

Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments

Hart E. Posen
Ross School of Business
University of Michigan
701 Tappan Street, R5340
Ann Arbor, MI 48109-1234
Ph. 734-764-1349
e-mail: hposen@umich.edu

Daniel A. Levinthal
Wharton School
University of Pennsylvania
2000 Steinberg-Dietrich Hall
Philadelphia, PA 19104-6370
e-mail: levinthal@wharton.upenn.edu

Forthcoming in *Management Science*

June 17, 2011

Acknowledgments: Hart Posen is grateful for financial support from the 3M Company. We thank seminar participants at Columbia, Duke, London Business School, Michigan, Minnesota, North Carolina, and Western Ontario for helpful comments on a prior version.

Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments

Abstract

A common justification for organizational change is that the circumstances in which the organization finds itself have changed, thereby eroding the value of utilizing existing knowledge. On the surface, the claim that organizations should adapt by generating new knowledge seems obvious and compelling. However, this standard wisdom overlooks the possibility that the reward to generating new knowledge may itself be eroded if change is an on-going property of the environment. This observation in turn suggests that environmental change is not a self-evident call for strategies of greater exploration. Indeed, under some conditions the appropriate response to environmental change is a renewed focus on exploiting existing knowledge and opportunities. We develop a computational model based on the canonical multi-armed bandit formulation of exploration and exploitation. We endeavor to understand the mechanisms by which environmental change acts to make purposeful efforts at organizational adaptation less (or more) valuable.

1. Introduction

A common justification for organizational change is that the circumstances in which the organization finds itself have changed. This idea is broadly reflected in the contemporary literature's focus on, for example, how the pace of environmental change accentuates the need for exploratory search (c.f., March 1991) and the importance and desirability of dynamic capabilities (c.f., Teece, Pisano, and Shuen 1997). Claims about the value of organizational action aimed at adaptation in the face of environmental change have a taken-for-granted quality and, on the surface, seem obvious and compelling. Yet, as we will argue, change in the external environment does not necessarily imply the need for, or benefit of, expanded efforts at organizational adaptation. Indeed, under some conditions the appropriate response to environmental change is a renewed focus on exploiting existing knowledge and opportunities.

In examining the merits of organizational adaptation, our work is closely related to March (1991), who identifies the need to allocate limited resources across both the exploitation of the known and exploration of the novel as a central strategic trade-off. If balancing exploration and exploitation is central to firm performance, as the consensus now suggests (Gupta, Smith, and Shalley 2006), then what factors affect the appropriate balance? To answer this question, March (1991) focuses on the role of environmental change, reflecting a long tradition of interest in the external environment and the ability and willingness of organizations to adapt (c.f., Benner 2009; Hannan and Freeman 1984; Henderson 1993; Lant and Mezias 1992; Sorenson and Stuart 2000; Tushman and Romanelli 1985). March (1991 p.80) concludes that “exogenous environmental change makes adaptation essential,” but he then tempers this conclusion by noting that it also “makes learning from experience difficult.”

The broader strategy literature is generally less circumspect. In recognizing the importance of Schumpeterian dynamics, the literature increasingly focuses attention on dynamic capabilities (Teece et al. 1997), high-velocity markets (Brown and Eisenhardt 1997), and hypercompetition (D'Aveni and Gunther 1994). While in substance, this work is nuanced in its treatment of strategic responses to environmental change, in practice it is often interpreted as suggesting a simple positive relationship

between the extent of environmental change and the merits of organizational adaptation. This idea is diffused in the conventional wisdom of the practitioner literature, and managerial sentiment seems to reflect Peters and Waterman's (1982) call for a "bias for action." While we agree with the Schumpeterian focus of much recent research, our analysis suggests that an exhortation to action in dynamic environments is *not always* appropriate. As such, we endeavor to understand the mechanisms by which environmental change makes purposeful efforts at organizational adaptation less (or more) valuable.

We anchor our formal development on the multi-armed bandit model. This model is the canonical representation of the exploration-exploitation problem across literatures as diverse as economics and statistics (Berry and Fristedt 1985; Gittins 1979; Robbins 1952) and computer science (Holland 1975; Sutton and Barto 1998). In the management literature, March (1991) originally formulated his analysis in terms of a specific form of a genetic algorithm;¹ subsequently, he reformulated his discussion of learning in terms of the bandit model (Denrell and March, 2001; March 1996, 2003, 2010).

Consider the bandit model in the context of a simple organization, where management consists of making choices under uncertainty. In each period, the organization must choose from a set of policy alternatives; it may choose to continue with the policy choice made in the prior period, or it may select a different policy choice. The payoff for choosing a particular alternative reflects a draw from a probability distribution with an unknown, alternative-specific, mean. In pursuing its activity, the organization seeks to maximize the flow of returns over time. The organization makes its choice in each period based on its beliefs about the expected returns to each of the alternatives. If these subjective beliefs were known to be objectively correct, then the organization would choose to exploit the alternative known to have the highest expected return. In the absence of such (perfect) knowledge, the organization employs a strategy that allocates effort between (a) exploiting the alternative currently believed to be superior, and (b)

¹ The model in March (1991) is a specific form of a genetic algorithm in the sense that individuals do not recombine attributes of other individuals directly, but rather they do so indirectly, mediated by the "organizational code."

exploring alternatives that currently seem less promising in hope of identifying a superior alternative to the current preferred alternative.

We examine this bandit model of organizational choice and focus on the implications of environmental turbulence that alters the returns to particular policy alternatives. Through our model, we highlight two main observations. First, the appropriate response to environmental change is not necessarily a strategic shift of effort towards exploration – indeed, a shift towards exploitation is sometimes superior.² Environmental change is a two-sided “blade” with respect to the benefits of organizational adaptation. Scholars have tended to focus on the diminution of the firm’s current knowledge stock and the value of existing practices, when change renders “a prior strategic orientation ... no longer effective” (Tushman and Romanelli 1985 p.180). Yet there is another, less thoroughly considered, side of the blade. It reflects the possibility that recurring, though not necessarily predictable, change erodes the future benefits of additional knowledge gained via exploration, suggesting a more circumspect conclusion as to the merits of organizational adaptation in dynamic environments.

Second, organizations are biased towards action in the face of environmental change, even in the absence of a strategic shift towards exploration. This action bias arises endogenously in the process of experiential learning where choices are based on beliefs about the merits of the alternatives. An organization exploits a particular policy alternative when it is believed to be markedly superior to others. Environmental change naturally undermines existing beliefs. Alternatives that were believed to be rather good are likely to seem worse. Moreover, there is the possibility that less examined alternatives may have improved. This compression of beliefs leads to an action bias in dynamic environments.

This paper proceeds as follows. In the next section, we begin by discussing the theoretical background of our model. Then we describe the multi-armed bandit computational model. In section 3,

² Keller and Rady (1999) developed a related result of a monopolist responding to a demand curve that shifts between two states, showing that monopolist may engage in less experimentation when the frequency of the shift between the two states increases.

we present the results of three simulation experiments. The first experiment examines the baseline properties of the model by focusing on a simple static environment. The second experiment adds turbulence to the baseline experiment. In doing so, we examine the mechanisms that underlay the action bias, and the appropriateness of purposeful efforts at adaptation in the face of change. The third experiment examines the implications of non-neutral turbulence that not only changes the relative merits of alternatives, but also makes the average better or worse. In section 4, we discuss the implications of these results for theory and practice.

2. Model

The multi-armed bandit model takes its name from a slot machine analogy. In the analogy, a gambler sits in front of a slot machine with N arms, each having different underlying payoff probabilities that are unknown to the player. The gambler's objective is to maximize the returns to playing the machine. This model has been the subject of significant study because its underlying structure closely resembles a variety of realist economic situations ranging from research settings, such as R&D (Hardwick 1992), to strategic issues, such as product pricing (Bergemann and Välimäki 1996), and consumer choice (Gans, Knox, and Croson 2007). Common to all of these settings is that information about the returns to an alternative can only be gathered by selecting it. As such, the primary feature of the bandit model is the process of search and learning (Denrell and March 2001; March 2010). As a consequence, as Holland (1975) notes, in the bandit model “the trade-off between gathering more information and exploiting it appears in the simplest terms.”

