



Contents lists available at ScienceDirect

Journal of Mathematical Psychology

journal homepage: www.elsevier.com/locate/jmp

Learning in noise: Dynamic decision-making in a variable environment

Todd M. Gureckis^{a,*}, Bradley C. Love^b^a New York University, United States^b The University of Texas at Austin, United States

ARTICLE INFO

Article history:

Received 14 January 2009

Received in revised form

12 February 2009

Available online 2 April 2009

ABSTRACT

In engineering systems, noise is a curse, obscuring important signals and increasing the uncertainty associated with measurement. However, the negative effects of noise are not universal. In this paper, we examine how people learn sequential control strategies given different sources and amounts of feedback variability. In particular, we consider people's behavior in a task where short- and long-term rewards are placed in conflict (i.e., the best option in the short-term is worst in the long-term). Consistent with a model based on reinforcement learning principles [Gureckis, T., & Love, B.C. Short term gains, long term pains: How cues about state aid learning in dynamic environments. *Cognition* (in press)], we find that learners differentially weight information predictive of the current task state. In particular, when cues that signal state are noisy, we find that participants' ability to identify an optimal strategy is strongly impaired relative to equivalent amounts of noise that obscure the rewards/valuations of those states. In other situations, we find that noise and noise in reward signals may paradoxically *improve* performance by encouraging exploration. Our results demonstrate how experimentally-manipulated task variability can be used to test predictions about the mechanisms that learners engage in dynamic decision making tasks.

© 2009 Elsevier Inc. All rights reserved.

Millions of Americans suffer from chronic illnesses such as heart disease and diabetes, and must carefully monitor their diet and exercise. However, making healthy lifestyle choices can be difficult. Due to the intrinsic variability in outcomes and observations, decision makers face uncertainty about both the actual state of their health and the costs and benefits associated with various dietary options, complicating the management of illness. For example, daily fluctuations in blood pressure and glucose levels can obscure estimates of actual health. This difficulty in separating signal from noise extends beyond understanding our physical state to the external environment. For example, it is difficult to obtain reliable information on the sodium, fat, and calorie content of foods. Under these circumstances, making effective decisions requires one to manage uncertainty from many sources.

In this paper, we explore how people learn effective decision-making strategies given similar kinds of uncertainty about what is signal and what is noise. Decision-making under uncertainty has played a central role in judgement and decision-making research. However, empirical attempts to understand this ability have often focused on decisions made in static, one-off situations based on

verbal descriptions of choice alternatives, as in gambles (Barron & Erev, 2003). More recently, researchers have adopted decision-theoretic approaches to understanding choice behavior in more realistic, online, and dynamic situations (Edwards, 1962, see Busemeyer, 2002 for a recent review). Like the real-world example of managing a chronic illness, participants in these tasks are asked to achieve a particular goal by making a sequence of decisions from one moment to the next based on their ongoing experience (Gureckis & Love, in press; Stanley, Mathew, Russ, & Kotler-Cope, 1989).

In our experiments, we examine human learning in a dynamic decision-making task, called the "Farming on Mars" task, where the experienced reward structure continually evolves in response to the actions of the individual (Gureckis & Love, in press). A key feature of our task is that the strategy that returns the most reward over the course of the experiment requires participants to forego immediately attractive short-term options in favor of a long-term beneficial strategy (Herrnstein, 1991; Herrnstein & Prelec, 1991). Just like the dilemma between deciding whether to eat a healthy meal or indulge in a higher calorie dessert, participants who seek to maximize their long-term well-being must forego immediately attractive alternatives. The extension presented here is to consider how decision maker's ability to uncover a non-obvious reward-maximizing strategy is impacted when relevant outcomes in the task are obscured by variability or noise (i.e., decisions are made with increasing uncertainty about the value of the mean).

* Corresponding address: Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, United States.

E-mail address: todd.gureckis@nyu.edu (T.M. Gureckis).

