Skilled Bandits: Learning to Choose in a Reactive World

Jared M. Hotaling^{1,2}, Danielle J. Navarro¹, and Ben R. Newell¹

¹University of New South Wales

²University of Illinois at Urbana-Champaign

Please send all correspondence to:

Jared Hotaling

Department of Psychology

University of Illinois, Urbana-Champaign

603 E. Daniel St. Champaign, IL 61820

jaredhotaling@gmail.com

Abstract

In uncertain environments we must balance our need to gather information with our desire to reap rewards by exploiting current knowledge. Achieving this balance is further complicated in reactive environments where actions produce long-lasting change to the system. In four experiments, we investigate how people learn to make effective decisions from experience in a dynamic multi-armed bandit task. In contrast to the typical exploitation-dependent diminishing rewards found in previous studies, options were framed as skills that developed greater rewards the more they were chosen. In Experiment 1, we provide a proof of concept, and in Experiments 2-4 we explore the boundaries of participants' sensitivity to reactive dynamics. Our results suggest that most individuals can learn effective strategies for coping with these reactive environments. A two-part comparison of several competing psychological models supports several conclusions: 1) a sizeable minority of individuals learned that their environment was reactive, 2) evidence suggests several distinct groups of individuals employed unique decision strategies, and 3) testing models with the simulation method reveals qualitative misfits that motivate future theory development.

Keywords: decision making; dynamic environments; explore-exploit dilemma; decisions from experience

1. Introduction

Many of the important choices we face in life lack the detailed descriptions necessary for effective decision making. Instead, they require that decision makers learn about the alternatives on offer through experience. In these situations, we must balance our need to reduce uncertainty by gathering information through exploration, with our desire to reap rewards by exploiting our current knowledge (see Mehlhorn et al., 2015 for a recent review). However, in many situations helpful information is not easy to obtain.

Consider the dilemma of a student just commencing a university education who needs to choose which classes to take. The “mathematical principles of continental knitting” class sounds like fun but may not provide her with a clear career path. In contrast, “principles of commercial law” and “introduction to COBOL programming” both sound incredibly boring, but potentially lucrative over a longer time frame. Perhaps a good choice might be to invest in one’s long term future by choosing the law class or the programming class, but which one to take? Careers in law or IT could both be rewarding, but it might depend on one’s aptitude – the skills involved are rather different, and if she turns out to be better at programming than at law it would be sensible to focus on the computer science classes.

This scenario presents a difficult decision making problem: the student needs to estimate her long run rewards, based on how her skills are likely to develop over a long timeframe (in programming or in law), using only the information about how well she performs at the very beginning of the learning sequence. This is not always straightforward. Some skills can be acquired quickly and as a consequence can be rewarding after very little investment, whereas other skills mature slowly and the reward is greatly delayed. The eventual reward from the slow-maturing option may be greater – even allowing for time-discounting – but this may be difficult for the decision maker to discern.

This decision making problem provides an example of a *reactive* environment. Not only do the rewards associated with a particular option change over time (e.g., salaries for lawyers fluctuate), they

depend on the history of choices (e.g., how much effort has one spent learning the intricacies of corporate law). There are many other examples of this kind of reactivity: consumer demand for a product – and hence its profitability – is shaped by the way it is advertised, the output from a mine decreases as the resources are extracted, the output of a farm increases as cultivation techniques are optimized, and so on. The critical point is that in a reactive environment our actions may produce long-lasting changes in the world, or in the rewards that the world delivers. In such an environment it is important to learn what options exist and how one's actions influence these options over time.

Although many real-world environments are reactive, previous research on repeated decisions from experience has tended to focus on static (unchanging) environments. Previous work has shown that in such environments the optimal strategy is to “front-load” one's exploration because this maximizes expected long-term rewards (Tversky & Edwards, 1966). However, this is not the case when the reward associated with options can change over time. In these “dynamic” environments, it is important to periodically monitor the evolution of each option in case dramatic changes occur (e.g. Navarro, Newell, & Schulze, 2016; Neth, Engelman, & Mayrhofer, 2014). However even in those cases when the dynamics reflect a reactive world – rather than one in which the environmental changes are independent of the decision maker – research has tended to focus on environments characterized by exploitation-dependent *diminishing* rewards, and often framed in terms of resource depletion (Gureckis & Love, 2009). In these tasks, participants may be foraging for scarce resources, such as berries in the forest. When they find a patch with many berries, they stop to harvest, but at some point the patch will be depleted and they must find a new patch. Thus, the rewards associated with an option decrease each time it is chosen.

The student dilemma described above is a scenario where the opposite pattern holds: rewards increase each time an option is chosen. In this paper, we focus on sequential decision making in a reactive environment that is inspired by this scenario. This environment, framed in terms of choosing

which skills to develop, involves a different set of task features and psychological capacities. Decision makers must now strategize their choices so that they quickly find the options with the most potential for growth. With enough experience in the same environment, they can learn to identify promising options more quickly based on previously encountered growth trajectories, allowing more time for exploitation and growth.

1.1. The Skilled Bandit Task

Our experimental situations rely on a bandit task paradigm, in which participants are presented with choice alternatives with unknown reward distributions that must be learned from a sequence of consequential choices. The bandit task framework is relatively well studied in situations where the reward distributions do not change over time (Steyvers, Lee, & Wagenmakers, 2009) or change in a purely random way (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Speekenbrink & Konstantinidis, 2015; Yi, Steyvers, & Lee, 2009). In addition, there are studies that rely on bandit style problems in a “resource depletion” scenario, where the rewards decrease as a function of participant behavior (Bogacz, McClure, Li, Cohen, & Montague, 2007; Gureckis & Love, 2009)¹, or change as a function of the presence of contextual features (Schulz, Konstantinidis, & Speekenbrink, 2018), or the availability of options changes over time (Ejova, Navarro, & Perfors, 2009; Navarro, Tran, & Baz, 2018; Shin & Ariely, 2004). Comparatively little is known about how people solve a “resource cultivation” problem where the rewards associated with each option increase as a function of repeated usage.

To provide an intuitively plausible context for these problems, we framed our study as a simulation of a “skill development” scenario in which participants repeatedly chose which skills they wanted to develop. Each time a skill was chosen it produced an immediate reward, allowing a participant to judge the typical rewards associated with each option. However, with each action the

¹ Note that Gureckis and Love (2009) and Bogacz et al. (2007) did include one option whose rewards grew each time it was chosen.

chosen skill would incrementally develop so that it would offer a higher expected reward in the future. Options developed at different rates, with some “skills” initially offering low rewards but growing substantially with repeated investment. Thus, participants learned about the rewards associated with each option by sampling (i.e. choosing) each and observing the rewards they received. Since the number of choices in each game was fixed, efficient sampling was important. Participants completed five games in the same environment, with only response labels being changed. Experiment 1 serves as a proof of concept, in which we use an environment offering two alternatives that each grew at a constant rate to demonstrate that individuals can learn about reactive growth in simple environments. In Experiments 2-4 we extend these findings to investigate the learning in more complex environments. Also, to provide a modest degree of verisimilitude to the choice scenario, Experiments 2-4 use curvilinear acquisition curves to roughly mirror the empirically-observed effects of practice on performance (Anderson, 1982; Heathcote, Brown, & Mewhort, 2000; A. Newell & Rosenbloom, 1981).

Our goal in this article is to better understand how people learn to make good decisions about developing resources and skills in dynamic environments. We begin our investigation with four experiments using the skilled bandit task. Later, we use a two-pronged analysis to examine several competing models for how individuals completed the task. Some of these employ simple reinforcement-learning mechanisms, while others make use of environmental dynamics to predict future rewards, or use choice heuristics to simplify the learning process. Our aim here is to compare models of different types in order to gain new insights into the strategies people use. Given the complexity of the task, we are interested in the level of sophistication with which participants approach the problem as well as what they learned about environmental dynamics. We use models to compare participants’ behavior in environments with more simplified dynamics (Experiments 2 and 3) to those with more complex dynamics designed to interfere with people’s ability to track values and map the relationship between actions and rewards (Experiment 4).

To foreshadow our results, we find evidence suggesting substantial individual differences, which imply differences in both the beliefs people have about their environment and in the decision making strategies they use. We also find that our two model comparison methods lead to different insights, and we show that using models to simulate (rather than fit) behavior exposes systematic prediction failures, with no models capturing some theoretically important aspects of human performance.

2. Experiment 1

The choice environment in Experiment 1 involved two options defined by the rate at which each skill's rewards developed. The *slow* option (S) initially offered the higher reward, but the *fast* option (F) grew twice as fast and offered the best long-term rewards. We predicted that participants would initially prefer S, but would eventually learn to choose F. Because we were specifically interested in the degree to which individuals learned the reactive nature of the environment, we interrogated participants expectations for future rewards. After completing each game, participants were asked to predict the rewards that they expected to receive if they had made five additional choices. By comparing predictions across different sequences of imagined choices we test whether individuals understood that rewards did not merely grow over time, but rather as a function of choosing.² To preview our design, half of participants made predictions for the sequence SSSSF, while the other half made predictions for the sequence FFFFS. Comparing, for example, the fourth S in the former sequence to the final S in the latter, we can test whether predicted rewards grow as a function of the number of times *each* option is chosen or as a function of the number times *any* option is chosen.