Formalizing the discussion above, the bandit model reflects a sequential choice problem where, at each point in time, t , an organization must choose between N alternatives. The realized outcome of a particular choice is a draw from an unknown (alternative specific) probability distribution. If the process is Bernoulli, then the choice of alternative i results in an outcome of $\sigma = (1, 0)$ – with $\sigma = 1$ reflecting a positive outcome with probability p_i and $\sigma = 0$ reflecting a negative outcome with probability $1-p_i$. As

such, the state of the environment can be described by the vector of payoffs to the alternatives, $P_t = [p_{1,t}, \dots, p_{N,t}]$. In the typical bandit model formulation, the environment is stable across time such that, for all t , $P_t = P$. We relax this assumption to allow for environmental turbulence.³ Consistent with the definition of turbulence as reflecting both the frequency and unpredictability of environmental change (Baum and Wally 2003; Dess and Beard 1984; Miller 1987; Wholey and Brittain 1989), we model the level of turbulence as the probability of a shock to the payoff values of alternatives.

We assume that an organization is a boundedly rational actor engaged in a process of search and learning (Denrell and March 2001; March 2010). Choice behavior in the bandit model is driven by the relationship between two key organizational constructs: beliefs and strategies. We discuss each in turn below.

In the bandit model, the first key feature of organizational behavior is that an organization is assumed to make its choice (select an alternative) based on its subjective assessment of the merits of the alternatives. In employing beliefs to facilitate choice, the model portrays a boundedly rational organization as possessing a mental model or cognitive map that encapsulates its understanding of the merits of the available set of choices (Barr, Stimpert, and Huff 1992; Gavetti and Levinthal 2000; Nelson 2008; Weick 1969). In particular, an organization's belief about the merits of a given alternative is $q_{i,t}$, indexed by alternative i at time t , such that its set of beliefs across alternatives is characterized by $Q_t = [q_{1,t}, \dots, q_{N,t}]$ with $0 \leq q_{i,t} \leq 1$.

Two dimensions of beliefs are particularly relevant. First, subjective beliefs may (or may not) have objective validity. To the extent that beliefs accurately characterize the state of the world and reflect the true underlying value of the alternatives, then they embody knowledge of the choice opportunities that the organization faces. Thus, we characterize knowledge based on the sum of squared errors in beliefs relative to the true alternative payoffs, as $1 - \sum_1^N (p_i - q_i)^2$. Second, beliefs may embody strong or weak

³ A static environment is typically assumed in the economics literature because of the challenges associated with obtaining closed form solutions. The rare exception here is Whittle's (1988) "restless bandit" model.

opinions. The strength of opinions reflects the extent to which an organization believes that some alternatives are strongly superior to others. We characterize this as the variance across beliefs, $\sum_1^N (\bar{q} - q_i)^2$.

To refine its beliefs and maximize the value of its stream of rewards, an organization engages in learning. In the initial period $t=0$, the organization has prior beliefs, $Q_0=[q_{1,0}, \dots, q_{N,0}]$, about the value of the N reward outcome probabilities $P=[p_1, \dots, p_N]$. In each subsequent period, the organization makes a choice from the set of alternatives. By acting – making a choice of alternative i – a firm receives feedback from the environment in the form of an outcome $\sigma_{i,t}=(1, 0)$ associated with a reward of $r_{i,t}=(r^+, r^-)$. This reward may, for example, take the form of sales and profits (or lack thereof), which augments the firm's stock of assets such that $S_t = \sum_1^t r_t$.

Equally important, from the learning perspective, performance feedback provides new information about action-outcome linkages (the value of alternatives) that serve to update beliefs. In particular, beliefs are updated as the organization chooses one option i and observes the outcome, $\sigma_t=(1, 0)$. We assume that beliefs at any given point in time reflect the average reward over an organization's entire history of trials of a given alternative (March 1996, Sutton and Barto 1998). This simple average updating is a special case of the more general Bush-Mosteller updating methodology. In a given period, the firm updates its beliefs, $q_{i,t}$, of the selected alternative following a Bush-Mosteller fractional adjustment methodology where $q_{i,t+1} = q_{i,t} + \delta(\sigma_t - q_{i,t})$ and δ is the rate of updating. Average updating is the special case of $\delta=1/(k_t+1)$, where k_t is the number of times alternative i has been tried at time t . For alternatives that are not selected in a given period, $q_{i,t+1} = q_{i,t}$.

The second key feature of organizational behavior in the bandit model is that an organization is assumed to possess a search strategy. The organization would like to select, in each period, the alternative with the highest expected value, $\max(p_1, \dots, p_N)$. However, the expected payoffs are unknown to the organization. Thus, conditional on a given set of beliefs about the relative merits of the alternatives, an organization requires a mechanism for choosing among the alternatives. An organization's search strategy

can be conceived of as a rule or routine (Cyert and March 1963; Nelson and Winter 1982) that maps the organization's beliefs about the merits of alternatives to specific choices from the set of alternatives.

Under an appropriate set of assumptions, Gittins and Jones (1974) demonstrate that an optimal strategy exists. The basic property of this rational choice strategy in the bandit model is as follows. Consider a slot machine with two arms, arm 1 has an unknown payoff probability, and arm 2 has a known payoff probability (or a fixed payoff). The rational choice solution involves identifying the optimal stopping time – the number of trials to sample arm 1 before deciding to select one or the other arm in perpetuity. In the context of a multi-arm bandit (with N unknown arms), the solution is not a stopping rule, rather it is a dynamic allocation index (Gittins index) that defines the optimal arm choice in any given period (Gittins 1979).

While the optimal strategy described by the Gittins index is sufficiently computationally intensive such that it is intractable except under a set of restrictive assumptions (Gans et al. 2007), it is worth pointing out its central characteristic. The “index for an arm is equal to the expected immediate reward from the arm, with an upward adjustment reflecting any uncertainty about the prospects of obtaining rewards from the arm, and the possibilities of resolving those uncertainties by selecting that arm” (Gittins and Wang 1992 p.1625). The index suggests a pattern of choice over time that is sensitive to both the mean of beliefs, and the uncertainty surrounding that mean.

Behavioral approaches that seek to identify an organizationally plausible, although not necessarily optimal, strategy date back to Schmalensee (1975) and Radner and Rothschild (1975). While there is a wide variety of plausible strategies, perhaps the simplest and best-known strategy is that of selecting, at each point in time, the alternative with the highest belief, $\max(q_1, \dots, q_N)$, reflecting the highest expected reward (Auer, Cesa-Bianchi, and Fischer 2002; Sutton and Barto 1998). This rule is referred to as greedy because at each point in time it myopically maximizes the expected reward in the next period. The greedy rule always exploits current knowledge, but never explores alternatives that are apparently inferior, but may in fact be superior.

As Brezzi and Lai (2002 p.105) note in their discussion of behavioral approximations to the Gittins index, a good rule “introduces adjustments into the myopic (greedy) rule so that some active experimentation is used to generate information about the unknown parameters.” Thus, a simple variation on the greedy rule is to behave greedily a majority of the time, but with some small probability ε select an alternative at random, independent of beliefs (Sutton and Barto 1998). This rule is typically referred to as ε -greedy. However, the ε -greedy rule has an important undesirable feature. The rule is both overly sensitive to beliefs and completely insensitive to beliefs. When the organization exploits (1- ε percent of the time), it chooses the alternative with the highest belief, even if that differential to the second highest belief is trivially small. In contrast, when it explores (ε percent of the time), it is equally likely to select among all alternatives, independent of beliefs about the merits of a particular alternative. As such, in exploring, the ε -greedy rule is equally likely to pick the worst alternative, as it is to pick the next-best alternative.