The problem of learning an effective response strategy in our task is one of adaptive control – participants interact with a system and are asked to make a continuous sequence of decisions which guide the dynamics of the system (Brehmer, 1992). We model human behavior in the task using a framework based on Markov Decision Processes (MDPs) and an approximate solution method known as Reinforcement Learning (RL). RL is an agent-based approach to learning through interaction with the environment in pursuit of reward-maximizing behavior (Sutton & Barto, 1998). The focus of RL research is to understand how a situated agent interacting with a responsive environment can arrive at effective strategies, making RL an excellent tool for studying human learning and decision-making in dynamic tasks (Bussemeyer & Pleskac, in this issue; Fu & Anderson, 2006; Sun, Slusarz, & Terry, 2005). Interestingly, the RL model that we develop in the later section makes specific predictions about the effects that different sources of noise may have on learning performance, which we then test in our experiments.

0.1. Noise as a Curse and Noise as a Tool

Anyone who has tried talking on a bad telephone connection would likely view noise as a curse that obscures important signals and increases the uncertainty associated with measurement. In most systems, decreasing signal-to-noise ratios results in lower performance (Green & Swets, 1966). Likewise, noise or variability, when interpreted by learners as signal, can guide behavior in surprising (and potentially maladaptive) ways. For example, Skinner (1948) showed how pigeons in an operant conditioning experiment, which were given reinforcement at random intervals, enacted a number of repetitive but “superstitious” behavioral patterns. Similarly, failures of statistical reasoning such as the regression to the mean and the gambler’s fallacy, reflect the human tendency to inappropriately view normal variation (i.e., noise) as signal (Kahneman & Tversky, 1973; Tversky & Kahneman, 1971). However, the negative effects of noise are not universal. For example, noise can help overcome local minima such as in systems based on simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983). In addition, in some environments, the inability to effectively separate signal and noise may, paradoxically, lead to improvements in behavior by encouraging alternative strategies or by focusing learners on alternative information sources (Burns, 2004). In this paper, we examine situations where variability or noise actually help people make decisions that maximize their long-term benefit.

In the laboratory, experimenter-manipulated noise is used as a tool to illuminate the structure of human perceptual and cognitive systems, since the way the performance degrades in noise can reveal aspects of system architecture (Gold, Sekuler, & Bennett, 2004; Green & Swets, 1966; Lu & Doshier, 1999; Pelli & Farell, 1999). For example, techniques such as equivalent input noise (Pelli & Farell, 1999) estimate an observer’s internal perceptual noise level by systematically degrading external stimuli. We adopt a similar perspective of “noise as a tool” in the experiments that follow. In particular, we parametrically degraded the signal-to-noise characteristics of feedback given to participants in the Farming on Mars task in order to estimate how such variability impacts performance. Before describing our experimental manipulations in more detail, we begin by describing the basic version of the Farming on Mars task and briefly review some previously reported findings using this paradigm.

0.2. The farming on mars task

In the “Farming on Mars” task utilized in our experiments, participants interact with a repeated, two-choice decision-making




Fig. 1. The payout function for the Farming on Mars task. The horizontal axis is the number of choices out of the last ten in which the Long-Term robot was selected. The vertical axis is the number of oxygen units generated as a result of choosing one of the robots on a trial. The two diagonal lines show the reward associated with each robot for each state. By design the Short-Term robot is better for every state (i.e., trial), but the best long-term strategy is to exclusively choose the Long-Term robot because the selection of the Short-Term robot transitions the state to the left, whereas selection of the Long-Term robot transitions the state to the right.

task presented as a simple video game. The cover story for the game is that two agricultural robots have been sent to the planet Mars in order to establish a farming system capable of generating oxygen for later human inhabitants. Participants are informed that each robot specializes in a different set of farming practices, but that only one robot can be active at a given moment. Participants’ job as controller is to select, on each trial, which robot should be employed in order to maximize the total oxygen generated over the entire experiment. Participants indicate which robot should do the farming, and are given feedback about how much oxygen was generated on that trial as a result of their choice.