2.1. Method

2.1.1. Participants. Two hundred four consenting US-based participants (74 female, 130 male, 0 other/undisclosed) were recruited via Amazon Mechanical Turk. The average reported age for

² It should be noted that Experiment 1 was chronologically the last experiment we ran. It is presented first because this simple task is the *logical* precursor to the more complicated designs used in Experiments 2-4, not because it was chronologically first.

participants was 38.28 years ($SD = 10.61$). The experiment lasted approximately 10 minutes, and all participants received \$1.70 (US), regardless of performance.

2.1.2. Procedure & Materials. Participants completed the experiment by visiting a dedicated website. They were shown an example choice screen from the game (see Figure 1), and were told the following:

This experiment takes the form of a simple game in which you are presented with two options that you can choose between. The options might be labelled LUY and REB, for instance, and you can click on the buttons to select one. When you make a choice you will receive some reward (points). Your goal when playing the game is to win as many points as possible, and you have a total of 30 clicks ("actions") in which to win them. During the game, your points total and actions remaining are displayed on screen.

After clicking on a button, participants received a reward. The amount was displayed inside the clicked button for 1000ms before disappearing. Participants were told:

There are some rules to the game and some ways of playing will work better (i.e. win more points) than others. In total, you'll be asked to play 5 different versions of the game, and each time your goal is to win as many points as possible across the 30 clicks.

The participant's points total and number of remaining actions were displayed on screen at all times. At the beginning of each game options were randomly assigned a three-letter label. At the end of the experiment, participants were given a completion code to claim payment via Amazon.

Figure 1. The choice screen in Experiment 1 after making three choices.

The expected reward of each option was determined by the number of times it had been chosen. Figure 2 shows the reward functions for S and F in Experiment 1. Unadjusted (see below) initial values of S and F were 10 and 1, respectively, and the slopes of their payoff curves were 2 and 4, respectively. To make the differences between options slightly less obvious, a small perturbation was applied to the reward, increasing the amount received by 0, 1 or 2 points, chosen uniformly at random. To also prevent simple strategies based on learning the initial value of each option, for each game we adjusted the initial value of both options by a sampling a value (uniformly with replacement) from a fixed set of possible offsets: 0, 11, 16, 29, 41 and 45. The equation and parameters for all reward functions are provided in the Appendix.

Figure 2. Reward functions for the two options in Experiment 1.

To explain the prediction component of the experiment participants received the following instructions:

At the end of every game you'll be asked to imagine that you have continued playing the game for five more clicks. We'll ask you to imagine (for example) that you chose LUY, then REB, then LUY again, and so on. For each of these hypothetical future choices we'll ask you to guess how many points that choice would earn.

After completing each game, participants were presented with a sequence of five hypothetical choices and were asked to type a number into a text box below each choice indicating the number of points they expected to receive. An image at the top of the predictions screen (see Figure A1) reminded participants of the available options and their previous locations on screen (left or right). Participants

were also reminded of the number of points received after last choosing each option³. After entering values for each imagined choice, the next game began with a new set of choice alternatives. Half of participants were placed in a condition where they predicted future rewards for the sequence SSSSF. The other half of participants made predictions for the sequence FFFFS. These sequences remained the same after each game.

2.2. Results

Figure 3 shows the mean choice proportions across trials and games in Experiment 1. As expected, participants initially preferred choosing S, but eventually learned to predominantly choose F. In Game 1, participants chose F on only 26.62% of trials ($SD = 6.89$). In Game 5, this increased to 57.66% of trials ($SD = 7.72$) indicating that participants had learned to select the option that offered the greatest long-term payoff.

Figure 3. Mean choice proportions in each game of Experiment 1.

³ If the participant had never chosen a option during the game, this value was listed as 'undefined'.

Our primary analyses focused on participants' predictions, which were converted into change scores by subtracting the final observed reward from the predicted reward. For example, consider a participant who finishes a game and receives 62 points after last choosing DAX and 23 points after last choosing WUG. When asked to predict the rewards they would receive if they continued to choose DAX, DAX, DAX, WUG, this participant responds with the values 67, 72, 76, 82, 26. In this case, change scores for these five predictions would be 5, 10, 14, 20, 3.

Figure 4 shows the mean predicted change in each between-subject condition. To check that participants understood the basics of the prediction task, we compared the first and fourth prediction in each condition. Since these involved choosing the same option, predictions for fourth reward should be higher than for the first reward. This was indeed the case, confirming that participant understood that reward grew over time ($M_1 = 3.91$, $SD_1 = 16.40$, $M_4 = 12.33$, $SD_4 = 17.50$, $t(203) = 22.52$, $p < .001$).

Figure 4. Mean predicted change in reward values for Experiment 1. Error bars indicate standard errors.

However, to further test if participants understood the reactive nature of growth, we compared predictions across conditions. If participants believed that rewards grew as a function of time – which we operationalize as number of choices for any option – their prediction for the final choice in the SSSSF conditions should be greater than that for the fourth choice in the FFFFS condition because the former occurs later in the sequence of imagined choices. In contrast, if participants believed that rewards grew as a function of the number of times a particular option was chosen, they should predict a greater reward for the fourth choice in FFFFS than for the fifth choice in SSSSF because the former occurs after three additional F choices. Figure 4 shows strong evidence that participants understood the reactive nature of the environment, with the fourth prediction in condition FFFFS being significantly greater than the fifth prediction in condition SSSSF ($M_4 = 12.89$, $SD_4 = 14.80$, $M_5 = 3.52$, $SD_5 = 20.53$, $t(202) = 3.75$, $p < .001$). The same is true when comparing the fourth prediction in condition SSSSF to the fifth prediction in condition FFFFS ($M_4 = 11.76$, $SD_4 = 19.94$, $M_5 = 1.94$, $SD_5 = 18.92$, $t(202) = 3.61$, $p < .001$). In further support of this conclusion, we found that predictions for the first instance of each option were similar across conditions. The final prediction in condition SSSSF did not significantly differ from the first prediction in condition FFFFS ($M_5 = 3.52$, $SD_5 = 20.53$, $M_1 = 4.10$, $SD_1 = 13.60$, $t(202) = .237$, $p = .813$). The same holds when comparing the first prediction in condition SSSSF to the final prediction in condition FFFFS ($M_1 = 3.73$, $SD_1 = 18.91$, $M_5 = 1.94$, $SD_5 = 18.92$, $t(202) = .674$, $p = .501$).

2.3. Discussion

Results from Experiment 1 allow several important insights. First, they show that people can learn to make good decisions in a simple dynamic growth environment. Although participants initially preferred to choose Option S – which provided the greatest immediate reward on Trial 1 – in Game 1, with experience they eventually preferred the superior Option F in Game 5. More importantly,

predictions for imagined future choices provide strong evidence that participants understood the reactive nature of growth. If participants believed that rewards grew independent of choice, rewards from a given option should depend on sequential position alone. However, we find clear evidence that predicted rewards depended on the number of times each option was chosen.

These findings serve as a proof of concept, indicating that participants can acquire knowledge of reactive environmental dynamics; knowledge crucial to achieving good outcomes in dynamic decision making. In Experiments 2-4 we extend these insights to progressively more difficult environments involving a greater number of alternatives with more complex reward functions.

3. Experiment 2

The choice environment in Experiment 2 introduced two additional complexities that should complicate learning. First, it presented participants with four options rather than two, therefore requiring more learning to determine the best alternative. Second, growth was characterized by curvilinear reward functions, which meant that participants could not linearly extrapolate growth rates across trials. These changes posed additional challenges to learning and promoted more extensive exploration.

Choice alternatives in Experiment 2 were defined by the amount and rate by which each skill's rewards developed: *low slow* (LS), *low fast* (LF), *high slow* (HS), and *high fast* (HF). We predicted that most participants would eventually learn to choose HF, as this offered the highest rewards. However, we also expected that some individuals would be attracted by the initial gains of LF and would not discover the best option. Across games, we expected these errors to diminish as individuals who had previously exploited HF would be less likely to settle for an inferior option because past experience suggested a better alternative. To further aid participants, half received additional instructions warning that some options may seem good at first but might fail to develop in the long term. We hypothesized

that participants receiving these additional instructions might make better choices and earn more points than those who did not.

3.1. Method

3.1.1. Participants. Two hundred one consenting US-based participants (85 female, 114 male, 2 other/undisclosed) were recruited via Amazon Mechanical Turk. The average reported age for participants was 35.35 years ($SD = 11.56$). The experiment lasted approximately 10 minutes, and all participants received \$1.70 (US), regardless of performance.

3.1.2. Procedure & Materials. Participants completed the experiment by visiting a dedicated website. They were shown an example choice screen from the game (see Figure 5), and were given instructions similar to those in Experiment 1. Options were labeled A to D. After clicking on a button, participants received a reward of between 1 and 99 points. The amount was displayed inside the clicked button for 800ms before disappearing. Participants were again told that the goal of the game was to “win as many points as possible”, but were now given a total of fifty actions with which to win them. As in Experiment 1, points total and number of remaining actions were displayed on screen at all times. To aid them in this more complex environment, participants were given additional instructions. Half of participants received the following *vague* instructions:

There are some rules to the game and some ways of playing will work better (i.e. win more points) than others.

The remaining half of participants in the *detailed* instruction condition were also told:

You should think about the game as like learning new skills. Each of the buttons is like a different skill that you can learn. Each time you practice a skill (i.e. click on a button) you get better at it

(i.e. win more points). But the rate at which you improve might differ across the different skills. Just like in life, some skills are quick to learn - you improve rapidly - but in the end they may not be as rewarding as other skills which are slower to learn but, ultimately, deliver larger rewards. To win the most points you need to find out which skill (button) will provide the best rewards overall.

