

Managing on Rugged Landscapes*

Steven Callander[†]

Stanford University

Niko Matouschek[‡]

Northwestern University

December 11, 2014

Abstract

An emergent theme in the study of organizations is the broad differences in managerial practices and performance across firms. We develop an explanation for these phenomena that turns on the complexity of the environments that firms operate in. We construct a model that formally captures the difficulty of the manager's problem and show how managers search for good managerial practices by combining theoretical knowledge with practical experience, learning as they go. In this setting the evolution of firms is path dependent, marked by numerous failures, successes, and reversals. Nevertheless, patterns emerge. We show in particular how initial differences in performance persist and grow in expectation over time. We then apply the model to several long-standing questions in the study of organizations, exploring how imitation and coordination interact with the difficulty of the manager's problem and impact the performance of firms. We also apply the model to the growth and development of nations, showing how the performance dynamics that emerge resonate with historical experience.

Keywords: organizations, learning, persistent performance differences, imitation, coordination, complementarities, complexity.

JEL classifications: D21, D83, L25, M1

*We are very grateful for comments from Wouter Dessein, Luis Garicano, Bob Gibbons, Marina Halac, Richard Holden, Emir Kamenica, David Kreps, Gilat Levy, Dan Levinthal, Derek Neal, Andrea Prat, Debraj Ray, Luis Rayo, Jan Rivkin, John Roberts, Mike Ryall, Balazs Szentos, Steve Tadelis, Tim Van Zandt, and participants at the NBER Organizational Economics meeting, the CEPR workshop on Incentives, Management, and Organization, the MIT Organizational Economics workshop, the Organizational Research group in Munich, LSE, Oxford, Purdue, Warwick, and several conferences. All remaining errors are our own.

[†]Graduate School of Business, Stanford University, sjc@stanford.edu.

[‡]Kellogg School of Management, Northwestern University, n-matouschek@kellogg.northwestern.edu.

1 Introduction

The corporate histories of even the most successful firms are littered with failed products, abandoned initiatives, and reorganizations that didn't pan out (see Harford (2011) for many examples). Often these failures reflect less the failings of the managers themselves than the sheer complexity of the environments that firms operate in. In such environments managers need to choose among many alternatives yet possess only a tenuous understanding of how these choices map into outcomes. As a result, they are left with little choice but to try an action to see if it works, learning as they go.

Understanding the manager's problem is an important step to understanding organizations and their role in the economy. Progress on this question, however, has been impeded by the absence of a model that captures formally the difficulty of the problem that managers face. The standard economic approach of smooth objective functions and convex action spaces leaves little scope for failures and missteps, describing a learning problem that is tractable yet overly simplistic. As Brynjolfsson and Milgrom (2013, p.14) observe: “[...] *as noted by Roberts (2004), much of the standard economic treatment of firms assumes that performance is a concave function of a set of infinitely divisible design choices, and the constraint set is convex. Under these conditions, decisionmakers can experiment incrementally to gradually identify an optimal combination of practices.*”

In contrast, the influential management literature on “rugged landscapes” takes a very different approach (Levinthal 1997; Rivkin 2000). In particular, it relaxes the assumption of concavity and generates “rugged” production functions with many peaks and troughs that can ensnare a manager's search. The downside of this richness, however, is that strategy itself is removed from the manager's toolkit. Instead, managers search over the landscape via exogenously imposed local search rules that take the form of hill climbing algorithms. This approach can explain why search may yield inefficient outcomes but it cannot address why managers would ever search in the manner prescribed. Indeed, managers are automatons in these models, lacking any understanding of the underlying process that generates the environment they face. As such, they are unable to form beliefs to support rational strategic action or even consciously decide in which direction they should search. The trade-off at the heart of this approach is concisely critiqued by Roberts and Saloner (2013, p.819) who observe: “*This newer model has pluses and minuses. On the one hand, dropping unwarranted but conventional assumptions [...] is clearly desirable. On the other hand, the approach assumes that there is no logic or theory that can be used to guide the selection. This conclusion seems to be too nihilistic.*”

In this paper we tread a middle ground between the rugged landscapes literature and the traditional economic approach to modeling organizations. We provide a formal model of the manager's

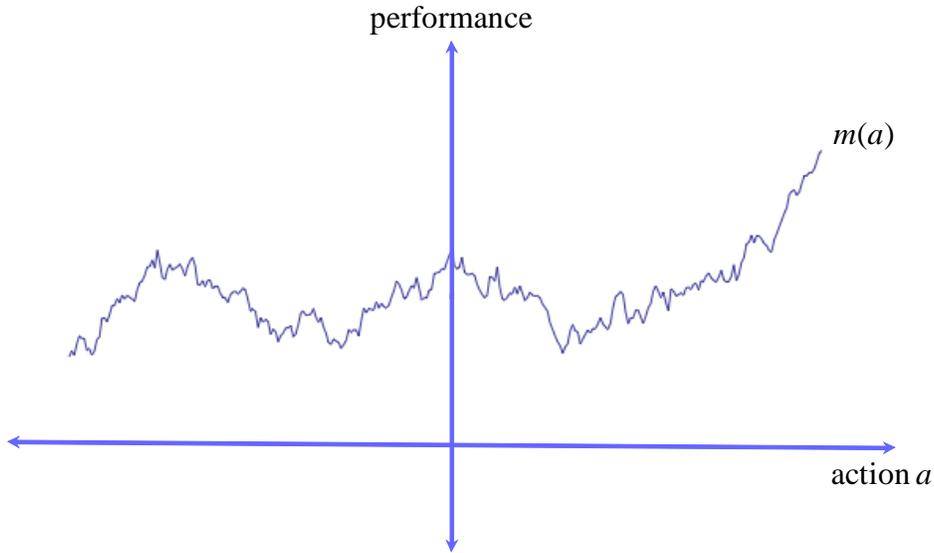


Figure 1: The production function $m(a)$

problem that captures the inherent difficulty of the task whilst allowing managers to search optimally over the space of actions. To this end, we model the production function as the realized path of a Brownian motion, as illustrated in Figure 1. An advantage of this approach is that managers possess an understanding of the underlying environment that they face. Knowledge of the drift and variance parameters of the Brownian motion provide the manager with theoretical knowledge that they can then combine with the lessons of experience to inform their beliefs and guide their search behavior. We characterize for this environment the optimal search rule and show that it takes a simple form when agents possess standard risk aversion and maximize their utility on a period-by-period basis. We then use this result to explore several long-standing questions in the economics of organizations.