A well-defined and examined strategy that addresses the above deficiency is found in the softmax choice rule attributable to Luce (1959), and employed widely (Camerer and Ho 1999; Gans et al. 2007; Sutton and Barto 1998; Vermorel and Mohri 2005; Weber, Shafir, and Blais 2004). This strategy formulation takes the form of random choice based on a Gibbs (Boltzmann) distribution. Here, the probability of selecting alternative i , m_i , is defined as:

$$m_i = \frac{e^{(q_i/(\frac{\tau}{10}))}}{\sum_{i=1}^N e^{(q_i/(\frac{\tau}{10}))}} . \quad (1)$$

where q_i is the belief regarding alternative i , and τ is a parameter reflecting the organization’s strategy.⁴ A strategy, τ , maps the organization’s beliefs $Q=[q_1, \dots, q_N]$ to choice probabilities $M=[m_1, \dots, m_N]$ where $\sum_{i=1}^N m_i = 1$. When $\tau \rightarrow 0$ even small differences in beliefs are weighted heavily. Thus, if one alternative

⁴ Note that the “/10” in Equation 1 acts simply as a scaling factor.

is viewed even marginally more favorably than the others, it will be selected with near certainty, and thus we observe a very greedy strategy (highly exploitative). For example, consider a three-armed bandit with beliefs across arms of $Q=[0.7, 0.8, 0.7]$. A purely exploitive strategy would map this to a set of choice probabilities of $M=[0.0, 1.0, 0.0]$ so that the firm chooses with probability one the alternative with the highest beliefs. In contrast, when $\tau \rightarrow \infty$, differences in beliefs are shrunk to the point that the organization is indifferent between alternatives (even if it believes that one alternative is far superior). This leads to exploration as a random choice amongst all alternatives, where each alternative is chosen with probability $1/N$. Continuing the example, at very high τ , the beliefs $Q=[0.7, 0.8, 0.7]$ would map to the choice probabilities $M=[0.33, 0.33, 0.33]$. In such a setting, the organization would be insensitive to the fact that its current beliefs suggest that alternative two is superior, and with probability 0.66 would choose an alternative other than the one currently believed best. At moderate levels of τ , alternatives that appear to be good (high beliefs) are more likely to be sampled. At intermediate levels of τ , alternatives that appear to be good (high beliefs) are more likely to be selected.⁵

Before proceeding with the analysis, we wish to highlight the relationship between the bandit model and that of the genetic algorithm model in March (1991). We focus on two key features of strategy in these models (a) the mechanism governing organizational choice, and (b) the meaning of exploration. First, March's (1991) genetic algorithm and the bandit model are fundamentally concerned with the trade-off between exploitation of existing knowledge and exploring for new knowledge in the process of organizational learning. Nevertheless, the level of analysis is different. March (1991) generates differences in the balance between exploration and exploitation by altering internal characteristics of the organization. The strategy features of his model are tunable parameters reflecting turnover and learning rate at the level of the individual. In our bandit model, a strategy τ encapsulates these internal organizational dynamics in a single parameter reflecting the sensitivity of organizations to their beliefs.

⁵ Organizations are, by design, risk neutral. They select an alternative solely based on the expected value of an alternative. Although, per March (1996) and Le Mens and Denrell (2011), risk aversion emerges endogenously in a model of experiential learning.

Yet, regardless of the level at which the tunable strategy parameter is implemented, the implications for tuning the explore-exploit trade-off remain consistent.

Second, March (1991) argues that exploration not only suggests a willingness to forego the opportunity to engage in an action that one believes superior, but also to privilege alternatives with less certain beliefs. A search strategy in the bandit model implies a similar understanding of the notion of exploration, because choice must focus on the mean value of an alternative with “an upward adjustment reflecting any uncertainty” (Gittins and Wang 1992 p.1625). It is less obvious that the softmax formulation of strategy in our bandit model, while explicitly only incorporating the first moment of beliefs (q_i), endogenously weights the mean belief by uncertainty. We discuss the characteristic of our bandit model in the analysis of the baseline characteristics of the model in the subsequent section.

3. Analysis

We use the multi-armed bandit model described above as an analytical tool to engage in an analysis of the efficacy of alternative adaptation strategies under turbulence. The structure of the simulation requires initializing two sets of variables associated with: (a) the environment and turbulence rate; and (b) organizational strategies and beliefs.

The opportunity structure of the environment is defined by initializing the payoff probabilities for each alternative. We formulate a ten-arm bandit model such that the vector of alternatives' payoff probabilities is $P_t = [p_{1,t}, \dots, p_{10,t}]$. The choice of a particular alternative i leads to outcome, positive/negative such that $\sigma_{i,t}=(1, 0)$, with probabilities $(p_{i,t}, 1-p_{i,t})$ and an associated reward of $r_{i,t}=(1,-1)$. Each alternative is allocated an initial payoff probability, $p_{i,0}$ that is taken to be a draw from a beta distribution ($\alpha=2$ and $\beta=2$). This produces a distribution of payoff probabilities (across alternatives) that is symmetrical with mean of 0.50 and standard deviation of 0.22. The distribution resembles in shape a

normal distribution, but has a distribution over probability values that it is bounded between zero and one.⁶

Models incorporating turbulence employ a probabilistic shock in a given period, where the probability of a shock, η , ranges from 0.0 to 0.32. When $\eta=0$, the environment is stable, and $P_t=P$. For $\eta>0$, conditional on the occurrence of a shock, the payoff to each of the alternatives is reset with probability 0.5. To reset an alternative, we redraw the payoff, $p_{i,t}$, from the same initial beta distribution ($\alpha=2$ and $\beta=2$). As a consequence, such turbulence is neutral in the sense that the expected value of the set of alternatives remains unchanged after a shock, while the payoff to individual alternatives change.

All simulations are seeded with 25,000 organizations, each facing a unique environment with alternatives drawn from a beta distribution ($\alpha=2$ and $\beta=2$). We set the initial asset stock to zero, and run the simulation for 500 periods. We specify initial beliefs across these alternatives, $Q_0=[q_{1,0}, \dots, q_{10,0}]$, to be homogenous and set the value to 0.5, which is equal to the mean value of the actual distribution of payoffs. The organizations are divided equally among five different strategies. The strategies, τ , reflect different balances between exploration and exploitation. The strategies implemented, which range from pure exploitation to high exploration, are $\tau=(0.0, 0.25, 0.5, 0.75, 1.0)$.⁷

The analysis proceeds over three experiments in which we examine the performance implications of alternative strategies of balancing exploration and exploitation across different environmental contexts: (1) stable world, (2) turbulent world (mean neutral), and (3) turbulent world (non-neutral: mean devaluing or mean enhancing).

⁶ We examined alternative initial arm distributions. In particular, we re-ran the models in this paper with arms drawn from a uniform distribution. The uniform is a special case of the Beta distribution with ($\alpha=1$ and $\beta=1$). The results are qualitatively the same.

⁷ Equation 1 does not permit the use of a value of zero for the strategy. As τ approaches zero, the value of Equation 1 approaches infinity. Therefore, we set as a pure exploitation strategy the computationally feasible value of $\tau=0.02$. This value is sufficiently close to zero that the strategy reflects pure exploitation in the sense that the alternative with the highest expected payoff is chosen in all runs. As such, the figures and discussion in the remainder of the paper simply refer to this as $\tau=0$.

Experiment 1: Basic Properties of the Model - Adaptation Strategies in a Stable World

We begin our analysis by examining the bandit model in a stable environment. Our objective is to examine the basic properties of the model and to demonstrate that it generates results that are consistent with the prior literature, in terms of both the March (1991) model and also the broader bandit literature.