Unknown to participants at the start of the task, there is a contingency between recent robot selections and the oxygen points received on the next trial. In particular, the current payoff depends on the relative allocation of responses to one of the two choice options presented on each trial. For expository purposes we will refer to one of the choice options as the “Short-term robot” and the other as the “Long-Term robot” although these labels were not provided to participants in the experiment. Fig. 1A shows an example of the payoff structure used in the task. The horizontal axis in Fig. 1 measures the participant’s current allocation to the Long-Term robot over the last ten trials (ranging from 0 to 10). The two diagonal lines describe the function relating the current choice history to the reward at any point in time. The upper diagonal line illustrates the reward received from selecting the Short-Term robot as a function of recent choice history, and the lower diagonal line shows the reward from selecting the Long-Term robot.

Note that the Short-Term robot always generates more oxygen than the Long-term robot on any given trial (i.e., the reward function for the Short-Term robot reward function is larger than for the Long-Term robot in every situation). However, each time the Short-Term robot is selected, the output of both robots is lowered on the following trial (i.e., the state of the system shifts to the left along the horizontal axes in Fig. 1). Selections of the Long-Term robot behave in the opposite fashion. Each time this robot is selected, the output of both robots is *increased* on the next trial (shifted to the right in Fig. 1A). Critically, over the window of the last ten trials, the reward received from repeatedly selecting the Long-Term robot exceeds that from always selecting the Short-Term robot (i.e., the highest point of the Long-Term robot curve is above the lowest point for the Short-Term robot curve in Fig. 1A). As a result, the optimal strategy is to select the Long-Term robot on every trial, even though selecting the Short-Term robot would

earn more on any single trial. In the experiment, participants are not given any relevant information about the differences between the robots, and thus can only arrive at the optimal strategy by interactively exploring the behavior of the system (cf. Berry and Broadbent (1988) and Stanley et al. (1989)).

The structure of the Farming on Mars task borrows from a growing literature looking at choice behavior in situations where actions that lead to long-term rewards conflict with those that yield immediate rewards (Herrnstein, 1991; Herrnstein & Prelec, 1991; Neth, Sims, & Gray, 2006; Tunney & Shanks, 2002). Interestingly, the conclusion from much of this work has been that both humans and other animals often fail to inhibit the tendency to select an initially attractive option even when doing so leads to lower rates of reinforcement relative to other strategies, a phenomena referred to as *melioration*. While melioration is sometimes taken as evidence that people fail to maximize their long-term expected utility (see Tunney and Shanks (2002) for a similar discussion), in practice, it is a significant challenge for learners to discover the optimal strategy in an unknown environment.

0.3. Why is the “Farming on Mars” task so difficult? Effective exploration and the problem of perceptual aliasing

The Farming on Mars task is a challenging task for a number of reasons. First, from the outset of the task, participants may have little sense of how their own actions are influencing the behavior of the system. Effective learning requires exploration of hidden contingencies between the agent’s past actions and future prospects. For example, work with a similar task found that manipulations that encourage subjects to “explore” the system by lowering the costs of making early choices can lead to improvements in later performance (Tunney & Shanks, 2002). Similarly, recent work has shown that motivational manipulations can influence the degree to which participants are willing to adopt exploratory strategies in sequential choice tasks, which can in turn translate into improvements in performance in tasks with non-obvious solutions (Worthy, Maddox, & Markman, 2007).

In light of these findings, one hypothesis is that, under certain conditions, performance may actually improve when outcomes in the task are more variable. Imagine trying to find the “globally” best restaurant out of a large set. If you have an enjoyable meal at the first place you visit, you may be less willing to sample other places. However, if each time you visit this restaurant, the quality varies due to random noise, it may coincidentally help encourage more extensive sampling of other options (Denrell, 2005). A bad experience one day might lead you someplace else for your next meal. In Experiment 1, we set up an analogous situation in the Farming on Mars task by parametrically varying the amount of variability associated with the reward signal in the task and assessing its impact on participants’ ability to discover the reward-maximizing strategy.