The optimal search rule is intuitive and simple to describe, yet the performance dynamics that it generates are far from smooth. The sequence of actions that managers take is highly path dependent, exhibiting a mixture of successes, failures, and strategy reversals. A key property is that efforts by a manager to improve performance can be counter-productive and actually lead to worse performance. We show that the fear of these missteps restrains managers and slows down learning. Whilst in moderation these missteps present little more than speed bumps that wash out over time, significant dips can have permanent effects on performance, so much so that a sufficiently poor outcome can cause a manager to abandon the search for better performance, derailing growth

altogether. The ups and downs of the growth path also imply that firms can overtake—and be overtaken by—other firms. This richness captures the perplexing regularity that a market leader one day can retreat to the middle of the pack the next, and disappear altogether the day after that.

We then turn to aggregate properties of the growth path to understand broader patterns in firm performance. Our first main finding is that good performance begets good performance. A well-known stylized fact in the economics of organizations is that firms exhibit persistent performance differences, or PPDs (Syverson 2011; Gibbons and Henderson 2013). We show that, despite the richness of individual growth paths, initial advantages persist in expectation and feed on themselves, growing over time. In complex decision environments, therefore, performance does not converge across firms. The difficulty in finding good actions restrains the ability of under-performing firms to catch up. More surprisingly, better performance actually eases pressures on managers, enabling them to experiment more boldly. This explains how non-trivial differences in performance can emerge and persist in seemingly similar organizations.

The persistence of performance differences raises several subsequent questions. From where do the performance differences originate? And why do they not disappear due to imitation? We devote the second half of the paper to these questions. The answers we provide – based on complementarities and coordination – connect our model, and the difficulty of the manager’s problem, to several long-standing areas of interest in organizational economics.

We begin with imitation. In practice, organizations do not exist in isolation, and the existence of PPDs creates a puzzle as to why under-performing firms don’t simply imitate the actions of their better performing rivals. Surprisingly, the failure to do so is not only a matter of observability as the actions of better performing firms are often well known.¹ Nevertheless, successful imitation does frequently occur, and the puzzle, at its root, is why imitation is sometimes possible and sometimes not. Milgrom and Roberts (1995) argue that a key element of the difficulty is complementarities across firm actions. We incorporate complementarities into the model in the spirit of Milgrom and Roberts. This, combined with the uncertainty inherent in our model, renders imitation problematic. Our main result then builds on this. It fixes the degree of complementarity and shows how the decision to imitate – and the effectiveness of imitation – then depends on firm strategy and performance, and we quantify this dependence. By offering a model with a conception of distance in the space of managerial actions, we are able to show how imitation is profitable only

¹“Such firms as Lincoln Electric, Walmart, or Toyota enjoyed sustained periods of high performance. As a result, they were intensively studied by competitors, consultants, and researchers, and many of their methods were documented in great detail. Nonetheless, even when competitors aggressively sought to imitate these methods, they did not have the same degree of success as these market leaders” (Brynjolfson and Milgrom 2013, p.14).

when competitors are ‘close’ to each other. Significantly, the relevant measure of closeness is not with respect to performance as, indeed, the better a competitor performs the more attractive is imitation. Rather, the measure of closeness that matters is in the action space. We show how it is possible for a competitor to simply be ‘too far ahead’ in the action space – that the managerial practices of the firms are simply too different – for imitation to be profitable.

Another practical reality of modern management is that large organizations require a degree of decentralized decision making. Decentralization, in turn, creates the need to coordinate actions across divisions. An open question – one that we take up here – is how the need to coordinate actions interacts with learning and incentives to experiment. Surprisingly, we show that coordination failures can lead to *more* experimentation than is optimal for the firm as well as to less. This result upends an intuition of Roberts (2004, pp. 60) who, writing for practicing managers, argues that the need to coordinate will only dampen the willingness of individual managers to search for better practices.

Most strikingly, however, we provide an explanation for why coordination problems emerge at all. Coordination problems create multiple routes forward for firms. In providing an explanation for the endogenous emergence of this multiplicity, therefore, our theory provides an explanation for the endogenous adoption of different strategies – and different performance levels – for firms that are otherwise identical.

This result is of broader interest, in particular to the analysis of coordination games. A significant feature of our contribution is that we do not build coordination difficulties into the problem directly. Indeed, at the beginning of search the managers face a unique equilibrium that is efficient and coordination problems play no role. Rather, coordination problems emerge *endogenously* as a function of firm strategy and performance and we characterize the conditions for when they emerge. Notably, we find that the main predictive factor to the emergence of coordination problems is the consistency of firm performance. Specifically, a firm that always exceeds performance encounters no coordination problems, always facing a unique equilibrium that is efficient. In contrast, a firm that falls short of expectations – even if previous performance levels are exceeded – faces coordination difficulties from then on, leading to more or less experimentation than otherwise would occur and potentially derailing the firm from its optimal growth path. The endogeneity of coordination problems suggests that some firms do indeed lead a seemingly charmed existence, performing well and immune from coordination problems, whereas other firms both struggle in performance and succumb to coordination mistakes.