In Figure 1, we examine the implications of altering adaptation strategy, τ , across three measures: (a) exploratory choice, (b) performance, and (c) knowledge generation. Each is measured at the end of the simulation. First, the dashed line (right axis) represents the extent to which the choices taken by organizations are exploratory, measured as the probability of choosing a different alternative in period $t+1$ than was chosen in period t . As τ increases, the level of exploration increases. At the most exploratory level of search strategy in our model, $\tau=1$, organizations select a new alternative roughly every second period.⁸

----- Insert Figure 1 about here -----

Second, the solid line (left axis) represents performance measured as the accumulated asset stock. The graph highlights that the best strategy, which displays a moderate level of exploration ($\tau=0.5$), accumulates on average 280 units, which reflects an average frequency of positive payoffs across periods of 78 percent. Third, the dotted line (right axis) reflects the accuracy of knowledge embodied in firms' subjective beliefs about the value of alternatives at the end of the final period. We measure knowledge as 1 minus the sum of squared error of beliefs relative to the true arm payoffs.⁹ The graph highlights that the level of knowledge obtained is increasing with the extent of exploration, from 0.55 at $\tau=0$ to 0.72 at $\tau=1$.

⁸ We also examined alternative metrics of exploration. For example, the probability of not choosing the arm that is believed to offer the highest payoff. The qualitative results on exploration behavior do not change with this, or other, alternative measures.

⁹ This measure of knowledge can be interpreted as reflecting an explicit comparison to the level of knowledge that would be obtained by a firm with beliefs chosen randomly. Knowledge is zero if beliefs are set randomly from a Beta (2,2) distribution matching reality, and knowledge is one if beliefs perfectly reflect reality such that SSE=0. Thus, this measure of knowledge has the attractive property that knowledge is greater than 0 when the organization has knowledge above that of an organization with random beliefs. Organizations pursuing a search strategy that embodies a moderate degree of adaptiveness, $\tau=0.5$, accumulate, on average, knowledge over the 500 periods of the simulation of 0.59.

Taken together, the inverse u-shaped performance curve across strategies, and the near linear increase in knowledge demonstrates the trade-off between generating new knowledge and leveraging existing knowledge. At low levels of exploration, little knowledge is attained and there is little knowledge to exploit. In contrast, at high levels of exploration, significant knowledge is attained, but it is never exploited to enhance performance. Thus, in this static environment, the results reflect March's (1991 p.71) conclusion that organizations that "engage in exploration to the exclusion of exploitation are likely to find that they suffer the costs of experimentation without gaining many of its benefits," while organizations that "engage in exploitation to the exclusion of exploration are likely to find themselves trapped in suboptimal stable equilibrium."

While our model replicates the main results of the March (1991) model, we demonstrate that that it also replicates key results in the bandit literature. We focus on two features of the model: (a) identifying the extent to which exploration in our model embodies choice based on uncertainty in beliefs, and (b) docking our results to a bandit model with an ϵ -greedy strategy.

Gittins and Wang (1992 p.1625) note that in selecting an alternative, it is appropriate to consider the "expected immediate reward from the arm, with an upward adjustment reflecting any uncertainty about the prospects of obtaining rewards from the arm, and the possibilities of resolving those uncertainties by selecting that arm." This intuition, which arises from the Gittins index, is consistent with the notion of exploration in March (1991) and the broader management literature. Nevertheless, while the softmax choice strategy we employ considers only the mean of beliefs, neglecting direct consideration of uncertainty, we demonstrate that it endogenously weights the level of uncertainty in its choice behavior.

In the bandit literature, strategy is defined as greedy if, in each period, it chooses the alternative with the highest mean belief. Analogously, we construct a measure of greediness based on experience with an alternative (as a proxy for uncertainty), rather than on the first moment of beliefs. Greedy-on-experience measures the probability that an organization will choose an alternative with which it has the

most experience. The results of these measures, at period 500, are presented in Figure 2. As exploration strategy, τ , increases, organizations become less greedy, on the basis both of mean and of uncertainty in beliefs. As such, organizations are more likely to explore alternatives that are presently believed to be inferior and for which the beliefs are relatively imprecise.

----- Insert Figure 2 about here -----

The intuition behind this result rests on the observation that alternatives that appear to be good are more likely to be sampled in the process of experiential learning (March, 1996; Denrell and March, 2001). Moreover, the more an alternative is sampled, the more certain the estimate of its mean value. As a result, there is a positive correlation between the mean belief regarding the payoff to a given alternative and the level of experience with that alternative. When an organization chooses a low mean belief alternative, it is also likely to select an alternative that is known with less certainty. Exploration strategies (high τ) are not only more likely to forgo alternatives with high mean belief expectations, they are likely to forgo more certain alternatives in favor of less certain alternatives.

For the purposes of model docking, we rerun our simulation employing a simple ϵ -greedy strategy, rather than our softmax strategy. This enables comparisons across the two choice mechanisms. We expect, and observe, that as $\tau \rightarrow 0$, the performance results converge to the greedy strategy with $\epsilon=0$. Moreover, as expected from the above discussion of the characteristics of the softmax choice rule, it performs somewhat better than the ϵ -greedy choice rule. Additionally, Sutton and Barto (1998 p.41) examine an $\epsilon=0.1$ greedy strategy on a 10-arm Gaussian (rather than Bernoulli) bandit model, and measure the percent optimal action. This measure reflects the probability that the alternative chosen in a given period is the alternative with the highest payoff probability. We examine this metric at period 500. In our implementation, both the $\tau=0$ and $\epsilon=0$ strategies exhibit the same optimal action result of 39 percent. With $\epsilon=0.1$, our model exhibits 63 percent optimal action, consistent with the Sutton and Barto (1998) result. Thus, the softmax strategy of our bandit model docks to the ϵ -greedy variant of the bandit, and our ϵ -greedy implementation docks to that of Sutton and Barto (1998).

In sum, in a static environment, the properties of the bandit model are consistent with March's (1991) claim that superior performance arises from a strategy that balances exploration and exploitation.

Experiment 2: Adaptation in a Changing World

In the second experiment, we add turbulence to the model, at levels ranging from very infrequent to near continuous (0 through 32 percent).¹⁰ We examine how and why turbulence alters the optimal level of adaptiveness embodied in a search strategy.

----- Insert Figure 3 about here -----

The main results of the turbulence model are depicted in Figure 3. This figure plots the optimal search strategy calculated by fitting a third order polynomial to the results (that is, Figure 1 replicated across levels of turbulence) and extracting the strategy, τ , associated with that peak performance (maximum stock of accumulated assets). The zero turbulence result reflects the conclusions from experiment 1: peak performance occurs at a search strategy of $\tau=0.56$, reflecting a moderate degree of adaptiveness. As the probabilistic frequency of turbulence increases, the optimal level of exploration takes an inverse u-shaped form. The curve is highly skewed in the sense that over most of the range of turbulence examined, the optimal degree of adaptiveness embodied in a search strategy is decreasing with increasing frequency of turbulence.