Feedback was given at the end of each game, specifying the total points earned, relative to the maximum possible (approximately 2600 points) and minimum possible (approximately 1400 points) scores⁴. Participants would receive the maximum score if they always chose HF and the minimum score if they always chose LS, so these values merely specified the total range of possible scores rather than the range of plausible outcomes. At the beginning of each game options were randomly assigned to buttons A-D. At the end of the experiment, participants were given a completion code to claim payment via Amazon.

⁴ Note that 1400 is the approximate score that would result from fifty consecutive choice for LS, i.e. the option with the lowest total reward. Calculating the true minimum score requires the computation of an optimal decision policy, where points are treated as costs to be minimized.

Figure 5. The choice screen in the skilled bandit task after making six choices.

The expected reward of each option was determined by the number of times it had been chosen. The left panel of Figure 6 shows the reward functions for the four options in Experiment 2. To make the differences between options less obvious, normally distributed random noise ($M = 0$, $SD = 3$, rounded to the nearest integer) was added to the expected rewards. The equation and parameters for all reward functions are provided in the Appendix⁵.

⁵ Experiments 2-4 did not use the initial value offsets from Experiment 1.

Figure 6. Reward functions for the four options in Experiments 2 and 3.

3.2. Results

Figure 7 shows the number of points earned across games and conditions. Contrary to predictions, participants in the detailed instruction condition earned only slightly more points per game ($M = 2,078$, $SD = 152$) than those in the vague instruction condition, ($M = 2,075$, $SD = 159$), $t(199) = 0.11$, $p = 0.91$. We tested the impact of instruction condition on choice using a multinomial logistic regression model and found no significant effects (see Table 1). In light of this result, we collapse across instructions conditions for all subsequent analyses. The top left panel of Figure 8 shows the proportion of choices of each option for trials 1 to 50, averaged across Games 1 to 5. The data indicate that participants began by choosing from each option equally but soon developed a strong preference for HF. As expected, choice proportions for HF increased across trials, at the expense of the other options. It is noteworthy that choices for LF showed a smaller decrease, indicating that some participants were drawn to it early on, presumably because its rewards initially increase rapidly.

Figure 7. Mean points earned in each game and experiment. Error bars indicate standard errors.

Figure 8. Results from each experiment (columns), across trials (x-axis). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions. Note that for all measures, data are displayed averaged across the 5 games completed by participants. The Appendix contains figures displaying the same measures for each game separately.

Table 1

<i>Multinomial Logistic Regression</i>							
	<i>Odds</i>	<i>Effect</i>	<i>Estimate</i>	<i>SE</i>	<i>t</i>	<i>df</i>	<i>p</i>
<i>Experiment 2</i>	$\overline{P(HF)}$	Intercept	1.068	0.044	24.548	50,244	< .001
	$\overline{P(LS)}$	Condition	0.025	0.028	0.899	50,244	.369
	$\overline{P(HS)}$	Intercept	0.208	0.05	4.164	50,244	< .001
	$\overline{P(LS)}$	Condition	0.057	0.032	1.815	50,244	.07
	$\overline{P(LF)}$	Intercept	0.522	0.047	11.052	50,244	< .001
	$\overline{P(LS)}$	Condition	0.029	0.03	0.978	50,244	.328
<i>Experiment 3</i>	$\overline{P(T)}$	Intercept	1.896	0.057	33.272	49,494	< .001
	$\overline{P(B2)}$	Condition	-0.092	0.037	-2.484	49,494	.013
	$\overline{P(L)}$	Intercept	1.643	0.057	28.902	49,494	< .001
	$\overline{P(B2)}$	Condition	0.146	0.037	3.966	49,494	< .001
	$\overline{P(B1)}$	Intercept	-0.008	0.073	-0.109	49,494	.913
	$\overline{P(B2)}$	Condition	0.073	0.047	1.541	49,494	.123
<i>Experiment 4</i>	$\overline{P(T)}$	Intercept	1.485	0.04	37.27	74,994	< .001
	$\overline{P(B2)}$	Condition	-0.087	0.027	-3.253	74,994	.001
	$\overline{P(L)}$	Intercept	1.227	0.04	30.841	74,994	< .001
	$\overline{P(B2)}$	Condition	0.136	0.026	5.139	74,994	< .001
	$\overline{P(B1)}$	Intercept	0.061	0.048	1.253	74,994	.21
	$\overline{P(B2)}$	Condition	0.089	0.032	2.785	74,994	.005

To understand the strategies that people used to solve the skilled bandit task, we also analyzed people's exploration and exploitation behavior. The middle row of Figure 8 shows the probability of *switching*, defined as choosing a different option from the one chosen on the previous trial. Clearly participants began each game with a large degree of exploration, switching between options often. However, this exploration quickly gave way to more exploitative behavior, characterized by lower switching rates. To test the effects of trial on switching we used a multilevel generalized linear mixed-effects model with a Bernoulli distribution and logit link function. In the model *trial* was a fixed effect, with random intercepts at the subject level. We found a significant effect of trial ($\beta = -0.145$, $SE = .008$, $p < .001$), indicating reduced exploration across trials. Finally, Figure 8 also shows the distribution of *stopping points* across participants, defined as the trial on which all exploration ceased and the participant chose the same option on every subsequent trial. As the plot illustrates, in the large majority

of cases people ceased exploration by trial 30 at the latest – we return to this point later in the paper when discussing the behavior of various models.

3.3. Discussion

The results from Experiment 2 build on those from Experiment 1, and demonstrate that people are capable of learning to make good decisions in a more complex dynamic multi-alternative risky choice environment characterized by resource growth. With experience, the majority of participants discovered and exploited the option offering the highest rewards. Across trials participants improved their performance, earning more points with less exploration as they accumulated information about the dynamics of the choice environment.

The switching and stopping point distributions shown in Figure 8 indicate that participants front-loaded their exploration, focused their efforts on distinguishing the options from one another at the beginning of the game. Later on, participants reduced exploration and committed to exploiting their preferred option. Such front-loaded exploration would also explain why some participants continued choosing LF. Like HF, LF's rewards increased rapidly. This would have immediately attracted participants, leading them to invest early actions in this option, further increasing its rewards. It is unclear whether those who continued choosing LF mistakenly thought it was the best option or whether they eventually realized their mistake but decided it was too late in the game to shift their focus to developing another option (see Rich & Gureckis, 2018 for a similar 'learning trap'). Figure A2 in the Appendix provides support for the idea that the increased performance observed across games partially stems from participants learning to distinguish LF from HF.

These findings, together with those from Experiment 1, paint a picture of sophisticated, if imperfect, learning. The absence of a clear effect of instructions suggests that, as was found in previous studies (Navarro et al., 2016; B. R. Newell, Koehler, James, Rakow, & Van Ravenzwaaij, 2013; B. R. Newell & Rakow, 2007), participants learned primarily through their experience making choices in the

experiment. If participants' behavior was initially influenced by the instructions they received, it was quickly overshadowed by their direct experience playing the game, which was the same across conditions.

4. Experiments 3 and 4

In Experiment 2, all options had relatively similar growth patterns and offered rewards within a somewhat narrow range for a given number of actions. In Experiments 3 and 4 we increased these differences to make the options more distinct. A *target* option, T, offered the highest total rewards but was virtually indistinguishable from a *lure* option, L, early on. This option also offered fast increasing rewards for the first ten actions, after which its rewards plateaued. To compensate for the increased task difficulty produced by the lure, the remaining two bad options, B1 and B2, were made clearly inferior. Thus, the reward functions used in Experiments 3 and 4 improve on those used in Experiment 2 in two ways: 1) no longer does one option dominate throughout the game, always offering the highest rewards, and 2) two inferior options are easily identifiable from the beginning of each game, therefore allowing participants to focus their attention on the two good options.

The environmental dynamics in Experiment 3 matched those from Experiment 2; each time an option was chosen its expected reward incremented according to its reward function, otherwise options remained unchanged. In Experiment 4 we explore a different type of environment, where unchosen options decrement. Inspired by the “use it or lose it” nature of many real world skills, we incorporated a simple forgetting function (e.g. Averell & Heathcote, 2011; Navarro, Pitt, & Myung, 2004; Rubin, Hinton, & Wenzel, 1999) in which each time a participant failed to practice a skill (i.e. choose an option) the expected rewards for the skill decreased, moving one step backwards on the reward function in Figure 6 (right). Thus, focusing one's efforts on one skill meant becoming “rusty” at other neglected skills. We expected this additional complication to increase the difficulty of the task, especially at the beginning when participants may have a prior expectation that unchosen options are static. Additionally, these

dynamics complicate the mapping between actions and outcomes, and increase the difficulty of tracking option values in memory. This should particularly affect the efficacy of simpler learning strategies that assume independence between options, and might lead participants to adopt a more sophisticated approach that uses memory for past actions to inform valuation and choice. Doing so would allow participants in Experiment 4 to adapt their behavior by switching between options less and stopping exploration earlier, so as to minimize decrements.

4.1. Method

4.1.1. Participants. One hundred ninety-eight people (69 female, 128 male, 1 other/undisclosed) participated in Experiment 3, and three hundred people (121 female, 178 male, 1 other/undisclosed) participated in Experiment 4. Each was based in the US and was recruited via Amazon Mechanical Turk. The average reported age for participants was 34.98 years ($SD = 11.15$) and 36.40 ($SD = 11.88$), respectively. As in Experiment 2, Experiments 3 and 4 each lasted approximately 10 minutes and participants received \$1.70 (US) each.