Throughout the paper we cast our model in the context of firms, although the connection to

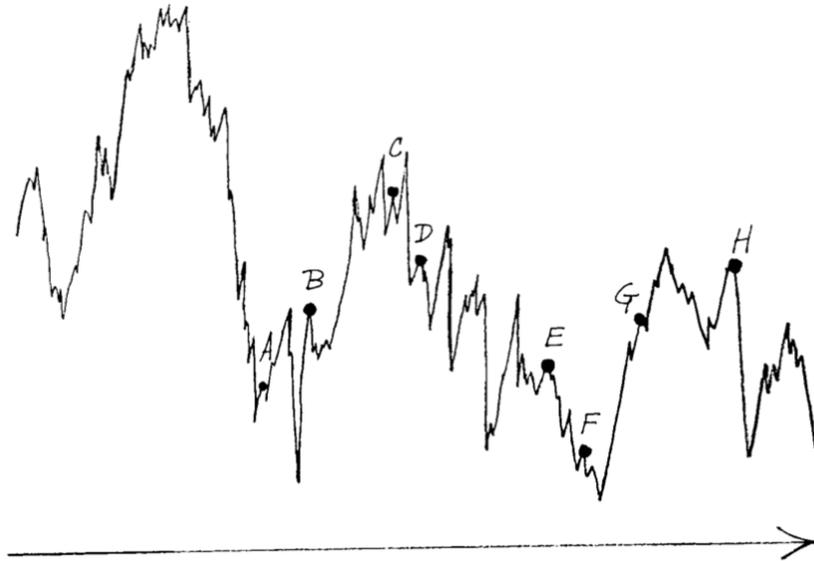


Figure 2: Reproduction of Figure 5.2 in Matsuyama (1996): “The horizontal axis represents the space of all possible economic systems, and the graph represents their performance levels.”

other types of organizations is readily apparent. One natural application is to economic development and the growth of nations. Societies face many decisions about how to structure their economies, decisions that are fraught with uncertainty and that necessitate trial-and-error learning. For instance, the futility of communism and *laissez faire* as ways to organize an economy were revealed only through painful experience. Even with this experience—and agreement that better performing systems lay between these extremes—identifying the optimal arrangement remains a difficult task.

Applied to the problems of economic development, our results on PPDs, path dependence, and coordination failures, provide an explanation as to why nations do not converge in wealth over time and why the practice of poor countries imitating rich countries has such a checkered history. The tight connection between the problems facing firms and those facing nations is laid out particularly clearly in a policy paper by Matsuyama (1996) who imagined that it is a highly non-monotonic function that best captures the performance of economies and the “ruggedness” of the development problem. Matsuyama writes “*The rugged nature of the graph [reproduced here as Figure 2] captures the inherent complementarity of activities in each system; the performance of an economy can change drastically by a small change in selection of activities. [...] There are a large number of locally optimal systems, and each society has evolved into one of them. There is*

no way for society to search in a systematic way for the global optimum, or other local optima that are more efficient” (Matsuyama 1996, pp.145).

The similarity between this description of economic development and the problems facing firms, as well as between Figures 1 and 2, is striking. The connections, we expect, will be apparent as we proceed through the text. For simplicity in presentation, however, we relegate an explicit treatment of development and growth to the concluding discussion.

2 Related Literature

The rugged landscapes approach was introduced to the management literature by Levinthal (1997). Levinthal sought to resolve a long-standing debate in the field of organizational change as to whether change predominantly originated from organizational adaptation (Cyert and March 1963) or from population-level selection effects through the birth and death of organizations (Hannan and Freeman 1977). The rugged landscapes approach allowed Levinthal to demonstrate that these arguments are complementary and can coexist. By modeling organizations as locally-adapting, Levinthal showed how a firm’s performance does depend on its own actions, yet at the same time the ruggedness of the landscape makes achieving global optima prohibitively difficult, opening the door for population level selection effects. Levinthal’s argument made a significant contribution to the debate on organizations in management and sociology. However, as we noted earlier, it does not address the question of why organizations follow the assumed search rules. This is the question of interest in our paper.

The rugged landscapes approach was further developed by Rivkin (2000) who applies it to a question more in line with the focus of our paper. In particular, Rivkin (2000) uses the rugged landscapes approach to demonstrate the difficulty of imitation in complex environments, especially when actions are complementary, using the theory of NP-completeness. He argues that when the set of actions that must be undertaken grows large, and the complementarities between actions is unknown, it is computationally infeasible for a firm to exactly imitate the strategy of a more successful rival. Our focus, in contrast, is on the deliberate choice of managers whether to imitate rather than on the feasibility of imitation *per se*.

The notion of problem difficulty in the management of organizations has appeared in the economics literature in Garicano (2000). We share this emphasis with Garicano (2000) but otherwise diverge in approach and application. Garicano conceives of organizations as facing a stream of problems that must be solved. He then presumes that, due to specialization and talent, workers differ in their ability to solve various problems. The animating question of his analysis is to show

how a hierarchical organizational structure most efficiently handles the tasks. In contrast, we endow our manager with standard economic capabilities and focus on a single problem to solve, exploring how the manager addresses this problem, and the dynamics of the learning rule and performance that is generated.

Formally, the Brownian motion framework corresponds to a bandit model with a continuum of correlated, deterministic arms. This contrasts with the standard assumption in the experimentation literature of stochastic and independent arms (see, for instance, Bolton and Harris 1999). The correlation across arms is important as it captures learning across alternatives. With independent arms a failed experiment is just that and nothing more: the failed alternative is discarded and it does not inform subsequent choice. In contrast, in the Brownian motion framework an experiment that fails nevertheless yields information that informs and guides future choices. This information dictates whether to continue experimenting, which direction to experiment in, and how bold to be. It is this path dependence that is the focus of our analysis. To be sure, this increase in realism comes with an analytic cost as we model decision makers who maximize utility on a period-by-period basis. The substantive conclusions we draw from the model do not obviously depend on this assumption, although further exploration is clearly warranted. We return to this point later in the discussion section.

The use of the Brownian motion to model experimentation was first introduced in Callander (2011). We depart from that paper in several key respects. Callander (2011) operates exclusively with quadratic utility to exploit the separability of mean and variance that it delivers. He also studies problems with an internal optimum (or ideal outcome). The agents know what this maximum possible outcome is and how close they are to it at any point in time. They also possess directional knowledge, whether their current action is to the left or right of the peak, and this knowledge structures their search behavior. His setting does not capture the typical economic environment in which income is unbounded and more income is preferred to less. This is the environment we study here. As is well known, quadratic utility is incompatible with this setting and, therefore, we identify conditions on utility necessary to support well-defined search. Rather than singling out a particular functional form, we work with a broad class of risk preferences that satisfy these requirements. Surprisingly, this richness leads to an optimal search rule that is qualitatively different yet at the same time simpler, which in turn allows broader insight and comparative statics. We then exploit this capability to explore several long-standing questions of complementarities, coordination, and imitation in organizations that have not previously yielded to economic methods.