In addition to altering the optimal level of exploration strategy, turbulence may also alter the extent of the cost associated with failing to match the optimal exploration strategy. To examine this issue, we calculate the flatness of the performance vs. τ curves by plotting the second derivative of performance at the optimal τ value (see Figure 3). If the curve is sufficiently flat, because little or no knowledge can be accumulated via experiential learning, then the distinction between exploration and exploitation becomes

¹⁰ At higher levels of turbulence, performance becomes invariant to exploration strategy.

less relevant. Over the interesting range of turbulence (0 through 8 percent), the curve is significantly concave. At very low levels of turbulence (below 0.5 percent), the second derivative of performance at the optimal τ value is increasingly negative, implying the performance versus τ curve is increasingly concave at the optimal τ . This suggests that increasing turbulence enhances the rewards to choosing the optimal τ . As turbulence increases further (at levels above 0.5 percent), the second derivative starts to increase (becomes less negative), implying that the performance versus τ curve at the optimal τ is becoming flatter. This implies that identifying the optimal strategy becomes less important for performance at higher levels of turbulence.¹¹

Because the discussion of turbulence and exploration in March (1991) is so prominent in the literature, it is worthwhile to directly compare these results to his. In contrast to March (1991), the results in Figure 3 suggest a contingency in the relationship between turbulence and the level of exploration. However, March (1991) examines only a single, relatively modest level of turbulence (2 percent). His results, like our own, show that organizational performance is improved by increasing the level of exploration as turbulence increases from zero through two percent. Like March (1991), we find that under conditions of very infrequent environmental change, a strategy stressing adaptation via exploration and new knowledge accretion is optimal, but as change becomes more frequent, the optimal response is, in fact, one of greater inertia.

We engage in additional analysis to understand the mechanisms driving this response to turbulence. First, we focus on the tension between the obsolescence of existing knowledge and the accretion of new knowledge. Second, we focus on the impact of turbulence on the organizations strength of opinions surrounding its beliefs about the relative merits of alternatives. We examine each mechanism in turn below.

¹¹ Even a flat curve performance versus τ curve does not imply that an organization should be indifferent between the levels of exploration strategy. In our model, exploration and exploitation have no explicit cost (only opportunity costs). If, in contrast to our model, exploration has a direct cost that is higher than exploitation then erring on the side of too little exploration is superior to erring on the side of too much exploration under high levels of turbulence.

Our working hypothesis, addressing the tension between knowledge obsolescence and accretion, is that the characterization of preferred search strategies in Figure 3 stems from the fact that turbulence erodes knowledge asymmetrically. That is, turbulence acts to erode performance via two knowledge mechanisms. First, it reduces the future value of existing knowledge because it changes the underlying state of the world. Second, turbulence reduces the value of efforts to generate new knowledge because the lifespan of returns to new knowledge is reduced in a world in which change is more frequent. We examine the possibility of an asymmetry in the effect of turbulence on these two mechanisms in Figures 4 and 5.

----- Insert Figures 4 and 5 about here -----

We engage in a differences-in-differences analysis of knowledge levels. Holding beliefs constant at period 400 levels, we examine the resulting level of knowledge erosion over the subsequent 100 periods. We then compare the period 400 knowledge results to the period 500 knowledge results (assuming constant beliefs), and the period 500 results assuming learning at a strategy of moderate adaptation, $\tau=0.5$. The dashed line in Figure 4 reflects the level of knowledge at period 400, across levels of turbulence. In the absence of turbulence, the knowledge level is 0.59 (as per the analysis in experiment 1). By the time turbulence reaches 8 percent, knowledge is reduced to 0.44. The solid line presents the results of using period 400 beliefs through period 500 (that is, period 400 beliefs are used, but not updated). As such, the difference between the solid and dashed lines reflects the erosion of existing knowledge over the 100-period time span. Finally, the dotted line reflects knowledge levels at period 500 if learning (updating of beliefs) continued through to period 500. The difference between the solid and dotted lines reflects the improved knowledge associated with efforts at new knowledge accretion over a 100-period time span.

In Figure 5, the solid and dashed lines reflect, respectively: (a) existing knowledge which is eroded by turbulence – the difference between the dashed and solid lines in Figure 4; and (b) new

knowledge accretion - the difference between the solid and dotted lines in Figure 4. The dotted line in Figure 5 depicts the net knowledge change.

Over the range from no turbulence through a frequency of one percent, existing knowledge is eroded rapidly. The more rapidly existing knowledge is eroded, other things being equal, the greater the impetus to explore. The alternative to employing additional efforts at exploration is using existing knowledge. As such, **the erosion of existing knowledge defines the opportunity cost of exploring** - which is decreasing rapidly over this range. At the same time, additional exploration is generating new knowledge at an increasing rate. With decreasing opportunity cost and increasing generation of new knowledge, there is a substantial net knowledge benefit to exploring (positive slope of the dotted line) and, as such, the appropriate response is one of increased exploration at low levels of turbulence.

Over the range of turbulence from one through two percent, the existing knowledge erosion level is flattening and, indeed, slightly decreasing. Relative to this low turbulence range, this reflects a substantial decrease in the *rate* at which opportunity cost is decreasing. At the same time, the rate of new knowledge accretion is increasing. However, the increase in the rate of new knowledge accretion is substantially less than the decrease in the rate of existing knowledge erosion (and increase in opportunity cost). Thus, the appropriate response, as reflected by the negative slope of the net knowledge change curve, is one of decreased exploration over the one to two percent turbulence levels.

Over the range of turbulence in excess of two percent, existing knowledge continues to erode less strongly (opportunity cost is increasing), and new knowledge accretion levels flatten, then decline. As a result, **the appropriate response at higher levels of turbulence is one of further decreases in exploration**. At the very highest levels of turbulence, there is little knowledge available to erode, and very limited potential for new knowledge accretion.

In sum, the appropriate adaptation strategy depends on the effect of turbulence on both the payoff to utilizing current knowledge and the payoff to generating new knowledge. The value of current knowledge reflects the possibilities of continuing to use existing knowledge into the future and, thereby,

drives the opportunity costs associated with engaging in additional exploration. However, turbulence does not simply alter the value of existing knowledge. It also alters the payoff to efforts at gathering new knowledge. The effect of increasing turbulence depends crucially on the relative consequences of turbulence along these two dimensions.

The second mechanism underlying the results in this experiment is that turbulence endogenously gives rise to an action bias, even when the exploration strategy itself is fixed. In our discussion of organizations' beliefs, we have focused on the extent these subjective beliefs embody knowledge. In addition, there is another important dimension of beliefs – the strength of opinions. The strength of opinions reflects how strongly a firm believes one alternative is superior to other alternatives. We measure this as the variance across an organization's beliefs. At the start of the simulation, firms have homogeneous beliefs about the alternatives; they simply believe all alternatives have a 50 percent payoff probability and, thus, the variance in beliefs across alternatives is zero. However, organizations quickly start to form divergent beliefs about these alternatives.

----- Insert Figure 6 about here -----

In Figure 6, we plot, at period 500, the accuracy of knowledge, strength of opinions, and extent of exploration across the range of turbulence for the $\tau=0.5$ strategy. The dotted line highlights the devaluation of knowledge with increasing turbulence. The dashed line plots the strength of opinions (variance across beliefs). The variance in beliefs across alternatives is reduced by nearly one-third as we move from a stable world to one in which turbulence is set at eight percent. It appears turbulence acts not only to devalue knowledge, but also to homogenize beliefs. 

The reduction in the strength of opinions in the face of turbulence has important consequences for the nature of organizational choice. As the strength of opinions is reduced, organizations endogenously increase the extent to which their choices are exploratory – even holding search strategy constant. The mechanism driving this action bias rests on the fact that in the experiential learning process, in making choices, organizations are sensitive to their beliefs about the merits of the alternatives available to them.

Organizations exploit a given policy choice when they believe it is markedly superior to others. Environmental change naturally undermines existing beliefs; alternatives that were believed to be rather good are likely to seem worse. Moreover, there is the possibility that less examined alternatives may have gotten better. As the variance across alternatives in beliefs is reduced by turbulence, organizations have fewer bases on which to differentiate among alternatives. As a consequence of this compression in beliefs, organizations exhibit an action bias in dynamic environments.

In sum, turbulence acts to alter the optimal level of adaptation by altering the returns to utilizing existing knowledge and generating new knowledge. Yet, as we demonstrate, turbulence does not necessarily give rise to a requisite shift towards a more exploratory strategy. Indeed, as the level of turbulence increases, the appropriate balance between exploitation and exploration shifts in the direction of exploitation. In addition, we noted a second mechanism at work – turbulence tends to endogenously generate exploratory behavior even in the absence of an overt change in adaptation strategy. In the next experiment, we examine further implications of this action bias.