Effective exploration is one of multiple challenges facing learners in the Farming on Mars task. A second, but related challenge arises from participants’ mental representation of the task. Each time a participant makes a choice in the task, the system changes so that the reward received on the next trial is different than it was on the previous trial. Recognizing how the current situation or “state” of the world is changing as a result of the one’s actions is essentially a categorization problem. The learner must identify when the current situation is different, and how to generalize their experience from one situation to the next (Redish, Jensen, Johnson, & Kurth-Nelson, 2007; Veksler, Gray, & Schoelles, 2007). In the standard version of the task, identifying these changes is difficult given that there are few direct cues available indicating to subjects that the world is changing on

each trial. As a result, participants must overcome the problem of *perceptual aliasing* in which relevant states in the world are poorly differentiated (Whitehead & Ballard, 1991). Returning to the restaurant example, imagine trying to find the best restaurant when not only the average quality of the food is obscured by randomly distributed noise, but there is also uncertainty about which restaurant one is dining at because all restaurants are identical inside. Failing to differentiate situations associated with different rewards makes learning the true reward structure a difficult task.

The issue of how human learners use cues in the environment to disambiguate distinct states or “situations” remains an active question in reinforcement learning research (Daw, 2003; Gureckis & Love, in press; Redish et al., 2007; Veksler et al., 2007). However, in a recent set of studies (Gureckis & Love, in press), we evaluated how simple visual cues can impact participants’ performance in the Farming on Mars task. In one experiment (Gureckis & Love, in press, Experiment 2), displays were augmented with a row of indicator lights. The position of the active light in the display served as a cue about the current state of the system. Participants given cues which correlated with the underlying system state performed better than participants attempting to learn without these cues. In addition, cues which allowed generalization from one situation (i.e., state) to other states performed best. These studies demonstrate the importance of state representation in dynamic and complex task environments. The state representation that the learner adopts may act as a “framework” for effectively structuring, integrating, and generalizing experience. However, to the degree to which participants use information about the current state to structure and integrate their learning experiences, variability on such cues should have a dramatic influence on performance. Variability in the cues that signal the current task state effectively increase the degree of perceptual aliasing by making one unclear at any point in time how state knowledge should be updated given current rewards. In the restaurant example, a diner may know the meal is enjoyable, but be confused about which restaurant served the meal.

0.4. Overview and summary

Previous studies establish that two somewhat similar sources of information may jointly influence performance in sequential decision-making tasks. The first is the structure of the rewards in the task (i.e., the payoff function). The second is information about functionally distinct task states (any cue which can help distinguish or categorize different states or situations). These two sources of information are both functionally and psychologically distinct. State information indicates the agent’s place in the overall system and can help link future actions, whereas reward information provides the valuation of those actions. Given that these two sources of information or signals play a distinct role in learning, we predicted that learners would respond differently to noise (i.e., variability) associated with either signal.

The remainder of the paper is structured as follows. We begin by describing the results of two novel experiments. In the first study, participants attempted to uncover a reward-maximizing response strategy given different amounts of noise that obscured the valuation of particular actions. In contrast to the conventional view that noise has a uniformly negative impact on system performance, we find that moderate amounts of noise actually *improves* the ability of participants to find an optimal response strategy by increasing their tendency to explore alternative strategies. In the second study, we compare the effect that different *sources* of noise have on performance in the task. In particular, we compare participants’ performance when comparable levels of noise obscure either the reward or state

signals. We find that variability that impairs participants' ability to identify the current task state hurts performance more than noise that obscures the valuation of those states. This result is surprising given that attending to information about the current task state is, in some ways, optional for accomplishing the primary task of maximizing reward. Simulations show that a simple model based on RL principles, including the balance of exploration and exploitation, the appreciation of future outcomes, and the identification of distinct task states, provides an excellent account of behavior across the two experiments. Our results illuminate how experimentally-manipulated task variability can be used to test predictions about the mechanisms that learners engage in dynamic decision making tasks.