3 The Model

There is a single manager. At the beginning of every period $t = 1, 2, \dots$, the manager takes an action that determines his performance level and, thus, his income. After the manager has consumed his income, time moves on to the next period. Our aim is to characterize the manager's optimal actions given the technology, preferences, and information structure that we describe next.

Technology: The manager's action $a_t \in \mathbb{R}$ determines his income level $m_t \in \mathbb{R}$ according to the production function $m(a_t)$, where $m : \mathbb{R} \rightarrow \mathbb{R}$. We follow Callander (2011) and model m as the realized path of a Brownian motion with drift $\mu > 0$ and variance $\sigma^2 > 0$. For reasons that will become apparent, we interpret the variance σ^2 as a measure of the complexity of the production process. Moreover, we refer to $a_0 = 0$ as the status quo action and denote status quo income by $m_0 = m(a_0)$. The realized path of the Brownian motion is determined by nature before the start of the game and does not change over time. Figure 1 depicts one possible realization of the Brownian motion.

Preferences: The manager's utility is given by $u(m)$, where m is his income. We assume that this function is four times continuously differentiable and satisfies $u'(m) > 0$ and $u''(m) < 0$ for all $m \in \mathbb{R}$. The first condition implies non-satiation and the second risk aversion.

We further assume that the utility function satisfies *standard risk aversion* which requires that a risk cannot be made more desirable by the imposition of an independent, loss-aggravating risk (Kimball 1993). Formally, suppose there are two independent random variables x and y such that

$$\mathbb{E}[u(m+x)] - u(m) \leq 0 \quad \text{and} \quad \mathbb{E}[u'(m+y)] - u'(m) \leq 0.$$

The risk x is undesirable since the manager would turn it down if someone offered it to him. And the risk y is loss aggravating since it increases the manager's expected marginal utility; it therefore makes a sure loss more undesirable and is itself more undesirable if the manager is exposed to a sure loss. The utility function then satisfies standard risk aversion if and only if

$$\mathbb{E}[u(m+x+y) - u(m+y)] \leq \mathbb{E}[u(m+x) - u(m)],$$

that is, if and only if the loss aggravating risk y does not make the undesirable risk x more desirable.

Most commonly used utility functions that exhibit non-satiation also exhibit standard risk aversion, including exponential, logarithmic, and power functions. Standard risk aversion, however, is stronger than decreasing absolute risk aversion in that the former implies the latter but the reverse

does not hold (Kimball 1993). We assume more than decreasing absolute risk aversion since this type of risk aversion is not enough to ensure that independent risks are substitutes. As such, decreasing absolute risk aversion is less useful in settings such as ours in which an agent has to choose between multiple, risky payoffs than it is in settings in which an agent has to choose between a single, risky payoff and a safe one (see, for instance, Chapter 9 in Gollier (2001)).

We denote the coefficient of absolute risk aversion by $r(m) = -u''(m)/u'(m)$. Since standard risk aversion implies decreasing absolute risk aversion we have $r'(m) \leq 0$ for all $m \in \mathbb{R}$. We will see below that for a non-trivial solution to the manager's problem to exist, $r(m)$ has to cross the ratio $2\mu/\sigma^2$, where μ and σ^2 are the drift and the variance of the Brownian motion. Unless we explicitly say otherwise, we therefore assume that $r(m)$ does cross $2\mu/\sigma^2$ and we denote the largest income level for which $r(m) = 2\mu/\sigma^2$ by \hat{m} . An example of a standard utility function that satisfies our conditions is given by $u(m) = \alpha m - \exp(-\beta m)$, where $\alpha > 0$ and $\beta > 2\mu/\sigma^2$.

Information: In any period, the manager knows the income generated by the status quo action and by any action he took in any previous period. We refer to these actions as “known actions” and to all other actions as “unknown actions.” In addition to the known actions, the manager knows that the production function was generated by a Brownian motion with drift μ and variance σ^2 . The manager does not, however, know the realization of the Brownian motion. In any period t , the manager's information set is therefore given by $I_t = \{\mu, \sigma^2, (a_0, m_0), \dots, (a_{t-1}, m_{t-1})\}$.

Optimal Learning Rule: For simplicity we assume that the manager maximizes expected utility on a period-by-period basis. An optimal learning rule is therefore given by (a_1^*, a_2^*, \dots) , where

$$a_t^* \in \arg \max_{a_t} \mathbb{E}[u(m_t) | I_t].$$

Our goal is to characterize the set of optimal learning rules.

4 Beliefs and Expected Utility

We start by examining the manager's beliefs about the income generated by any unknown action. For this purpose, consider any period t and let l_t and h_t denote the left-most and right-most known actions. Consider now an unknown action a_t that is to the right of h_t . For any such action, the manager believes that income $m(a_t)$ is drawn from a normal distribution with mean

$$\mathbb{E}[m(a_t)] = m(h_t) + \mu(a_t - h_t) \tag{1}$$

and variance

$$\text{Var}(m(a_t)) = (a_t - h_t) \sigma^2, \quad (2)$$

where μ and σ^2 are the drift and the variance of the Brownian motion. The manager therefore expects an action to generate higher income, the further it is to the right of h_t . At the same time, however, the further an action is to the right of h_t , the more uncertain the manager is about the income generated by that action. The beliefs for actions to the left of the left-most action l_t are analogous.

Notice that the manager's beliefs about any action to the right of h_t depend only on h_t and that his beliefs about any action to the left of l_t depend only on l_t . Similarly, the manager's beliefs about any action between l_t and h_t depend only on the known actions closest to that action. To see this without having to introduce more notation, suppose that there are no known actions between l_t and h_t . For any action $a_t \in [l_t, h_t]$, income $m(a_t)$ is normally distributed with mean

$$\text{E}[m(a_t)] = \frac{a_t - l_t}{h_t - l_t} m(h_t) + \frac{h_t - a_t}{h_t - l_t} m(l_t) \quad (3)$$

and variance

$$\text{Var}(m(a_t)) = \frac{(a_t - l_t)(h_t - a_t)}{h_t - l_t} \sigma^2. \quad (4)$$

The manager's expected income is therefore a convex combination of the income generated by l_t and h_t . Moreover, the further the action is from the closest known action, the more uncertain the manager is about the income generated by that action.