Experiment 3: Non-Neutral Turbulence

Turbulence alters the appropriate adaptation strategy because it alters organizational knowledge. In the implementation of environmental change in experiment 2, turbulence reflects a shock to the payoffs to the alternatives, but the shock is neutral. That is, the pre-shock and post-shock distribution of payoffs across alternatives is drawn from the same beta distribution and, as such, the expected value of the set of alternatives is held constant at 0.5. In this experiment, we examine turbulence that is non-neutral in that it makes the world more (or less) munificent (increases or decreases the expected value of the set of alternatives). We hold turbulence at the level of a five percent probability of a shock and examine

turbulence that changes the mean payoff values (cumulatively) over the range from ± 20 percent of the neutral turbulence level.¹²

To illustrate munificence-altering turbulence, consider the emergence of (sufficiently fast) flash based computer memory and the consequent changes in the digital data storage industry. This change reduced the payoff to producing magnetic spinning disks. It also enhanced the returns to a wide variety of applications that were, in the past, deemed to be inferior. These newly enhanced choices include not only solid-state disk drives as a direct replacement for spinning hard drives, but also alternative applications of long-term storage outside the realm of the traditional PC (e.g. music player, smart phone, auto). If this turbulence had been neutral, then to the extent that the returns to selling spinning hard-drives had been reduced, the expected returns to the alternatives would have been similarly enhanced. In this sense, neutral turbulence shuffles the deck but does not change the expected reward. Yet the turbulence in the long-term digital storage industry was non-neutral. The decrease in the expected returns to engaging in designing and selling spinning hard-drives was more than offset by the enhanced returns to the many new uses for solid state drives. As such, this turbulence was munificence-enhancing because it not only reshuffled the deck, it also increased the expected return to the set of alternatives available to organizations in the industry.

Figure 7 displays the optimal search strategy across the range of the munificence of turbulence, holding the level of turbulence constant. Note that the zero percent point reflects the neutral turbulence result from experiment 2. The main result is that if turbulence increases the level of munificence, the optimal level of adaptiveness embodied in a search strategy increases. Because we hold the level of turbulence constant, the rates of erosion of existing knowledge and accretion of new knowledge is held fixed. As such, a different mechanism must be driving this result.

¹² We do so by shifting the parameters of the beta distribution. These shocks are cumulative in that a subsequent shock resets the payoff probability of an alternative from a progressively worse (better) distribution.

----- Insert Figure 7 about here -----

To identify the mechanism, we plot in Figure 8 the strength of organizations' opinions (dashed line) and the level of exploration (solid line) across the levels of the munificence of turbulence. With the search strategy held constant at $\tau=0.5$, increasing munificence of turbulence increases the strength of opinions. This occurs because the payoff to all alternatives is, on average, increasing with munificence enhancing turbulence. However, in the process of experiential learning, organizations are more likely to select an alternative that is known to be good (and as such, the set of alternatives that are considered tends to be decreasing over time). This raises the possibility that they may fail to realize, on average, all alternatives have improved. Thus, in the experiential learning process, organizations may interpret the enhanced performance results as information about the quality of the alternative that they currently favor.

----- Insert Figure 8 about here -----

This increase in the strength of opinions under munificence-enhancing turbulence leads to a decrease in the extent of exploration (per the discussion in experiment 2). But as turbulence, neutral or munificence-enhancing, continues to devalue the returns to existing knowledge and to new knowledge accretion in the same manner, the appropriate level of τ has not changed. Thus, this endogenous decrease in exploration requires a countervailing change in adaptation strategy.

In the face of munificence-enhancing turbulence, organizations are subject to an "inaction bias." This bias is not driven by any inherent behavioral bias, but rather by the nature of belief formation in the experiential learning process.¹³ An increased effort at exploration (a more exploratory strategy) is needed to overcome the endogenous reduction in exploration that is associated with the mechanism described above. Indeed, the opposite is true for munificence-reducing turbulence, which tends to exaggerate the bias towards action.

¹³ Nor is it driven by performance shortfalls relative to aspiration levels, which may also lead endogenously to exploration in the face of turbulence that erodes performance (Lant and Mezias 1992).

In sum, this experiment serves to highlight the observation that turbulence devalues knowledge embodied in beliefs. This affects not only the return to an explicit strategic shift towards exploration, as discussed in experiment 2, but also the nature of organizational action itself. When organizations are engaged in learning, and they are sensitive in their choices about the beliefs regarding the relative merits of alternatives, turbulence can engender organizational action even in the absence of a strategic shift. Indeed, this experiment suggests a rather counter-intuitive conclusion. The appropriate organizational response to turbulence is particularly sensitive to the nature of turbulence – not only the severity, but also the extent to which turbulence is enhancing or degrading the state of the world.

Sensitivity Analysis

We examine the sensitivity of our results to (a) alternative specifications of the key parameters of the bandit model, (b) episodic turbulence, and (c) selection. In each case, the results are robust. First, two key elements of the bandit model have been altered. We examine a bandit model in which belief updating is implemented following a Bayesian methodology, rather than the simple average updating methodology employed in the discussion above. We also examine a model where the alternatives are drawn from a uniform distribution, rather than the near-normal Beta ($\alpha=2$, $\beta=2$) distribution.¹⁴ Second, we examine a model in which turbulence is episodic. In particular, we examine turbulence that occurs in the middle third of the 500-period simulation, across a wide range of severity. We also examine turbulence that occurs in a 50-period window, where the window occurs earlier or later in the simulation. Third, we examine a model in which there is selection. We examine non-competitive selection in which an organization exits when its asset stock falls below zero. We also examine competitive selection, where an organization's probability of failing is a function of its performance relative to other organizations. In each of these additional experiments, the results are robust.

¹⁴ Note that the uniform distribution is a special case of the Beta distribution parameterized as $\alpha=1$ and $\beta=1$.

4. Discussion

We formally consider the effect of changing environmental conditions on firm search strategy and, in particular, the extent to which an organization should pursue a strategy of adaptation under conditions of environmental change. On the surface, the claim that organizations should respond to changing circumstances by generating new knowledge seems obvious and compelling. This logic is reflected in a broader exhortation to action in both the academic and popular literatures – and its validity deserves examination. Is the exhortation to action in a dynamic environment always appropriate? Our modeling efforts focus attention on the implications of turbulence for organizational beliefs about the relative merits of alternatives, and suggest a more circumspect conclusion about the value of organizational adaptation.

Our model points to an inverse u-shaped relationship between the optimal degree of adaptiveness embodied in a search strategy and the frequency of turbulence. This result stems from two opposing mechanisms. On one hand, intuition suggests environmental change devalues existing knowledge. The failure to respond to change in the environment implies decreased performance. This mechanism is well known in the existing theoretical literature and underlies much of the wisdom regarding the appropriate response to change. On the other hand, less well developed and understood in the literature is a mechanism that acts in the opposite direction. Environmental change erodes the rewards to exploratory efforts at accumulating new knowledge. Our analysis suggests that the delicate balance between these two mechanisms, and the possibility of their asymmetrical response to turbulence, determines the appropriate organizational response to environmental change.

Environmental change has a second, arguably more subtle, implication. We conceive of strategies as reflecting managerial and organizational attempts to understand the world and act appropriately. This view is reflected in the bandit model where a strategy is a function that maps knowledge about the world onto particular choices. Thus, a search strategy does not determine a specific action, but rather, action

conditional on beliefs about the state of the world. For a given search strategy, belief structures that are more highly skewed, that place greater weight on some alternatives and lesser weight on others, will generate choices that are more exploitive. By implication, factors that alter organizational beliefs may lead to changes in choice behavior even if the search strategy is fixed and unchanging.