1. Experiment 1

Experiment 1 evaluates the impact of variability in the reward signal on task performance. Using the Farming on Mars task, participants were assigned to one of four conditions which were identical with respect to the number of trials and the underlying payoff function, but which differed in the way in which variability was added to the experienced rewards. In each condition, the dependence of current reward on past choices was determined by the number of choices made to the Long-Term robot over the last ten trials (as illustrated in Fig. 1). In the *no-noise* condition, there were no additional sources of variability. In contrast, in the *low-noise*, *medium-noise*, and *high-noise* conditions, on each trial, a normally distributed random noise (with mean equal to zero and standard deviation σ_r) was added to the payoff indicated in Fig. 1. This variability in the experienced rewards obscured the underlying structure of the task. One straightforward prediction is that performance should degrade as the amount of noise increased (consistent with work finding a decrease in performance as the discriminability of two options is lowered, Bussemeyer and Myung (1992)). On the other hand, moderate amounts of noise might, in some circumstances, actually help participants adopt appropriate exploration strategies. The inconsistent feedback that arises in a noisy environment might rule out simple hypotheses and encourage sampling of alternatives.

1.1. Method

1.1.1. Participants

Ninety-two undergraduates from New York University and the University of Texas participated for course credit and a small cash bonus which was tied to performance. Participants were randomly assigned to one of the four conditions: the no-noise (NN) condition, a low noise (LN) condition, a medium noise (MN) condition, and a high noise (HN) condition. Twenty-three participants were included in each condition.

1.1.2. Materials

The experiment was run on standard desktop computers using an in-house data collection system written in Python. Stimuli and instructions were displayed on a 17-inch color LCD positioned approximately 47 centimeters away from the participant. Participants were tested individually in a single session. Extraneous display variables, such as which robot corresponds to the left or right choice option, were counterbalanced across participants. In addition, no other relevant information distinguished between the two options other than the rewards that the participant received for their choices.

Table 1

Summary of the Conditions in Experiment 1. The column titled *d'* Between Reward Curves measures the discriminability of the two reward curves as a function of the increasing variability. Similarly, the column titled *d'* Between States measures the discriminability of two adjacent states (i.e., between two adjacent points in the same reward curve).

Condition	Noise level	<i>d'</i> Between Reward Curves	<i>d'</i> Between States
No noise (NN)	$\sim N(0, \sigma_r = 0)$	–	–
Low noise (LN)	$\sim N(0, \sigma_r = 70)$	7.14	1.43
Medium noise (MN)	$\sim N(0, \sigma_r = 100)$	5.0	1.0
High noise (HN)	$\sim N(0, \sigma_r = 300)$	1.67	0.33