The assumption that the production function is generated by a Brownian motion therefore ensures that the manager's beliefs take a simple form that satisfies several intuitive properties. First, beliefs are normally distributed. Second, the manager knows more about an action, the closer the action is to a known action, and the less complex is the production process. Third, the manager engages in directed search, that is, he knows where he can expect better actions and, as we will see below, he focuses his search in that direction. And finally, even if, over time, the manager learns the income generated by a very large number of actions, he can never infer the entire production function. In this sense, there is a limit to theoretical knowledge and thus a deep need to learn by trial-and-error.

Now that we have examined the manager's beliefs, we can specify his expected utility. The simplicity of Equations (1)-(4), combined with normally distributed uncertainty enables a key simplification of the problem. It does this by allowing the problem to be recast in mean-variance space, which allows us, in turn, to leverage several existing results. Specifically, suppose that the

manager believes that income is normally distributed with mean M and variance V and let z denote a random variable that is drawn from a standard normal distribution. We then have the following lemma, which is proven in “Hilfsatz” 4.3 in Schneeweiss (1967) and Theorem 1 in Chipman (1973).

LEMMA 1 (Schneeweiss 1967 and Chipman 1973). *Suppose that $|u(m)| \leq A \exp(Bm^2)$ for some $A > 0$ and $B > 0$. Then the expected utility function*

$$W(M, V) = \mathbb{E} \left[u \left(M + \sqrt{V}z \right) \right]$$

exists for all $M \in (-\infty, \infty)$ and $V \in (0, 1/(2B))$.

The restriction in the lemma ensures that expected utility is integrable, and for the remainder of the paper we assume that it holds. Notice that since we are free to choose any positive parameters A and B , this restriction is mild. With this translation in hand, we can then establish the concavity of expected utility.

LEMMA 2. *The expected utility function $W(M, V)$ is concave.*

The lemma follows from Theorem 3 in Chipman (1973) and Theorem 2 in Lajeri and Nielsen (2000). Specifically, Theorem 3 in Chipman (1973) states a condition that ensures concavity of the expected utility function and Theorem 2 in Lajeri and Nielsen (2000) shows that this condition is equivalent to absolute prudence $-u'''(m)/u''(m)$ being decreasing in m for all $m \in \mathbb{R}$. The lemma then follows from the fact that in a setting such as ours, in which income is unbounded, decreasing absolute prudence is equivalent to standard risk aversion (Proposition 3 in Kimball (1993)). In the next section we will see that in the relevant range both expected income and its variance are linear in the manager’s action. Standard risk aversion therefore ensures that the manager’s problem is concave.

5 Managerial Learning

We begin by focusing on the optimal action in the first period and then turn to subsequent periods.

5.1 The First Period

In the first period, the manager faces a continuum of possible actions, with each action representing a different expected income and risk level. The set of possibilities resembles an investment line, albeit in mean-variance space, with the mean and variance increasing linearly in the manager’s action. If the manager remains at the status quo he encounters no risk and realizes status quo income m_0 . All other possible actions involve risk. However, it is clear that the manager will never

choose an action to the left of the status quo, $a_1 < a_0$, as this lowers expected income and adds risk. Actions to the right of the status quo, on the other hand, offer potentially attractive trade-offs of expected income against risk.

Proposition 1 below establishes that managers do engage some risk in the first period, experimenting in search of better performance. However, a manager experiments only if his starting income level is sufficiently high. Moreover, the willingness of a manager to experiment boldly is strictly increasing in their status quo income. Higher starting income, therefore, favors experimentation unambiguously in the first period.

To understand the result (and state it completely) some formality is required. Let $\Delta_1 = a_1 - a_0$ denote the size of the step the manager takes in the first period, where $\Delta_1 \geq 0$ as the manager only searches to the right. We then know from (1) and (2) that the manager's expected income is $m_0 + \mu\Delta_1$ with variance $\sigma^2\Delta_1$. The manager's problem can then be written as

$$\max_{\Delta_1 \geq 0} W(m_0 + \mu\Delta_1, \sigma^2\Delta_1),$$

where $W(\cdot)$ is the expected utility function defined in Lemma 1. As observed above, Lemma 2 ensures that this problem is concave.

Next, by differentiating $W(\cdot)$ with respect to Δ_1 , we obtain

$$\frac{dW(m_0 + \mu\Delta_1, \sigma^2\Delta_1)}{d\Delta_1} = \mathbb{E} \left[u' \left(m_0 + \mu\Delta_1 + \sigma\sqrt{\Delta_1}z \right) \right] \frac{\sigma^2}{2} \left(\frac{2\mu}{\sigma^2} - R(m_0, \Delta_1) \right), \quad (5)$$

where

$$R(m_0, \Delta_1) \equiv - \frac{\mathbb{E} [u''(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z)]}{\mathbb{E} [u'(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z)]}$$

and where we make use of the fact that $\mathbb{E}[u'(\cdot)z] = \sigma\sqrt{\Delta_1}\mathbb{E}[u''(\cdot)]$.

The sign of expected marginal utility is therefore determined by the relative size of the ratio $2\mu/\sigma^2$ and $R(m_0, \Delta_1)$. This relationship plays a key role in our analysis. Notice that $R(m_0, \Delta_1)$ is the coefficient of absolute risk aversion for the indirect utility function $\mathbb{E}[u(\cdot)]$ which is, in general, different from the expected value of the coefficient of absolute risk aversion $r(\cdot)$.

