to the time horizon H . For example, a simulation for choosing option 1 on the first trial would produce a trace of states, actions and rewards visited upto the time horizon; e.g.

$$\text{Simulation Output} = (S_1^i, a_1^i = 1, r_1^i, S_2^i, a_2^i, r_2^i, \dots, S_H^i, a_H^i, r_H^i) \tag{S42}$$

The value of each simulation is simply the sum of the simulated rewards; e.g.

$$\tilde{Q}^i(a_1 = 1|S_1) = \sum_{t=1}^H r_t^i \tag{S43}$$

The model then aggregates these values to compute the two action values

$$\tilde{Q}(a_1|S_1) = \frac{1}{n} \sum_{i=1}^n \tilde{Q}^i(a_1|S_1) \tag{S44}$$

where n is the number of simulations. The model then chooses the option with highest simulated value.

A critical component of the model is exactly how simulations play out — that is, how rewards and actions are sampled during the simulations. For the simulated rewards, for each simulation i , the model samples an underlying mean for each bandit, j , from the posterior distribution over mean given the outcomes of the forced trials; i.e.,

$$\tilde{\mu}_j^i \sim p(\mu_j|D_t) \tag{S45}$$

where the posterior over the mean, $p(\mu_j|D_t)$, is given by Equation S33. Given the sampled means, similar to the Deep Exploration algorithm for the Infinite Bandits Task, whenever a bandit is chosen in the simulation the reward is taken to be equal to that mean; i.e.

$$r_t = \tilde{\mu}_{a_t^i} \quad (\text{S46})$$

For the simulated actions, the model chooses according to a Thompson sampling scheme. In particular, we assume that the model chooses action a_t in state S_t with probability equal to the probability that option a_t has the higher mean in this state; i.e.,

$$p(a_t^i = 1|S_t^i) \sim p(\mu^1 > \mu^2|S_t) \quad (\text{S47})$$

After the simulated action is taken, the corresponding reward is sampled and the simulated state is updated. For the Horizon and Parametric Uncertainty Tasks, the state update follows equation S48. For the full feedback task, because information about both rewards is revealed, the state updates to

$$S_{t+1} = (n_t^1 + 1, R_t^1 + r_t^1, n_t^2 + 1, R_t^2 + r_t^2) \quad (\text{S48})$$

6.3.1 Deep exploration can account for behavior in Horizon Task variants

We used the Deep Exploration model to simulate behavior in each of the three tasks. To illustrate the model’s behavior, we set the number of samples $n = 2$ and the standard deviation of the reward distribution, $\sigma = 16$. Although we did not explicitly fit these parameter values to behavior (because the goal of our analysis is to show that Deep Exploration can capture the qualitative properties of human behavior) we note that these parameter settings give values for the information bonus and decision noise that are comparable to those seen in human performance.

As shown in Figure S12, the model captures the main qualitative features of all three experiments. In particular, the model captures the horizon dependence of directed and random exploration (Figure S12A and B); the reduction in directed and random exploration in the full feedback task (Figure S12A and B); the horizon dependence of decision noise in the full feedback task (Figure S12B); and the uncertainty dependence of directed and random exploration (Figure S12C and D).

Figure S12: Deep exploration model of Horizon Task can account for the dependence of directed and random exploration on horizon, feedback and uncertainty.

References

- Bellman, R. (1956). A problem in the sequential design of experiments. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 16(3/4):221–229.
- Duff, M. O. and Barto, A. (2002). *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. PhD thesis, University of Massachusetts at Amherst.
- Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S., and Wyart, V. (2019). Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature neuroscience*, pages 1–12.
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173:34–42.
- Song, M., Bnaya, Z., and Ma, W. J. (2019). Sources of suboptimality in a minimalistic explore–exploit task. *Nature human behaviour*, 3(4):361–368.
- Speekenbrink, M. and Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2):351–367.
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294.
- Wilson, R. C. and Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, 8:e49547.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., and Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074.