4.1.2. Procedure & Materials. The experimental task, stimuli, and instructions were identical to those of Experiment 2. Although we did not find an effect of instructions in Experiment 2, the design changes made for Experiments 3 and 4 – e.g. the introduction of a lure option – led us to believe it may prove important. Again, approximately half of participants were randomly selected to receive either vague or detailed instructions⁶. The right panel of Figure 6 shows the expected reward functions for the four options in Experiment 3. Here, the expected reward of each option was determined by the number of times it had been chosen, as with Experiment 2. However, in Experiment 4, each time an option was chosen, its reward incremented, but the rewards for all other options decremented. That is, after choosing an option, its expected reward would increase (move to the right) according to the curves in

⁶ Due to a randomization error, Experiment 3 involved 93 participants in the vague condition and 105 participants in the detailed condition. For Experiment 4, 128 and 172 participants were in the vague and detailed conditions, respectively.

Figure 6. The expected rewards for all unchosen options would decrease (move to the left) according to the curves in Figure 6. The only restriction to these decrements were that options never decreased below their initial value. Normally-distributed random noise ($M = 0$, $SD = 3$, rounded to the nearest integer) was again added to the expected rewards to make the differences between options less obvious.

4.2. Results

Figure 7 shows that instructions did not significantly affect the number of points earned in Experiment 3 ($t(196) = 1.58$, $p = 0.12$) or Experiment 4 ($t(298) = 0.93$, $p = 0.34$). However, participants did earn significantly more in Experiment 3 ($M = 2,371.56$, $SD = 231.25$) than in Experiment 4 ($M = 2,174.29$, $SD = 362.35$), $t(496) = 6.80$, $p < 0.001$. In Experiment 3, our multinomial logistic regression analysis showed that instruction condition had a significant effect on the relative odds of choosing T and L relative to the reference category B2 (see Table 1). In Experiment 4, instructions had a significant effect on the relative odds of choosing each of T, L, and B1 relative to B2 (see Table 1). Despite finding evidence that detailed instructions in Experiments 3 and 4 did affect choices, in both cases the qualitative effect was modest, and for simplicity we again collapse across instruction sets in later analyses (see Figures A3 and A4 for details).

Figure 8 shows the major experimental results after collapsing across instructions conditions. Although behavior was quite similar across experiments, some differences are evident. In both experiments, participants developed a strong preference for the Target (T) and Lure (L) options, though participants chose these more in Experiment 3 ($M = 0.85$, $SD = 0.08$) than Experiment 4 ($M = 0.78$, $SD = 0.14$), $t(496) = 6.27$, $p < .001$. Note that this lower rate of choosing T and L in Experiment 4 contributes to the lower points totals observed in this experiment.

People's exploration was also similar across experiments. Figure 8 shows that the probability of switching decreased significantly across trials in Experiment 3 ($\beta = -0.134$, $SE = .006$, $p < .001$) and Experiment 4 ($\beta = -0.101$, $SE = .004$, $p < .001$) as exploration gave way to exploitation. However, it is also notable that switching rates were lower in Experiment 4 as participants attempted to avoid costly reward decrements. Replicating the results for Experiment 2, the stopping time for exploration tended to be fairly early, with the vast majority of participants choosing the same option for every trial after Trial 30, as shown in Figure 6. See Figures A5 and A6 in the appendix for a breakdown of behavioral results by game.

4.3. Discussion

Our goals with Experiments 3 and 4 were twofold. First, we simplified information search and learning by including two obviously inferior options. We hoped that this would free participants to focus their efforts on identifying and distinguishing the target option from the lure. In one sense, this was true. Participants developed a strong preference for Options T and L over B1 and B2. However, participants often preferred L most of all, rather than the optimal option, T. Potential explanations for this fall into two categories. First, it could be that participants who initially chose L never explored other options and therefore failed to learn that T offered better long-term rewards (see Rich & Gureckis, 2018). The greater frequency of T choices in the detailed instructions conditions supports this. Evidently this warning helped some participants to realize that they had been lured by a suboptimal option, leading them to eventually switch to option T. Another possible explanation for choosing L over T is that after choosing L several times participants realized they had been lured, but also understood that they had only a limited number of actions remaining, so they would be better off continuing with L than switching to an "under-developed" T. Such thinking may have led some to employ a front-loaded exploration strategy that intentionally neglected late exploration in favor of commitment to the most developed option.

Our second goal with Experiments 3 and 4 was to demonstrate that participants can adapt to environments characterized by the “use it or lose it” nature of many real skills. Not surprisingly, performance in Experiment 4 was adversely affected as participants initially struggled to understand the relationship between their choices and rewards. In contrast, participants in Experiment 3 had an easier time distinguishing the options early on. This would explain why more people were drawn to the lure — which offered the highest rewards at the beginning of each game — in Experiment 3 compared to Experiment 4, where the early advantage of L over T would have been harder to notice. However, despite being in a less advantageous environment, by Game 5 participants in Experiment 4 had improved substantially and earned points at rates similar to participants in Experiment 3 (see Figures A3 and A4).

The observed instruction effects show that those in the detailed condition were better able to avoid the lure. This is not surprising considering that the detailed instructions warned participants that some options may have reward functions like that for L. That instructions would affect choices in Experiments 3 and 4 but not Experiment 2 is also sensible because identifying and avoiding the lure in Experiments 3 and 4 was more difficult and more important. That is, in Experiment 2, HF was the best option for the entirety of the game, whereas in Experiments 3 and 4, T initially offered lower rewards than L. For this reason, a special warning that initially attractive options may change for the worse was especially well suited to the reward functions used in Experiments 3 and 4.

5. A Comparison of Psychological Models for Skilled Bandit Problems

Experiments 1-4 illustrate that people perform well at the skilled bandit task. More precisely, we observed that the majority of people eventually exploited the options that provided the highest (or second highest) rewards over the long term. More importantly, in Experiment 1 we obtained evidence for sensitivity to the reactive dynamics of the environment via the post-test prediction questions. To a similar end, we now use Experiments 2-4 to consider several computational models with different

assumptions about environmental dynamics. Our approach is a simple one; by comparing each model's ability to account for participants' behavior we gain new insights into the strategies people used, as well as beliefs they held about environmental dynamics. As a general rule, we assume that participants will seek to optimize their performance and use strategies that are well-suited to their environment. As a result, we expect that participants' behavior will be better captured by models whose assumptions and learning mechanisms are well suited to the environment in each experiment.

The model comparison involves two components. In the first, we conduct a quantitative analysis in which we fit several competing models to the choices observed in Experiments 2-4 at a trial-by-trial individual subject level. Here we find mixed results, with most individuals best accounted for by simple reinforcement learning models, and others by more sophisticated models that use environmental dynamics to predict future rewards. In the second part of the analysis, we take the best performing models and simulate their behavior in the task and find both similarities between models and systematic prediction errors, with all models failing to capture some theoretically important characteristics of human performance.

5.1. Skilled Bandit Models

We begin by outlining several models that might plausibly account for human behavior in the skilled bandit task. Several of these models can be construed as reinforcement-learning (RL) models with varying degrees of sophistication (Sutton & Barto, 1998) while others rely on decision making heuristics that are more typical in the judgment and decision making literature (Gigerenzer & Todd, 1999; Payne, Bettman, & Johnson, 1993). Each model operates through a small number of simple, iterative processes that take advantage of environmental regularities to produce seemingly complex behavior without extensive computation or knowledge representation.

5.1.1. The Naïve Mean Model. One might imagine that a simple strategy for the skilled bandit task would be to track the mean reward for each option and choose the one with the highest value.

With the *Naïve Mean* model we test the hypothesis that participants used this strategy in the experiments. To begin, each option's value is initialized to 19⁷. On each trial the model chooses probabilistically based on the mean reward for each option. After each choice, the mean for the chosen option is updated in light of the obtained reward, and a new choice is made according to a softmax rule:

$$p(j) = \frac{e^{v(j)/\tau}}{\sum_i e^{v(i)/\tau}}, \quad (1)$$

where j is an option, $v(i)$ is the mean reward obtained so far from the i -th option, and τ is a parameter controlling sensitivity to evidence strength. This model entirely ignores the dynamics of the environment because it reduces each option to a mean value, rather than using the history of rewards to predict the future.

5.1.2. The Recency Model. Given that expected rewards associated with each option change over time in the skilled bandit environment, a more appropriate strategy would be to choose based on recent rewards. According to the *Recency* model, on each trial the decision maker chooses probabilistically based on the most recent reward received from each option. Thus, the model is identical to the Naïve model, except that $v(j)$ is the value of the *last* reward from option j .