These measures are equated, however, for initial experimentation, and for $\Delta_1 = 0$ we have $R(m_0, 0) = r(m_0)$. Thus

$$\frac{dW(m_0, 0)}{d\Delta_1} \begin{cases} > 0 & \text{if } m_0 > \hat{m} \\ \leq 0 & \text{if } m_0 \leq \hat{m}, \end{cases} \quad (6)$$

where \hat{m} denotes the largest income level m for which $r(m) = 2\mu/\sigma^2$. This establishes the income threshold, below which a manager chooses to retain what they have and not experiment further. All firms face the same trade-off between risk and marginal return, yet it is only the high performing

firms that choose to search for better performance. The rich get richer, so to speak, whereas the lower performing firms face a performance trap. Managers who start with poor performance are more averse to risk, and this deters them from even trying to identify better managerial practices, condemning them to stagnant performance.

The complete characterization of first period behavior is then given by the following.

PROPOSITION 1. *The manager's optimal first period action is unique and given by*

$$a_1^* = \begin{cases} a_0 + \Delta(m_0) & \text{if } m_0 \geq \hat{m} \\ a_0 & \text{if } m_0 < \hat{m}, \end{cases}$$

where $\Delta(m_0)$ is implicitly defined by $R(m_0, \Delta(m_0)) = 2\mu/\sigma^2$ and satisfies $\Delta(\hat{m}) = 0$ and $\Delta'(m_0) > 0$ for all $m_0 \geq \hat{m}$.

The proposition reveals clearly two important properties of manager behavior: That low performance firms don't search – the performance trap – and that, among those that do search, the better performing a firm is the more assiduously the manager searches for even better actions. A third property that the proposition reveals, less evident but no less significant, is that an optimum action exists at all, particularly for firms that begin with high wealth. This is not obvious as ever bolder experiments produce a wealth effect. The expected outcome increases, and the coefficient of absolute risk aversion at this income declines, yet the marginal risk-reward trade-off remains at a constant rate. A seemingly reasonable conjecture, then, is that this wealth effect will inspire a manager who is willing to engage in some risk to engage in evermore risk, to the point that a finite optimal action does not exist. Proposition 1 demonstrates that this conjecture is false. Regardless of how high the starting income, the steady accumulation of risk eventually dominates the wealth effect of search, and that even the best performing firms and most bold managers reach a point where the risk is too much to bear and they experiment no further.

That a manager experiences a declining marginal benefit from each increment in the size of the experiment manifests formally in the concavity of expected utility. Yet concavity alone is insufficient to guarantee a finite and non-trivial optimum. Necessary too is the crossing condition that the coefficient of absolute risk aversion crosses $2\mu/\sigma^2$ for some income level. To confirm this point, Proposition 2 relaxes the crossing condition and establishes that in this case the solution to the manager's problem is either trivial or does not exist.

PROPOSITION 2. *If the coefficient of absolute risk aversion $r(m)$ satisfies $r(m) > 2\mu/\sigma^2$ for all $m \in \mathbb{R}$, the manager does not engage in search in the first period (or any subsequent period). If, instead, $r(m) < 2\mu/\sigma^2$ for all $m \in \mathbb{R}$, an optimal action does not exist. Finally, if $r(m) =$*

$2\mu/\sigma^2$ for all $m \in \mathbb{R}$, then in any period t the manager is indifferent between the right-most action and any action to its right.

When $r(m) > 2\mu/\sigma^2$ for all $m \in \mathbb{R}$ the manager is too risk averse to engage any risk, regardless of the income level. The result follows from the definition of \hat{m} in Proposition 1. If, instead, $r(m) < 2\mu/\sigma^2$ the manager is sufficiently risk tolerant to engage risk for any income level. However, as this implies that

$$-\mathbb{E} \left[u'' \left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z \right) \right] < 2\mu/\sigma^2 \mathbb{E} \left[u' \left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z \right) \right],$$

it follows from (5) that marginal expected utility is strictly positive for all $\Delta_1 \geq 0$. An optimum therefore doesn't exist. Finally, when $r(m) = 2\mu/\sigma^2$ for all $m \in \mathbb{R}$, the manager is indifferent whether to undertake risk at every wealth level, and an analogous argument to the previous case establishes the result.

5.2 The Second and Subsequent Periods

We just saw that if $m_0 \leq \hat{m}$ the manager takes the status quo action in the first period and learns nothing further about the mapping. In this case the manager then faces the same problem in any subsequent period and always chooses the status quo action again. Consequently, for the remainder of the paper we suppose that $m_0 > \hat{m}$ holds.

After experimenting in the first period, the manager learns an additional point in the mapping and now knows two actions: the status quo action a_0 and the manager's first period action a_1^* . From (3) it is clear that any action on the bridge between a_0 and a_1^* is dominated by one of the ends as it delivers a lower expected income with uncertainty. Thus, if the manager is to experiment further, he continues to search in the same direction to the right. In forming beliefs about what to expect to the right of a_1^* , however, the income at a_1^* is the only relevant information. Thus, in deciding whether further experimentation is preferred to action a_1^* , the manager faces the exact same problem as he did in the first period.

The second period differs from the first period in two key respects, however. First, the income at a_1^* will almost surely be different from the status quo income; it is likely to have increased but may have decreased. In fact, the income at a_1^* could be so low as deter further experimentation altogether. From Proposition 1 we know that this is the case if $m_1^* \leq \hat{m}$. The second difference between the periods is that, should the manager decide not to experiment, he has two known actions from which to choose. He may repeat action a_1^* or he may reverse course and revert to the status quo action a_0 . Reversing course is clearly optimal if second period income is so low as to fall below

\widehat{m} . Surprisingly, the manager will choose to reverse course and stop experimenting even for income levels above \widehat{m} . At these income levels the manager prefers to experiment rather than stay at a_1^* , yet as that is not his only option he instead prefers to backslide to the certainty of the status quo income rather than experiment further. This possibility is captured in the following lemma.

LEMMA 3. *There exists a threshold level of income $\widetilde{m}(m_0) \in (\widehat{m}, m_0)$ such that*

$$u(m_0) = W(\widetilde{m}(m_0) + \mu\Delta(\widetilde{m}(m_0)), \sigma^2\Delta(\widetilde{m}(m_0))), \quad (7)$$

where $\Delta(\widetilde{m}(m_0)) > 0$. Moreover, the derivative of $\widetilde{m}(m_0)$ satisfies $0 < \widetilde{m}'(m_0) \leq 1$.