When an organization is involved in a process of experiential learning, environmental change will tend not only to erode the extent of knowledge embodied in beliefs, but also lead to more diffuse, less distinct beliefs about the relative merits of alternatives. Once-attractive alternatives are likely to become less appealing and, as a result, once-positive beliefs about these initiatives will tend to be driven lower with repeated experience. By the same token, should the organization happen to explore alternatives it previously viewed as less attractive, it may find they have improved in the intervening periods. This homogenization of beliefs that arises in the experiential learning process under conditions of turbulence, leads to an action bias.

Consider the case of adaptation to extreme turbulence in the auto industry. In the post-2007 period, oil prices rose markedly and the US housing market collapsed. This led to rapid changes in consumer preferences; they shifted from large, heavy all-wheel drive vehicles to smaller, lighter vehicles with better gas mileage.

The taken-for-granted strategy during this period of turbulence was that of significant action aimed at adaptation. Existing knowledge about designing and building large, heavy, and powerful vehicles was devalued. The response entailed altering the balance between exploitation and exploration by shifting resources towards the development of less well understood technologies (e.g., electric vehicles or light weight carbon fiber bodies) which were expected to be a better fit with the new environment.

This strategic response is reflected in GM's effort to bring the all-electric Volt to market in the post 2007 period. In order to allocate resources to the Volt, GM reduced efforts to exploit existing

knowledge (cutting, at the margin, investment in mid-size fuel efficient vehicles).¹⁵ The process of developing this highly exploratory technology was very slow, not because of organizational impediments to change, but rather because of the limited rate of new knowledge accretion that has been an ongoing feature of this technology domain (batteries in particular). By the time the Volt reached the market in 2010, oil prices had dropped to the point where the vehicle was no longer economically viable, and an ongoing recession made consumers highly price sensitive (as assessed by Consumer Reports, February 2011).¹⁶ Thus, ongoing turbulence had devalued not only GM's return to deploying existing knowledge and technologies (embodied in large SUVs), but also the near-term returns to their exploratory investment in an electric vehicle.¹⁷

While the organizational impetus to act was strong, because this turbulence was seen as munificence-reducing, the returns to exploring were devalued by ongoing turbulence. Yet, there was an alternative course of action in the face of this environmental change. Ford, for example, chose to reduce its efforts at exploration, allocating a larger share of resources to exploitation. They implemented a strategy that focused more effort on refining existing technologies and designs, for example, deploying gas-electric hybrid technology in key models, and modifying the highly successful European Fiesta for the US market. As CNN reported in January 2009: "In the near term, the carmakers' (Ford's) so-called EcoBoost engines - turbocharged engines with highly sophisticated fuel injection systems - will provide the greater fuel economy Americans want as gas prices rise gradually."¹⁸ Ford did not forego exploration in its entirety, continuing an investment program in electric vehicle technologies, but whereas GM

¹⁵ www.businessweek.com/print/magazine/content/08_21/b4085036665789.htm

¹⁶ Consumer Reports, February 2011, as reported on <http://green.autoblog.com/2011/03/01/consumer-reports-chevy-volt-is-going-to-be-a-tough-sell-to-the/print/>

¹⁷ While subsequent increases in fuel prices may render GM's investments worthwhile, the key point is that the possibility of future turbulence reduces the expected returns from exploratory investment.

¹⁸ http://money.cnn.com/2009/01/12/autos/ford_electric_plans/index.htm

massively shifted resources towards exploration in the face of extreme turbulence, Ford shifted resources, at the margin, toward exploitation.

Our theory suggests that the appropriateness of such a strategy is a function of the extent to which existing knowledge is devalued by turbulence, as well as the extent to which ongoing turbulence may devalue investments in exploration aimed at the generation of new knowledge. The challenge for organizations is the ex ante assessment of the magnitudes of these two competing forces. Organizations should be aware of the bias to act, and view the decision to shift the exploit or explore balance towards exploration as a strategic choice rather than a strategic imperative.

An interesting implication of our theory – there are conditions under which the appropriate response to environmental change is to heighten the degree of inertia – is broadly consistent with arguments within the population ecology literature that question the plasticity of organizations. But the mechanism in our theory is quite different. Population ecology holds that organizations are subject to inertial pressures that may arise from embedded organizational structures, routines, and processes that inhibit adaptation (Hannan and Freeman 1984). As a result, inertia may not only be a salient property of organizations, but that it may be performance-enhancing in the face of environmental turbulence.

While population ecology questions the feasibility of organizational adaptation, our theory questions the desirability of adaptation, and points to the constraints on beliefs and their revision given some fixed set of experiences. This distinction between structural inertia, on one hand, and cognitive limitations on the rate of learning, on the other, is important because the two problems necessitate different solutions. If the limiting factor is structure, then the managerial implications are related to designing a more flexible organization, one amenable to modification commensurate with the new state of the world (Davis, Eisenhardt, and Bingham 2009). For example, Brown and Eisenhardt (1997 p.1) argue for a response that “blends limited structure around responsibilities and priorities with extensive communication and design freedom to create improvisation within current projects.” If, in contrast, the limiting factor is a bound on the rate of learning, then the managerial implications are related to actions

that enable organizations to more rapidly acquire and effectively process new information.¹⁹ For example, this may take the form of building enhanced knowledge-management capabilities through additional efforts at environmental scanning. In technology industries, this may include corporate venture capital investments that provide enhanced access to information on emerging technologies (Benson and Ziedonis 2009). Alternatively, this may take the form of purposeful organizational forgetting, selectively discarding old knowledge in order to foster more rapid accretion of new knowledge (Hatch and Dyer 2004; Hedberg, Nystrom, and Starbuck 1976). McNamara and Baden-Fuller (1999), in a case study of publicly traded biotechnology firms, find that newer firms may be more adaptable not only because they have more flexible organizational structures, but also because they learn more rapidly since prior (inappropriate) knowledge does not hinder the process of new knowledge accretion.



Our multi-armed bandit model opens up the possibility of a number of potentially interesting extensions. Two extensions are of particular interest. First, one might model competitive interactions across organizations. In the model presented here, organizations are effectively competing independently against a common environment and subsequently comparing their performance. More realistically, organizations compete directly with one another. Two types of competitive interactions are promising additions to our model. Competition may take the form of competitive selection in which an organization's probability of failing is a function of its performance relative to other organizations. Alternatively, competition may take the form of competitive performance feedback such that instead of receiving direct feedback on the success or failure of a trial, the organization receives feedback that is a function of its own trial success or failure, but also the extent to which others in the population were successful in a given period. This would make feedback about the true value of an alternative noisy, and in turn alter both the nature of beliefs and subsequent choices.

¹⁹ Of course, there may be some relationship between interventions that address the forces of structural inertia and those that facilitate learning.

Second, one might model organizational awareness of environmental change. Organizations in our model are not directly aware of the existence of environmental change, and form no inferences about it. Yet, the usefulness of our discussion is predicated on a limited notion of awareness in the sense that an organization must be sufficiently aware so that it can appropriately tune its adaptation strategy. One might consider a more complex model in which there is second-order learning about the existence, magnitude, and munificence of turbulence. Clearly organizations have beliefs regarding such factors, although there does not seem to be strong evidence that organizations reliably forecast such matters.

Finally, additional extensions include (a) heterogeneity across organizations in the capability to learn; (b) non-fixed search strategies that respond to either the stock of accumulated resources, the nature of existing beliefs, or the level of performance aspirations; (c) interdependence across choices; and (c) endogenous environmental change that derives from interdependence across organizations.