5.1.3. The Delta Rule Learning Model. The *Delta Rule Learning* model (DL; Samuel, 1959; Sutton, 1984) represents a reinforcement learning approach to solving the skilled bandit problem, and is of particular interest in neuroscience models of learning (Schultz, Dayan, & Montague, 1997). Its goal is to use error feedback to sequentially update expectations for each alternative. Before taking an action, the model computes an estimate, Q , of the value of each option (initialized to 19). After a choice is made, the obtained reward is used to compute an error,

$$\delta_t = r_t - Q(j)_{t-1}, \quad (2)$$

⁷ This value was selected to be slightly lower than the initial value of each option (20).

where r_t is the reward obtained on trial t after choosing Option j . This difference is then used to update the prediction for Option j according to:

$$Q(j)_t = Q(j)_{t-1} + \alpha \cdot \delta_t, \quad (3)$$

where α is a learning rate parameter. For small values of α , predictions will be based on a large number of past rewards and will be updated slowly. For large values of α , learning will occur quickly and predictions will depend primarily on the most recent rewards received. Finally, to derive choice probabilities Q values are used in place of v in Equation 1.⁸

5.1.4. The Eligibility Trace Model. The *Eligibility Trace* model (ET) takes learning a step further by using memory for past actions to inform value updating. Such models are designed to cope with delayed feedback, where the effects of one's actions are not immediately experienced. They do this by keeping a memory of past actions and assigning credit or blame for changes in rewards based on which options have been chosen recently. We use a version of the model presented by Bogacz et al. (2007) modified for four alternatives. For each alternative the model stores a memory trace, ε , initialized to 0. After making a choice, the memory trace for the chosen option, j , and each option, i , (including $i = j$) are updated according to:

$$\begin{aligned} \varepsilon_{t,i} &= \varepsilon_{t-1,i} \cdot \gamma, \\ \varepsilon_{t,j} &= \varepsilon_{t-1,j} + 1. \end{aligned} \quad (4)$$

γ is a memory decay parameter, between 0 and 1, controlling the rate at which each trace decreases.

Equation 2 is used to compute a prediction error, however unlike with DL, each option is updated according to:

$$Q(i)_t = Q(i)_{t-1} + \alpha \cdot \delta \cdot \varepsilon_{t,i}. \quad (5)$$

⁸ Note that the Recency model is a special case of the DL where $\alpha = 1$.

As with DL, Equation 1 is used to compute choice probabilities using Q values in place of v . Thus, the model describes a process of prediction-based learning similar to DL. It adds to this memory traces – which track the exponentially decreasing frequency of choosing each option – that modify the degree to which each prediction is updated.

Since the ET model relaxes the assumption (made by DL) that the results of one's actions are immediately experienced it should be less adaptive in the environments of Experiments 2 and 3, where the rewards for each option grow independently. However, ET is robust against delayed feedback, so we expect it to perform well in Experiment 4, where rewards do not simply grow independently with choice, but might instead depend on one's entire sequence of choices.

5.1.5. The Linear Model. The *Linear* model departs from those above, and draws inspiration from the judgment and decision making literature. It uses past rewards as cues for predicting future earnings. Using linear regression, it fits a line to each option's previous rewards, estimating a temporary intercept, b , and slope, m . These values are then used to compute the value of the next reward expected from each option. In this sense, the model can be thought of as an optimal – i.e. expected utility maximizing (assuming linear reward functions; see Payne et al., 1993) – weighted additive model, with previous rewards acting as cues for predicting the next outcome. The model begins by assigning an initial value of 19 to each option. After each action, a new temporary intercept (b) and slope (m) are calculated in light of the new reward. These are used to update the model's estimated intercept, B , (initialized to 19) and slope, M , (initialized to 0) values according to the delta rule. That is, both intercept and slope are learned via the delta rule with a single learning rate parameter. Choice probabilities are again described by Equation 1, but with $v(j)$ being the predicted future reward for option j according to the estimated regression line.

We tested two versions of the linear model. The first, *Linear*, assumes that rewards grow independent of an individual's choices. It therefore estimates the linear relationship between trial

number and the rewards received for choosing each option. The second, *Linear-R*, assumes a reactive environment in which one's actions determine growth, and estimates the rate at which rewards increase each time an option is chosen (i.e. it estimates the linear relationship between rewards and the number of times an option was chosen). Comparing these models – with respect to how well they account for behavior – provides evidence for whether or not a participant learned the reactive nature of the task environment.

5.1.6. The Threshold Model. The *Threshold* model implements a decision-oriented heuristic strategy consisting of a simple choice rule. According to the model, the decision maker explores the options randomly (selecting with equal probability) until one option produces a reward greater than a predetermined threshold, ϑ . At this point in time, all other options are ignored and the agent chooses the selected option for the remainder of the game.

5.1.7. The Cull Model. The *Cull* model is a variation of the Threshold model. Rather than searching for rewards above a fixed threshold, it operates through sequential elimination (e.g. Tversky, 1972), whereby the decision maker mentally removes clearly inferior options one at a time. It begins by choosing from each option with equal probability. As it received rewards, it compares the largest reward observed in the game so far to the largest reward produced by each option (initialized to 19). If the difference is greater than a predetermined threshold, θ , the option is removed from consideration and never selected again.

5.1.8. The Two-Component Model. Taking a slightly more sophisticated approach, the *Two-Component* model combines features from DL and Linear. It proposes that people attend to two sources of information about each option: the most recent reward received from that option, and the rate of

growth in the options rewards.⁹ The linear slope, M_i , for each option's reward growth (initialized to 0) is computed and combined with its Q value (initialized to 19) according to:

$$v(i) = w \cdot M_i \cdot s + (1 - w) \cdot Q(i). \quad (6)$$

In this expression, w is a free parameter between 0 and 1 that controls the relative weight placed on recent rewards versus the change in rewards, and s is an arbitrary scaling factor (fixed at $s = 10$ in our modeling) chosen to roughly equate the reward and slope scales.¹⁰ Thus, the model instantiates the intuition that participants in the skilled bandit task track which options are *currently* offering the best rewards *and* which options are *growing* the fastest. After each choice Equations 2 and 3 are used to independently update M and Q estimates via the delta rule, with both processes sharing a single α parameter. M is updated as in the Linear model, except that the temporary slope, m , is computed across the five most recent rewards. After calculating the value of each option, Equation 1 is used to determine the likelihood of choosing each option.

As with Linear and Linear-R, we tested two versions of the model that differed in how they calculated the slopes of reward functions. The *Two-Component* (2C) model assumes that rewards grow as a function of trial number – independent of an individual's choices – while the *Two-Component-Reactive* (2CR) model assumes that growth is a function of choosing and therefore depends on the number of times an option is chosen. We compare the relative performance of these models as means of assessing whether or not participants learned to exploit the reactive nature of the task environment.

5.1.9. Summary. In total, we examine a diverse set of ten models. With the Naïve and Recency models, we test the hypothesis that participants made choices by tracking simple summary measures. With the Threshold and Cull models, we test for the use of heuristics whereby participants compare

⁹ Alternatively, 2C can be viewed as a method of approximate planning, with slope and intercept used to compute an expected future reward. However, unlike with the Linear model, this framing is only approximately accurate due to the complicating effect of the w parameter.

¹⁰ Note that s can be absorbed into w and τ parameters if $w' = s \cdot w / (1 + (s - 1) \cdot w)$ and $\tau' = \tau \cdot (1 + (s - 1) \cdot w)$.

rewards (or reward differences) to a criterion value. The DL and ET models posit that learning occurs through sequential updating of values based on new observations. Finally, the Linear and 2C models test the hypothesis that participants use past reward growth to predict or anticipate future outcomes. Additionally, with the reactive variants of these models (Linear-R and 2CR) we test if participants understood that rewards grew in response to their actions.

5.2. Modeling Individual Subject Data

To fit the models to data at an individual subject level, we take an approach similar to Speekenbrink and Konstantinidis (2015), and use the “one step ahead” procedure from Yechiam and Busemeyer (2006). Under this approach, the model behavior is yoked to that of the participant. That is, the model is given the full history of an individual’s previous choices and rewards until trial $t - 1$, and is asked to predict the probability of making each response on trial t . To measure the overall fit of the model’s predictions to an individual’s behavior, we sum the log of the probability of making the same choice as the individual on all 250 trials in the experiment. Maximum likelihood parameter estimation was performed using a grid-search plus gradient descent routine¹¹.

Table 2

Model Fitting Results

Model	Parameters	-Log-Likelihood			AIC Weight			Individuals Best Fit		
		Exp. 2	Exp. 3	Exp. 4	Exp. 2	Exp. 3	Exp. 4	Exp. 2	Exp. 3	Exp. 4
Naïve	τ	174.78	182.08	222.2	0.01	0.01	0.01	2 (1%)	3 (1.52%)	3 (1%)
Recency	τ	179.74	187.56	217.04	0.01	0.02	0.01	2 (1%)	5 (2.53%)	5 (1.67%)
DL	τ, α	134.19	136.55	139.39	0.26	0.23	0.20	95 (47.26%)	82 (41.41%)	96 (32%)
ET	τ, α, γ	132.93	134.80	136.01	0.25	0.24	0.35	37 (18.41%)	38 (19.19%)	107 (35.67%)
Linear-R	τ, α	183.05	189.46	224.46	0.01	0	0	2 (1%)	0 (0%)	1 (0.33%)
Linear	τ, α	190.93	204.30	243.92	0	0	0	1 (0.50%)	0 (0%)	1 (0.33%)
Threshold	τ, ϑ	174.10	179.70	198.18	0.03	0.02	0.02	6 (2.99%)	5 (2.53%)	6 (2%)
Cull	τ, θ	159.88	158.51	171.25	0.08	0.11	0.10	16 (7.96%)	19 (9.60%)	29 (9.67%)
2CR	τ, α, w	132.88	134.01	137.70	0.20	0.23	0.17	23 (11.44%)	37 (18.69%)	32 (10.67%)
2C	τ, α, w	133.28	134.73	138.04	0.16	0.13	0.13	17 (8.46%)	9 (4.55%)	20 (6.67%)

¹¹ More precisely: to avoid the problem of local minima, a grid of equally spaced parameter start points was used. A bounded Nelder-Mead optimization routine (Nelder & Mead, 1965) was run from each of these points, and the best discovered parameters were used.