This implies that if first period income m_1^* is equal to $\widetilde{m}(m_0)$, the manager is indifferent between the status quo and further experimentation. It then follows that the manager strictly prefers experimentation to the status quo if $m_1^* > \widetilde{m}(m_0)$ and that he strictly prefers the status quo to experimentation if $m_1^* < \widetilde{m}(m_0)$. Thus, for income levels m_1^* strictly between $\widetilde{m}(m_0)$ and \widehat{m} , learning stops and the manager reverts to the status quo even though the marginal return from engaging in further search is positive.

The problem the manager faces in any period $t > 2$ is very similar to the one he faces in period 2. The only difference is that in any period $t > 2$, the manager does not necessarily compare his expected utility from engaging in further search with his utility from the status quo, as he does in the second period. Instead, the manager compares his expected utility from engaging in further search with his utility from whatever known action generates the largest income level, which may be the status quo action or some other known action.

To state the proposition that characterizes the manager's optimal action in all periods $t \geq 2$, let \overline{m}_t denote the largest known income level in period t , that is, let

$$\overline{m}_t = \max\{m_0, m_1^*, m_2^*, \dots, m_{t-1}^*\}.$$

Also, let \overline{a}_t denote the action that generates \overline{m}_t , that is, let

$$\overline{a}_t \in \{a_0, a_1^*, a_2^*, \dots, a_{t-1}^*\} \text{ such that } m(\overline{a}_t) = \overline{m}_t.$$

And finally, recall that h_t denote the right-most known action in period t . We can then state our next proposition.

PROPOSITION 3. *The manager's optimal action in period $t \geq 2$ is unique and given by*

$$a_t^* = \begin{cases} h_t + \Delta(m(h_t)) & \text{if } m(h_t) > \widetilde{m}(\overline{m}_t) \\ a(\overline{m}_t) & \text{if } m(h_t) \leq \widetilde{m}(\overline{m}_t), \end{cases}$$

where $\Delta(m) > 0$ is the Δ that solves $R(m, \Delta) = 2\mu/\sigma^2$ and $\tilde{m}(\cdot)$ is defined in (7).

In any period $t \geq 2$, the manager therefore engages in search if and only if the income level $m(h_t)$ associated with the right-most previously taken action h_t is above a threshold $\tilde{m}(\bar{m}_t)$, where the threshold is increasing in the largest known income level \bar{m}_t .

Together Propositions 1 and 3 characterize the optimal search rule and show that it takes a simple form. Figuratively, the manager's exploration of the rugged performance landscape depends on his starting point. If that point is too low, he just stays put. If his starting point is sufficiently high, he starts exploring the rugged landscape by taking discrete steps towards the right. He continues this rightward march indefinitely unless his income falls off a sufficiently large cliff. If this were to happen, continuing the search bears too much risk for the manager and he instead reverts to the highest peak he discovered during his exploration. Strikingly, this peak is only a local peak in expectation. The manager does not know if a better action lies in the vicinity of this action and it is not worthwhile for him to find out. Moreover, the manager believes with probability one that better actions exist far to the right. Nevertheless, he chooses not to seek out better performance as, despite the allure of higher income, the threat that his efforts will lower his performance is too great. It is this risk that compels him to stop searching altogether.

6 Persistent Performance Differences

A key feature of the optimal search rule is that it depends on the manager's status quo income. A natural question then is whether two managers with different status quo income levels can expect their incomes to converge or diverge over time. The answer is not immediate since there are forces that go in either direction: on the one hand, higher status quo income favors experimentation in the first period; on the other hand, however, it also makes it more tempting to stop experimentation and revert to the status quo in subsequent periods. The next proposition shows that on average the desire to experiment dominates.

PROPOSITION 4. *Suppose there are two managers, Manager H and Manager L. The production function of Manager $k = L, H$, is characterized by status quo income m_0^k , drift μ , and variance σ^2 , where $m_0^H > m_0^L > \hat{m}$. Then*

$$E_1 [m_t^*(m_0^H) - m_t^*(m_0^L)] > m_0^H - m_0^L \text{ for all } t = 1, 2, \dots, \quad (8)$$

where $E_1[\cdot]$ are the expectations taken at the beginning of the first period.

It is easy to show that there are realizations of the production functions for Managers H and L that are consistent with any dynamic: convergence, divergence, overtaking, and so on. The proposition, however, shows that on average, income levels diverge over time. In particular, it shows that if Manager H 's status quo income is one dollar above Manager L 's, then Manager H 's expected income will be strictly more than one dollar above Manager L 's expected income in every subsequent period. On average, therefore, income diverges rather than converges over time.

Proposition 4 provides an explanation for why a difference that exists between firms may persist and grow over time, but it doesn't explain where the initial difference originates from. The persistence and growth of differences are consistent with evidence on persistent performance differences in organizations. In Section 8 we show how coordination problems—and equilibrium multiplicity—can explain the origin of performance differences, even when firms are otherwise identical.

7 The Risk of Imitation

The model until this point is of a single firm searching over time. In practice, of course, organizations rarely operate in isolation. This naturally raises the question of imitation and why trailing firms do not simply imitate market leaders. That this doesn't occur as frequently as one may expect has been a long-standing puzzle in the study of organizations. The failure is both because firms often choose not to imitate their more successful brethren, and because attempts at imitation often fail, sometimes spectacularly so. The intriguing aspect of this phenomenon is that the strategies of market leaders are often readily observed by competitors, with the revelation sometimes offered willingly by the market leader. One such famous example is the case of Lincoln Electric Company (see also footnote 1).

Taking up the example of Lincoln Electric, Milgrom and Roberts (1995, pp.203-4) suggest informally that imitation is hampered by complementarities within organizations: “*An important puzzle is why Lincoln's successes have not been copied. [...] Our discussion suggests that Lincoln's piece rates are a part of a system of mutually enhancing elements, and that one cannot simply pick out a single element, graft it onto a different system without the complementary features, and expect positive results.*” We incorporate this intuition into our model, operationalizing complementarities by supposing that firm performance depends on two interdependent choices, a common, industry-wide choice as well as an idiosyncratic, individual firm choice. This, combined with the uncertainty in our model over the outcome of actions, renders imitation problematic when only knowledge of the industry-wide dimension is transferable.