1.1.3. Design

Participants were given a simple two-choice decision-making task (the Farming on Mars task described above) presented as a simple video game (see Fig. 4 for an example display). At the start of the experiment, participants were presented with instructions on the screen that described the basic cover story and task. Participants were informed that their goal was to maximize the total output from the Mars Farming system over the entire experiment by selecting of one of two robot systems on each trial. Unknown to participants, the number of oxygen units generated at any point in time was a function of their choice history over the previous ten trials. In addition, the payoffs associated with each robot system were manipulated so that one option was better than the other in the long-term, despite appearing worse in the short-term. A graphical depiction of the rewards in the tasks is shown in Fig. 2. If h represents the number of trials in the last 10 which were allocated to the Long-Term robot, then the payoff for any selection of the Long-Term Robot was $400 + 1000 * \frac{h}{10} + N(\mu = 0, \sigma_r)$. Alternatively, the payoff for the Short-Term robot was $900 + 1000 * \frac{h}{10} + N(\mu = 0, \sigma_r)$. At the start of the experiment, we initialized h to 5 (so as to not favor either option). In these rewards equations, $N(\mu, \sigma_r)$ indicates a normally distributed random number with mean μ and standard deviation σ_r . The parameter σ_r varied between conditions. In the NN condition $\sigma_r = 0$, in the LN condition $\sigma_r = 70$, in the MN condition $\sigma_r = 100$, and in the HN condition $\sigma_r = 300$. The effect of the random trial-to-trial noise was to degrade the signal-to-noise characteristics of the reward signal. Table 1 summarizes the difference between conditions in the discriminability of different task states and between the two reward curves while Fig. 2 illustrates the payoff functions in each condition, including the 95% confidence intervals. As is visible in the figure, as the variability increases the discriminability of the two reward curves (i.e., for the Short-term and Long-term options), as well as the transitions between successive task states is lowered.

1.1.4. Procedure

The 500 trials of the experiment were divided into five blocks of 100 trials each. At the end of each block, participants were given a short break and each successive block picked up where the last block left off. In order to maintain motivation, participants were told that they could earn a small cash bonus of \$2–5 which was tied to their oxygen generating performance in the task. However, participants were not told how oxygen points would translate into cash rewards, only that generating more oxygen would yield a larger bonus.

On each trial, participants were shown a control panel with two response buttons labeled either *System 1* or *System 2*. Between these two buttons was a video display where trial-relevant feedback and instructions were presented. Participants clicked one of the two response buttons using a computer mouse. After a selection was made, a short animation (lasting approximately 800 ms) indicated that the response was being sent to the Mars base. Following this animation, the amount of oxygen generated on that trial was shown in numerical terms (i.e., “New Oxygen Added: 800.00”). A short auditory beep was presented when the oxygen

Fig. 2. The reward functions used in each condition of Experiment 1 as a function of the current task state (i.e., the number of Long-term Robot selections in the last 10 trials). The shaded regions show the 95% confidence intervals for each noise condition. See Table 1 for a summary of the effect this noise has on the ability to discriminate the rewards from different states and actions.

Fig. 3. Overall results of Experiment 1. Panel A shows the average proportion of long-term (maximizing) responses made throughout the experiment as a function of noise condition. Panel B presents the average proportion of maximizing responses considered in blocks of 25 trials for all four conditions. In both figures, the horizontal line at 0.5 shows chance performance. In each graph, error bars are standard errors of the mean.

points display was updated, indicating that the reward for that trial had been received. The points display was shown for 800 ms, after which the screen reset to a “Choose” prompt that indicated the start of the next trial. No information about the cumulative oxygen generated across trials was provided.

1.2. Results

1.2.1. Overall performance

Fig. 3A shows the proportion of trials in which the Long-Term robot (i.e., reward-maximizing response) was chosen in each condition of the experiment. A one-way ANOVA on these proportions revealed a significant effect of condition, $F(3, 91) = 10.42$,

$p < .001$. In both the NN and MN conditions, the proportion of Long-Term choices did not significantly differ from chance ($M = .53$, $SD = .17$, $t < 1$ and $M = .42$, $SD = .28$, $t(22) = -1.24$, $p = .23$, respectively). However, in the LN condition, participants chose the Long-Term robot more often than the Short-Term robot, $M = .60$, $SD = .18$, $t(22) = 2.74$, $p = .012$. In contrast, in the HN condition, participants chose the Short-Term option more often, $M = .26$, $SD = .21$, $t(22) = -5.29$, $p < .001$. Planned comparison revealed that the proportion of maximizing responses did not differ between the NN and LN conditions, $t(44) = 1.32$, $p = .19$. However, a significantly larger proportion of maximizing responses was recorded in the LN condition compared to the MN, $t(44) = 2.51$, $p = .016$, and HN conditions, $t(44) = 5.79$, $p < .001$.