As illustrated in Table 2, the models differed in terms of the number of free parameters. To correct for this difference in complexity we computed the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for each combination of individual and model. Since our conclusions are the same for both measures, we focus on AIC for the remainder of this article, and to assist with interpretation we compute AIC weights that roughly estimate the probability of each model according to the observed data (Wagenmakers & Farrell, 2004)¹². Mean model fit measures are reported in Table 2, and Figure 9 plots the distribution of individual AIC weights for the top five models (calculated using the entire array of models). Each point represents an individual, while the solid and dotted lines indicate the mean and median values, respectively.

Figure 9. Akaike Information Criterion weights for each model in Experiments 2-4. Each dot represents an individual. Mean and median values are indicated by the solid and dotted lines, respectively. Shaded areas indicate 95% confidence intervals. Higher weights indicate a greater probability of each model according to the observed data.

¹² Some care is required when interpreting AIC weights. The AIC weights assume that the true model is included in the test set, and more generally the AIC is based on a large sample approximation (Akaike, 1998). A variety of more sophisticated methods are available (see, e.g. Myung, 2000 for a review) but are often computationally prohibitive to implement.

In Table 2, we can see that several models provide poor fits to data. Models that track simple summary statistics (Naïve and Recency) perform very poorly in all three experiments, as do most of the heuristic models (Linear, Linear-R, and Threshold). For the best models, Figure 9 reveals clear differences. The performance of Cull is interesting because it yields particularly extreme AIC weights that indicate very good fits to some individuals and very bad fits to most others. This result points to the presence of a distinct group of participants (10%) who were well characterized by the model, which is loosely inspired by Tversky (1972) theory of covert sequential elimination. DL performed the best overall, followed by ET, suggesting that the majority of participants focused on learning the current value of each option, rather than extrapolating from past observations to predict future outcomes.¹³ The success of 2CR and 2C in fitting some participants (10-20% and 5-8%, respectively) indicates that a number of individuals combined the DL's learning mechanism with a second focused on learning the rate of change in outcomes.

Several conclusions can be drawn from this analysis. Many of the simplest models perform poorly, suggesting that people do not merely track the mean or most recent rewards. The failure of the Threshold model likewise indicates that people do not simply stop exploring when they reach some goal. On the other hand, the similar Cull model gives a good account of some people's choices, so it would seem that sequentially removing poor options based on relative performance is a more plausible strategy than the fixed-criterion approach adopted by the Threshold model. 2CR and 2C were the best models for a sizeable minority of individuals, implying that many people attempted the more complex strategy of tracking both the current and expected future reward for each option. However, the success of DL and ET shows that most people neglected any extrapolation from past rewards.

¹³ Note that the asymptote in Figure 9 around .49 for DL stems from the fact that it is nested within each of ET, 2C, and 2CR. For many participants who are well fit by DL, there is a version of these more complex models which produces a similarly good fit, but with one additional parameter. These individuals are also represented in the cluster of points around .17 for each of ET, 2C, and 2CR.

Our model comparison also provides insight into participants' beliefs about the dynamics of the skilled bandit environment. A sizeable proportion of individuals were best fit by 2CR, which assumes a reactive environment where outcomes depended on one's choices. These individuals appeared to understand that by choosing an option they increased its future expected payoffs. A smaller group of participants were best fit by 2C, which assumes that the environment changes at a constant rate, independent of the choices one makes. The remaining participants were best characterized by models that implied either constant expected rewards or random drift in rewards. These findings paint a picture of several distinct groups, with different beliefs about the environments, applying diverse strategies to the skilled bandit task.

When we look at the modeling results in more detail, interesting patterns emerge across experiments. In Experiment 2, DL was the best fitting model (in terms of individuals best fit), followed by ET, 2CR, 2C, and Cull. This pattern was similar in Experiment 3, though 2CR's performance increased (largely at the expense of DL). Experiment 4 shows a different result, where ET is the best model, with DL and 2CR faring worse than in Experiment 3. This change is sensible because ET extends DL by allowing learning signals to cascade backward to earlier actions, assigning partial credit or blame to older actions. This is ill suited to the environment in Experiments 2 and 3 because each option grew independently from the others and so previous choices for one option did not affect the reward obtained after choosing another option. This contrasted with the environment in Experiment 4, where each choice affected the expected payoffs for all options. Here eligibility traces could provide a powerful method for improving performance. For instance, an individual might choose Option X, observe a Reward R, then explore the other three options before returning to X and receiving a reward less than R. DL would simply "blame" X for the reduced reward by decreasing its estimated Q value, but ET provides a method for also decreasing the value of the other options. By "blaming" these options for the decrement in X's payoff function, ET decreases the chance that they will be chosen, thus leading to less switching and

better performance. Also, that 2CR would fare better in Experiment 3 than Experiment 4 is understandable because it has no means of representing interdependencies between options. It is also reasonable that participants would be more likely to employ the more demanding two-component learning procedure in the simpler environment of Experiment 3. The more complex dynamics in Experiment 4 may have led participants to focus their efforts on simply learning the current value of each option.

Overall, the analysis provides some support for previous work arguing for the inclusion of eligibility traces in models of human reinforcement learning (e.g. Bogacz et al., 2007). Additionally, the success of 2CR, 2C, and Cull highlights the importance of testing and incorporating additional learning and choice mechanisms into future research. Indeed, Figure 9 shows very strong evidence for several models at the individual level, with some AIC weights near 1 for Cull, ET, 2CR, and 2C. That said, these conclusions depend on the one step ahead procedure used to estimate the models' parameters. Critically, by providing the model with an individual's full history of actions and rewards this method "yokes" the model to the data. This prevents errors from compounding over time because the model is "reset" after each action, regardless of whether or not its previous action was correct. This procedure is computationally efficient, but it ensures that model predictions never stray too far from data. In the next section, we take a different approach in order to examine how the models perform without this guidance.

5.3. Removing the Yoke Reveals Systematic Prediction Failures

In the second part of our analysis we focus on the top five models (Figure 9). The critical difference between this exercise and the previous is that, rather than fitting the models to response data using the one step ahead method, we examined the models' unyoked performance in the choice environment. More specifically, we took the best-fitting parameters estimated earlier, and used them to simulate the behavior of each model (sometimes referred to as the *simulation method*). Models did not

have access to the participants' reward and response histories, and instead received rewards based on their own actions. We were interested in how well each model could capture participants' behavior after removing this link to the data because previous research has demonstrated model fitting and simulation can produce divergent results (Steingroever, Wetzels, & Wagenmakers, 2014)¹⁴. To this end we used the best-fitting parameters for each individual and model to simulate 500 iterations of the skilled bandit task.

Table 3

Sum of Squared Errors

	Cull		DL		ET		2CR		2C	
	choices	switching	choices	switching	choices	switching	choices	switching	choices	switching
Experiment 2	0.476	0.815	0.104	0.410	0.084	0.438	0.122	0.297	0.110	0.309
Experiment 3	0.553	0.625	0.068	0.496	0.074	0.515	0.061	0.329	0.055	0.374
Experiment 4	1.319	2.224	0.223	2.320	0.181	2.111	0.159	1.745	0.162	1.851

As a measure of performance, we averaged the predictions of each model across iterations and individuals and compared this to data. We calculated two measures of fit summarized in Table 3. The first compares predicted choice proportions to participants' mean choice proportions (across games and individuals) for each trial. The second compares predicted switching rates to participants' mean switching rate (across games and individuals) on each trial. Examining the sum of squared errors (SSE) for these two measures reveals a pattern slightly different from that in Table 2. For none of the comparisons was DL the best model, despite being the best fitting model in two experiments according to the one-step-ahead method. ET produced the smallest deviation for choices in Experiment 2, but the two-component models produced the best fits for the five remaining measures. 2CR was the best predictor of switching in each experiment, with 2C coming second. For choices, 2C produced the best fit in Experiment 3, and 2CR was the best in Experiment 4. The strong performance of the two-component

¹⁴ Though see Halpern and Gureckis (2011) for a demonstration of successful parameter recovery using the Q-Learning model.

models lends more support for the idea that people attend to both immediate value and growth trends. That their relative performance would improve under simulation suggests a robustness to errors that may have aided them in mimicking participants' behavior when not yoked to data.

Still more can be learned by plotting model predictions. Figures 10-12 reveal a collection of successes and systematic misfits across experiments. In Figure 10, we see that for Experiment 2, while Cull fails to produce the amount of learning found in the data, the remaining models effectively reproduce the observed choice proportions and switching rates across trials. In fact, there is very little to distinguish them from one another with regard to choices, switching and stopping. The most notable misfits to data pertain to stopping. Stopping in the first ten trials was common for all of the models, except Cull, though participants rarely showed this. Participants almost never stopped in the final ten trials, but all models (especially Cull) frequently produced this behavior.

Figure 10. Simulation results for the top models (columns), across trials (x-axis) in Experiment 2. The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure 11 shows predictions for Experiment 3, and reveals a similar pattern. Again, Cull produces less pronounced changes in choice proportions, with choice rates for the bad options never falling much

below 0.10. The other models were again difficult to distinguish qualitatively. Switching rates and choice proportions for these models closely matched those of participants, though they predict a slightly exaggerated crossover between L and T. Also, as we found in Experiment 2, all models produced qualitatively incorrect predictions for stopping point, characterized by too much early stopping and far too much late stopping.