These possible extensions provide important opportunities to build on the basic insights provided here, but would not negate the basic tensions characterized in the current work. Our objective is to first and foremost understand how knowledge is eroded by turbulence, and the consequences of this erosion for the appropriate level of adaptiveness. Adaptive search is multifaceted. It is not only a strategy, but also a set of enacted choices. Strategies map beliefs about the relative merits of alternatives to enacted choices. Examining the interplay between these factors greatly enhances our understanding of the dynamics of exploration and exploitation, and the rewards to organizational adaptation. Further, the manner in which these factors play out, not surprisingly, depends on the context in which they operate, in particular the degree environments are changing. However, the nature of these dependencies is surprising. Environmental change is not a self-evident call for strategies of greater exploration. The question of the appropriate organizational response to environmental change is a central challenge for organizations and unpacking and understanding this tension remains a fertile and important line of inquiry for organizational theorists and strategy scholars.

Bibliography

- Auer, P., Cesa-Bianchi, N., and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* **47**(2-3).
- Barr, P., Stimpert, J., and Huff, A. 1992. Cognitive Change, Strategic Action, and Organizational Renewal. *Strategic Management Journal* **13**(Special Issue) 15-36.
- Baum, J., and Wally, S. 2003. Strategic Decision Speed and Firm Performance. *Strategic Management Journal* **24**(11) 1107.
- Benner, M.J. 2009. Dynamic or Static Capabilities? Process Management Practices and Response to Technological Change *Journal of Product Innovation Management* **26**(5) 473-486.
- Benson, D., and Ziedonis, R. 2009. Corporate Venture Capital as a Window on New Technologies: Implications for the Performance of Corporate Investors When Acquiring Startups. *Organization Science* **20**(2) 329-351.
- Bergemann, D., and Välimäki, J. 1996. Learning and Strategic Pricing. *Econometrica* **64**(5) 1125-1149.
- Berry, D. A., and Fristedt, B. 1985. *Bandit problems: sequential allocation of experiments*. Chapman and Hall, London.
- Brezzi, M., and Lai, T.L. 2002. Optimal learning and experimentation in bandit problems. *Journal of Economic Dynamics and Control* **27**(1) 87-108.
- Brown, S., and Eisenhardt, K. 1997. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Administrative Science Quarterly* **42**(1) 1-34.
- Camerer, C., and Ho, T.H. 1999. Experience-weighted Attraction Learning in Normal Form Games. *Econometrica* **67**(4) 827-874.
- Cyert, R., and March, J. 1963. *A behavioral theory of the firm*. Prentice-Hall, New Jersey.
- Davis, J., Eisenhardt, K., and Bingham, C. 2009. Optimal Structure, Market Dynamism, and the Strategy of Simple Rules. *Administrative Science Quarterly* **54**(3) 413-452.
- Denrell, J., and March, J. 2001. Adaptation as information restriction: The hot stove effect. *Organization Science* **12**(5) 523-538.
- Dess, G., and Beard, D. 1984. Dimensions of Organizational Task Environments. *Administrative Science Quarterly* **29**(1) 52-73.
- Gans, N., Knox, G., and Croson, R. 2007. Simple Models of Discrete Choice and Their Performance in Bandit Experiments. *Manufacturing Service Operations Management* **9**(4) 383-408.
- Gavetti, G., and Levinthal, D.A. 2000. Looking Forward and Looking Backward: Cognitive and Experiential Search. *Administrative Science Quarterly* **45**(1) 113-137.

- Gittins, J.C. 1979. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society. Series B (Methodological)* **41**(2) 148-177.
- Gittins, J. C., and Jones, D. M. 1974. *A dynamic allocation index for the sequential design of experiments*. In J. Gani (Ed.), *Progress in statistics*: 241-266. North-Holland, Amsterdam.
- Gittins, J.C., and Wang, Y. 1992. The Annals of Statistics. *The Annals of Statistics* 20(3) 1625-1636.
- Gupta, A., Smith, K., and Shalley, C. 2006. The Interplay Between Exploration and Exploitation. *Academy of Management Journal* **49**(4) 693-706.
- Hannan, M., and Freeman, J. 1984. Structural Inertia and Organizational Change. *American Sociological Review* **49**(2) 149-164.
- Hardwick, J.P., and Stout, Q.F. 1992. Bandit strategies for ethical sequential allocation. *Computing Science and Statistics* **23** 421-424.
- Hatch, N.W., and Dyer, J.H. 2004. Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal* **25** 1155-1178.
- Hedberg, B., Nystrom, C., and Starbuck, W. 1976. Camping on Seesaws - Prescriptions for a Self-Designing Organization. *Administrative Science Quarterly* **21**(1) 41-65.
- Henderson, R. 1993. Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *Rand Journal of Economics* **24**(2) 248-270.
- Holland, J. 1975. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control & artificial intelligence*. University of Michigan Press, Ann Arbor.
- Keller, G. and S. Rady. (1999). Optimal experimentation in a changing environment. *Review of Economic Studies*, 66: 475-507.
- Lant, T., and Mezias, S. 1992. An Organizational Learning Model of Convergence and Reorientation. *Organization Science* **3**(1) 47-72.
- Le Mens, G., and Denrell, J. 2011. Rational Learning and Information Sampling: On the Naivety Assumption in Sampling Explanations of Judgment Biases. *Psychological Review*. **18**(2) 379-392.
- Luce, R. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York, NY.
- March, J.G. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2**(1) 71-87.
- March, J.G. 1996. Learning to be risk averse. *Psychological Review* **103**(2) 309-319.
- March, J.G. 2003. Understanding Organizational Adaptation. *Society and Economy* **25**(1) 1-10.
- March, J.G. 2010. *The Ambiguities of Experience*. Cornell University Press, Ithaca, NY.

- McNamara, P., and Baden-Fuller, C. 1999. Lessons from the Celltech Case: Balancing Knowledge Exploration and Exploitation in Organizational Renewal. *British Journal of Management* **10**(4) 291-307.
- Miller, D. 1987. The structural and environmental correlates of business strategy. *Strategic Management Journal* **8**(1) 55-76.
- Nelson, R. 2008. Bounded rationality, cognitive maps, and trial and error learning. *Journal of Economic Behavior & Organization* **67**(1) 78-89.
- Nelson, R., and Winter, S.G. 1982. *An evolutionary theory of economic change*. Belknap Press, Cambridge, MA.
- Peters, T.J., and Waterman, R.H. 1984. *In search of excellence: Lessons from America's best-run companies*. Harper & Row, New York, NY.
- Radner, R., and Rothschild, M. 1975. On the allocation of effort. *Journal of Economic Theory* **10** 358-376.
- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* **58** 527-535.
- Schmalensee, R. 1975. Alternative models of bandit selection. *Journal of Economic Theory* **10**(3) 333-342.
- Sorenson, J., and Stuart, T.E. 2000. Aging, obsolescence, and organizational innovation. *Administrative Science Quarterly* **45**(1) 81-112.
- Sutton, R.S., and Barto, A.G. 1998. *Reinforcement learning: An introduction*. Cambridge, MA, MIT Press.
- Tushman, M., and Romanelli, E. 1985. *Organizational Evolution: A Metamorphosis Model of Convergence and Reorientation*. In L. L. Cummings and B. M. Staw (Editor.), *Research in Organizational Behavior*. 171-222.
- Vermorel, J., and Mohri, M. 2005. Multi-armed bandit algorithms and empirical evaluation. *Machine Learning: ECML 2005* 437-448.
- Weber, E., Shafir, S., and Blais, A.-R. 2004. Predicting Risk Sensitivity in Humans and Lower Animals: Risk as Variance or Coefficient of Variation. *Psychological Review* **111**(2) 430-445.
- Weick, K. 1969. *The social psychology of organizing*. Addison-Wesley, Reading, MA.
- Wholey, D., and Brittain, J. 1989. Characterizing Environmental Variation. *The Academy of Management Journal* **32**(4) 867-882.
- Whittle, P. 1988. Restless Bandits: Activity Allocation in a Changing World. *Journal of Applied Probability* **25** 287-298.

Figures

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8