Figure 11. Simulation results for the top models (columns), across trials (x-axis) in Experiment 3. The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Experiment 4 clearly poses the greatest challenge to the models, and Figure 12 shows notable misfits throughout. Cull's performance is again quite poor. For the remaining models, although the predicted choice proportions and switching rates resemble those of participants, all struggle to reproduce some aspects of the data. Each model overestimates people's preference for the lure over the target. This suggests that the models were more often led astray by the initially higher rewards offered by L, and struggled more than some participants to recognize the long-term advantage of T. The models also showed substantially higher switching rates throughout. In this respect, the models show less adaptation than participants to the use-it-or-lose-it dynamics of Experiment 4. Similarly, the models'

predictions regarding late stopping provide a particularly poor fit in Experiment 4 where most participants were careful to avoid this costly behavior. The prevalence of this misfit across all models suggests that the softmax decision rule does not adequately capture the aspect of behavior. Future work should therefore investigate alternative decision processes that better explain the strategically abrupt nature of individuals' exploration behavior.

Figure 12. Simulation results for the top models (columns), across trials (x-axis) in Experiment 4. The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

In sum, the results of our simulation analysis show that, across three unique environments, four models – 2CR, 2C, DL, and ET – behave in ways quite similar to participants. The similar performance of several models here fits with the mixed results found earlier in Section 5.2. However, several unique insights can be gleaned from these simulations. Departing from what we found in the one-step-ahead model fitting exercise, Cull performed poorly and was clearly inferior to the other models. Conversely, the 2CR and 2C were quantitatively superior to DL and ET (though these four models were qualitatively indistinguishable). We also found that the models failed to capture important aspects of human

behavior. Most notably, all models consistently showed later exploration and stopping compared to participants, particularly in Experiment 4 where this was very rare.

6. General Discussion

One of the key features of sequential decision making in real life is the need to plan in a way that takes account of the *reactivity* of the environment. That is, to get by in the real world we need some understanding of how our past actions shape the potential rewards associated with future actions. Here we have focused on the skilled bandit task, a resource cultivation problem in which people need to solve an explore-exploit dilemma in a reactive world. Across four experiments, we found that people tended to make adaptive choices. They used early exploration to identify high growth options, then exploited those options for the remainder of the task. In some cases (Experiment 2) providing detailed instructions had no effect on people's ability to learn, whereas in others (Experiments 3 and 4) there is evidence for a mild effect. Overall though, the vast majority of learning was self-guided, with people learning from experience (see also Weiss-Cohen, Konstantinidis, Speekenbrink, & Harvey, 2016).

Arguably, the most interesting findings (in our view) are those specifically pertaining to learning of reactivity. In Experiment 1, we directly measured beliefs about the relationship between choice, time, and reward by eliciting predictions for imagined future actions. These predictions clearly indicated that participants expected options' rewards to grow only when chosen. In Experiments 2-4 we posed participants with the more difficult task of learning reactive dynamics in an environment with four alternatives and more complex, overlapping reward functions. Here too, participants were successful in efficiently exploring the environment and eventually exploiting the best options. New and important insights can also be found by comparing behavior across experiments. For example, differences in behavior across Experiments 3 and 4 suggest that people adapt their strategies in response to the choice environment: in particular, when the reactive environment has a "use it or lose it" rule (Experiment 4),

over-exploration is costly and participants learned to minimize exploration in order to avoid reward decrements.

6.1. Modeling Insights

In addition to these behavioral findings, new insights emerge when we consider various models of learning and decision making. The approach we took to modeling human performance in this task drew upon a variety of different perspectives. We included simple learning models (Naïve and Recency), as well as more complex ones (DL and ET). We tested models inspired by the judgment and decision making literature (Threshold and Cull), and models that use the statistical structure of the environment (2CR and Linear). When fit to data in a standard one-step-ahead method, evidence for several distinct groups emerged. Across all three experiments, the largest number of participants were fit by DL, followed by ET, 2CR, Cull, and 2C. The success of 2CR and Linear-R compliments our findings from Experiment 1 and suggests that some individuals (14%) developed strategies that exploited the reactive dynamics of the task environment. On the other hand, the good performance of DL and ET implies that many people employed relatively simple learning mechanisms. We also found that some individuals employed unique decision procedures, such as the sequential elimination process described by Cull. All together, these findings show that a diverse set of mechanisms is needed to account for the diversity of behavior we observe.

With our second model analysis, we removed the yoke and simulated the models within the task environment. Interestingly, we found that the superiority of DL and ET over the other models disappears, and 2CR and 2C produced the best account, even if predictions from these four models were qualitatively similar. Additionally, simulations revealed systematic prediction failures in all models. This is particularly evident when we compare the stopping time distributions produced by models to those produced by the participants. At some point in the task (usually about 2/3 of the way through) participants appear to make a decision to cease *all* future exploration, and exploit until the task

deadline. Arguably this reflects an adaptation to the finite horizon nature of the decision task. None of the models here properly account for the finite horizon, suggesting that a more sophisticated decision policy is required (see, e.g. Navarro et al., 2016). More generally, this issue points to a missing element in the models. Given that it seems unlikely that people solve the notoriously difficult problem of computing the optimal partially-observable Markov decision process (POMDP; e.g. Kaelbling, Littman, & Cassandra, 1998) in real time, it is an open question how people make efficient, goal-directed choices during strategically critical periods of the task. We expand on this idea in the following section.

6.2. What Is the Goal of a Psychological Model?

Our work has several implications for the study of learning and decision making. On the one hand, we find that reinforcement learning models – particularly DL and ET – are robust tools for predicting and simulating behavior. In this sense, we join a growing number of researchers demonstrating the power of RL models to expand our understanding of complex choice behavior (Bogacz et al., 2007; Fu & Anderson, 2006; Gureckis & Love, 2009; Speekenbrink & Konstantinidis, 2015; Sun, Slusarz, & Terry, 2005). On the other hand, our simulation results highlight the limitations of this approach, with all models showing systematic failures. How, then, are we to overcome these limitations? The answer to this question is more complicated than it seems, and revolves around the issue of what we want in a model. Is our goal to formally articulate simple laws and general principles that characterize behavior? If so, DL provides a good model of behavior in Experiments 2 and 3. But perhaps we want our models to capture something deeper about the psychological processes at play. In this case, such simple RL models may be lacking.

Models with rich representations of cognitive processes have been successfully applied to many choice phenomena (see Pleskac, 2015) and have become increasingly influential across various fields of decision science (Busemeyer, 2015). Similarly, our results show that models that incorporate choice heuristics (Cull) and multi-attribute integration (2CR) match and sometimes exceed the performance of

RL models. However, complicated knowledge representations and decision mechanisms may be superfluous in many research contexts. For example, the DL model maintains a relatively simple representation of the task environment – only storing estimates of each option’s current value – but often gave a better account of behavior than models with additional learning mechanisms (e.g. ET, 2CR). Should we be bothered that our models ignore seemingly important components of the psychological process? If we regard the models as addressing the computational level of analysis (see Marr, 1982), then it is enough to describe the general characteristics of the decision problem people face, without addressing the algorithmic and implementation details. However, our simulation analysis exposes the limitations of this approach, with all models failing to capture key patterns of behavior. This suggests the need for some additional elements.

Absent from the models we consider here are higher-order cognitive elements to support functions like inference, reasoning, strategy selection, and adaptation¹⁵. Indeed, we find evidence that participants approach the skilled bandit task with a sophistication no model considered here can capture. For example, we find that performance improves across games, with participants making better choices, exploring less (see Figures A2, A5, and A6), and earning more points (see Figure 7). Most striking are the changes in stopping point, with late exploration becoming quite rare in later games. Another demonstration of adaptation across games can be seen in the degree to which individuals explored in *streaks* – selecting an option multiple times in order to assess its current value and growth rate – versus *sweeps* – cycling through options repeatedly. We find that the average streak length increases across games as people learn to explore more efficiently (see Figure A7).

We also find behavioral differences across experiments that indicate people’s ability to tailor their strategies to the particulars of their environment. In Game 1 of Experiment 4, participants switched

¹⁵ In a separate analysis, we attempted to model change across games by fitting models separately to Games 1 and 5. Unfortunately, fitting to only one game produced weak evidence, and the results were therefore uninformative. A report of this analysis can be found at <http://osf.io/p7bky/>.

more, stopped later, and explored in shorter streaks than their counterparts in Experiment 3 (Figure A7). This suggests that the greater complexity of the Experiment 4 environment spurred participants to investigate the dynamics relating their actions to rewards. However, the “use it or lose it” nature of the environment made excessive exploration costly. Not surprisingly, we find that by the end of the experiment, each of the above trends reversed. In Game 5 of Experiment 4, participants switched less, stopped earlier, and explored in longer streaks than those in Experiment 3. Apparently, with greater experience participants accumulated knowledge about the reward functions and environmental dynamics, and were better able to combine prior knowledge with proximal feedback to optimize their actions. These patterns suggest the need for additional mechanisms for modeling this kind of cumulative learning and adaptation. It remains a goal for future research to develop new models that account for this higher-order learning. Also, sequential mechanisms like Q-learning and SARSA (Sutton & Barto, 1998) may provide a method for modeling the more strategic elements of people’s behavior, though implementing these is complicated by the very large number of states in the skilled bandit task¹⁶.

Finally, the question of how to conceptualize individual differences is important. In the past, some have assumed a single learning mechanism, relying on stochastic decision rules to explain individual differences (e.g. Bogacz et al., 2007). Another approach is to also allow individuals different learning parameters within the same model framework. Our results suggest that neither approach is sufficient to account for the diversity of behavior observed across individuals. We find that participants are best characterized as groups of individuals applying fundamentally different strategies, with parameter differences accounting for behavioral differences within each group. That said, sampling error and model mimicry may contribute to these results, so more work is needed to establish the true level of strategic heterogeneity.

¹⁶ Modern methods for online planning in large POMDPs (Silver & Veness, 2010) also provide an avenue for future investigations.

6.3. Future Directions

Having laid the groundwork for further investigations into the strategies people use to learn and choose in dynamic growth environments, several interesting possibilities for future work arise. One particularly promising idea would be to directly compare learning and strategic adaptation in growth and depletion environments. Comparing performance across environments with that of an optimal model can shed light on the factors that facilitate and hinder learning. For example, are growth environments inherently easier to learn than depletion environments? How do the dynamics of the task interact with people's priors about the nature of the environment (cf. Schulze & Newell, 2015; Schulze, van Ravenzwaaij, & Newell, 2017)?

Another direction for future work is to explore the question of what knowledge is transferred from one game to the next. Consistent with Navarro et al. (2016) we found evidence for changes in behavior across games. Most notably, familiarity with the task allowed people to “do more with less”, earning more points with less exploration. This pattern of performance suggests that people learn the structure of the problem and transfer that knowledge appropriately. In previous work Acuña and Schrater (2010) argued that human performance on bandit problems does reflect some form of structure learning (Tenenbaum, Griffiths, & Kemp, 2006), albeit in a very different context to the skilled bandit task. With this in mind it would be interesting to investigate how people learn across tasks and not just within them.

6.4. Conclusion

Our findings suggest that individuals learned to make adaptive decisions over time in complex, dynamic environments. Across games, participants became more efficient on average, earning higher rewards with less exploration. We find strong and convergent evidence – from predictions, choices, and computational modeling – that individuals successfully learned about the reactive nature of reward growth. The success of DL and ET demonstrates the power and utility of these models, while the good

performance of 2C, 2CR, CULL indicates the prevalence of more sophisticated strategies. These results provide new and valuable insights into the learning and decision making process, and pave the way for further development of psychological theories that incorporate higher-level cognitive mechanisms without reducing generality or introducing excessive complexity.

Author Note

This research was supported by funding from the Australian Research Council (Grants No. DP160101186 and DP150104206). We thank Todd Gureckis and Maarten Speekenbrink for their valuable feedback regarding this research.

References

- Acuña, D. E., & Schrater, P. (2010). Structure Learning in Human Sequential Decision-Making. *PLoS computational biology*, 6(12), e1001003. doi:10.1371/journal.pcbi.1001003
- Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Caski (Eds.), *Selected papers of Hirotugu Akaike* (pp. 199-213). New York, NY: Springer.
- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, 89(4), 369-406.
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 25-35.
- Bogacz, R., McClure, S. M., Li, J., Cohen, J. D., & Montague, P. R. (2007). Short-term memory traces for action bias in human reinforcement learning. *Brain research*, 1153, 111-121.
- Busmeyer, J. R. (2015). Cognitive science contributions to decision science. *Cognition*, 135, 43-46.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876.
- Ejova, A., Navarro, D., & Perfors, A. (2009). When to walk away: The effect of variability on keeping options viable. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, 1258-1263.
- Fu, W.-T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, 135(2), 184-206.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gureckis, T. M., & Love, B. C. (2009). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition*, 113(3), 293-313.
- Halpern, D., & Gureckis, T. (2011, October). On the identifiability of parameters in reinforcement learning models [Blog post]. Web log post Retrieved from <http://gureckislab.org/blog/?p=3450>
- Heathcote, A., Brown, S., & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, 7(2), 185-207.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2), 99-134.
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman.
- 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(3), 191-215.
- Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology*, 44(1), 190-204.
- 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.
- Navarro, D. J., Pitt, M. A., & Myung, I. J. (2004). Assessing the distinguishability of models and the informativeness of data. *Cognitive psychology*, 49(1), 47-84.
- Navarro, D. J., Tran, P., & Baz, N. (2018). Aversion to option loss in a restless bandit task. *Computational Brain & Behavior*, 1(2), 151-164. doi:10.1007/s42113-018-0010-8
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7, 308-313.
- Neth, H., Engelman, N., & Mayrhofer, R. (2014). Foraging for Alternatives: Ecological Rationality in Keeping Options Viable. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, 1078-1083.

- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition*, 1, 1-55.
- Newell, B. R., Koehler, D. J., James, G., Rakow, T., & Van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, 41(3), 329-338.
- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, 14(6), 1133-1139.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Pleskac, T. J. (2015). Learning models in decision making. *The Wiley Blackwell Handbook of Judgment and Decision Making*, 629-657.
- Rich, A. S., & Gureckis, T. M. (2018). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General*, 147(11), 1553-1570. doi:10.1037/xge0000466
- Rubin, D. C., Hinton, S., & Wenzel, A. (1999). The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(5), 1161.
- Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3), 210-229.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(6), 927-943.
- Schulze, C., & Newell, B. R. (2015). Compete, coordinate, and cooperate: How to exploit uncertain environments with social interaction. *Journal of Experimental Psychology: General*, 144(5), 967-981.
- Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2017). Hold It! The Influence of Lingering Rewards on Choice Diversification and Persistence. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.
- Shin, J., & Ariely, D. (2004). Keeping doors open: The effect of unavailability on incentives to keep options viable. *Management science*, 50(5), 575-586.
- Silver, D., & Veness, J. (2010). *Monte-Carlo planning in large POMDPs*. Paper presented at the Advances in neural information processing systems.
- Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2), 351-367.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2014). Absolute performance of reinforcement-learning models for the Iowa Gambling Task. *Decision*, 1(3), 161.
- Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53, 168-179.
- Sun, R., Slusarz, P., & Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review*, 112(1), 159.
- Sutton, R. S. (1984). *Temporal credit assignment in reinforcement learning*. University of Massachusetts, Amherst, MA.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). Cambridge: MIT press.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *TRENDS in Cognitive Science*, 10(7), 309-318.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4), 281-299.

- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71, 680.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11(1), 192-196.
- Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, 135, 55-69.
- Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, 19(1), 1-16.
- Yi, M. S., Steyvers, M., & Lee, M. (2009). Modeling human performance in restless bandits with particle filters. *The Journal of Problem Solving*, 2(2), 5.

Appendix

In Experiment 1, rewards for both options grew linearly. Initial values for Options S and F were 10 and 1, respectively. With each new action, the expected reward for S and F increased by 2 and 4 points, respectively.

In Experiments 2-4, each option's reward function was defined by a starting point (20 for all options), asymptote (ψ), rate (γ), and the growth equation:

$$r_t = r_{t-1} + \gamma \cdot (\psi - r_{t-1}), \quad (1)$$

where r_t is the expected reward on Trial t . Table A1 gives the asymptote and rate values for each option in Experiment 2 and Experiment 3.

Table A1

<i>Payoff Function Parameters</i>								
	Experiment 2					Experiment 3		
Option	HF	LF	HS	LS	T	L	B1	B2
Intercept	20	20	20	20	20	20	20	20
Asymptote	40	30	40	30	70	30	30	30
Rate	0.10	0.10	0.05	0.05	0.06	0.20	0.02	0.01

Now that you have finished this game, we would like you to imagine what would happen if you were to continue playing the game for five more turns. Imagine that you were shown the same pair of options:

DAX

WUG

Specifically, we want you to pretend that your next five choices would have been: WUG, WUG, WUG, WUG, and then DAX. In the questionnaire below we would like you to make your best guess about how many points you would receive if you made these choices over the next five turns.

(In case you have forgotten: the last time you chose DAX you received 37 points and the last time you chose WUG you received 94 points).

Choice 1: WUG. What reward do you expect?

Choice 2: WUG. What reward do you expect?

Choice 3: WUG. What reward do you expect?

Choice 4: WUG. What reward do you expect?

Choice 5: DAX. What reward do you expect?

Figure A1. Predictions screen from Experiment 1.

A1, A4, and A5 show the major experimental results, for each game, after collapsing across instructions conditions in Experiments 2-4, respectively. Although behavior was similar in each game, incremental changes can be seen, particularly in the stopping point distributions.

Figures A2 and A3 show the major experimental results in the vague and detailed instructions conditions, respectively. Behavior was similar across conditions, though participants appeared to choose the target more in Experiments 3 and 4 when given detailed instructions.

Figure A6 plots the mean number of consecutive choices for the same option across a moving window of twenty trials. Streak length increased across trials and games, though game effects were more pronounced in Experiment 4.

Figure A2. Results from Experiment 2, across trials (x-axis) and games (columns). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure A3. Results from the vague instructions condition of each experiment (columns), across trials (x-axis). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure A4. Results from the detailed instructions condition each experiment (columns), across trials (x-axis). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure A5. Results from Experiment 3, across trials (x-axis) and games (columns). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure A6. Results from Experiment 4, across trials (x-axis) and games (columns). The top row shows mean choice proportions. The middle row shows switching probability. The bottom row shows stopping point distributions.

Figure A7. Mean number of consecutive choices for the same option within a moving window of twenty trials. X-axis values indicate the final trial in the moving window. Means were calculated at the individual level.