This construction yields the immediate, and obvious, implication that imitation is less viable

– and less profitable – the greater is the complementarity. Our main result, however, goes beyond this. We fix the degree of complementarity and show how the viability of imitation then depends, in turn, on firm strategy and performance, and we quantify this dependence. We also provide a richer account of what successful imitation means, both in terms of outcomes, and in the strategies that are chosen when a trailing firm ‘imitates’ a market leader.

Formally, we extend the model such that the manager takes two actions, denoted by a_t^X and a_t^Y for period t . These actions produce total income

$$m_t = m^X(a_t^X) + m^Y(a_t^Y),$$

where $m^X(\cdot)$ and $m^Y(\cdot)$ are the production functions for the two actions. Each production function is an independently realized path of a Brownian motion with drift $\mu/2$ and variance $\sigma^2/2$. The status quo actions are given by $a_0^X = a_0^Y = a_0$ and each status quo action generates status quo income $m^X(0) = m^Y(0) = m_0/2$.

To enable imitation, we then add a second firm, Firm B , whose actions on the two dimensions are denoted b^X and b^Y . We impose two requirements consistent with the discussion in Milgrom and Roberts (1995). First, we suppose that dimension X represents a common, industry-wide determinant of firm performance, whereas dimension Y is an idiosyncratic, firm-specific factor. The same set of actions may be available to all firms, but because of firm-specific factors the outcome on this dimension may be very different. Formally, Firm B ’s production function on the second dimension is a distinct function, $m^Z(\cdot)$. For instance, a competitor may be able to copy perfectly Lincoln Electric’s piece rate labor contract, but within the culture of the new firm this contract will produce a very different outcome to what it produces at Lincoln Electric. This formulation is consistent with the frequent observation in practice that firms in the same industry, with strategies that are indistinguishable, nevertheless produce very different performance.

The second requirement we impose is that the two dimensions of choice are complementary. In fact, to allow use of the results from the previous section, we assume that the actions are strict complements. The manager therefore always has to set the actions equal to each other, that is, he has to set $a_t^X = a_t^Y = a_t$. Specifically, before contemplating imitation, the manager’s problem in the first period is given by

$$W(m_0 + \mu\Delta_1, \Delta_1\sigma^2) = \mathbb{E} \left[u \left(m_0 + \mu\Delta_1 + \sigma\sqrt{\Delta_1}z \right) \right],$$

where $\Delta_1 = a_1 - a_0$ denotes the size of the manager’s first period step. Notice that this expression is equivalent to the manager’s first period expected utility in our baseline model and recall that the optimal action is characterized in Proposition 1.

The manager's problem is different to that underlying Proposition 1, however, as the manager also has available information about Firm B . Our interest is in the conditions under which this information changes the manager's behavior and drives convergence in the performance of firms. The general analysis of this question depends on the particular histories of each firm and is analytically complicated. To capture the essential forces at work, we consider the special case of imitation in the first period by Firm A , albeit when the firms are at different actions and performance levels. Specifically, for Manager A we suppose his income satisfies $m_0 \leq \hat{m}$ such that, in isolation, he would choose not to experiment. For Firm B we suppose that his status quo actions are $d > 0$ units to the right of Firm A 's, that is, $b^X = b^Y = a_0 + d$, and that his income on dimension X is D units higher, that is, $m^X(d) = D + m_0/2$.

Manager A faces a trade-off in that whilst Firm B 's performance provides valuable information on the common dimension, it offers neither help nor guidance on the second dimension. And because of the complementarity of actions, if Manager A imitates his competitor, he is compelled to simultaneously move his action on the second dimension from the safety of the status quo to a risky alternative. The next proposition shows how this trade-off depends on the distance between firms in both the performance space and action space.

PROPOSITION 5. *Any action a_1^* that is optimal for Manager A satisfies*

$$a_1^* \in \begin{cases} (d, \infty) & \text{if } D > D_1, \\ \{d\} & \text{if } D \in (D_2, D_1], \\ (a_0, d] & \text{if } D \in (D_3, D_2], \\ [a_0, d) & \text{if } D \in (D_4, D_3), \text{ and} \\ \{a_0\} & \text{if } D \leq D_4, \end{cases}$$

where $D_1 > D_2 > D_3 > D_4 > 0$. Further, D_1 and D_3 are strictly increasing in d and σ^2 and $D_4 \rightarrow \infty$ as $d \rightarrow \infty$. Finally, the optimal action is unique if $D \geq D_2$ or $D \leq D_4$.

The proposition establishes that imitation is possible, that information from the other firm's performance may induce a change in management practice where it otherwise wouldn't occur. It also establishes that imitation does not always occur. Despite possessing information about how to achieve higher performance, the manager may nevertheless ignore this information and proceed as if the other firm didn't exist. The key contribution of Proposition 5, however, is to expose the factors on which the decision to imitate turns and to quantify this dependence.

Consistent with the intuition of Milgrom and Roberts (1995), the cost of imitation is driven by complementarity between actions. Proposition 5 then shows how, with the degree of complementarity fixed, the imitation decision depends on firm strategy and performance. The main insight is

that imitation is only profitable when the firms are ‘close’ to each other. The closeness is not measured in performance, however, but rather it is measured in the action space. Thus, it is possible that the strategy and practices of one firm are simply too far away from the trailing firm to make imitation worthwhile. A larger gap in the performance space does, in fact, make imitation more attractive. Yet we show in the proposition that for any distance in the performance space—thus, regardless how much better one firm performs than the other—the trailing firm will not imitate the leader if their strategies and actions are too different.

A novelty of the imitation in Proposition 5 is that it is not all-or-nothing. The trailing firm can use the information from the action of the leading firm yet not imitate that action precisely. There are many possibilities: imitation can be precise, it can be partial, or imitation can even move beyond the strategy of the other firm. The final possibility is perhaps the most surprising as it implies that a firm that was initially unwilling to experiment at all, is inspired to not only imitate the leading firm but to move beyond that firm’s strategy and experiment further. The key to understand this result is to note that successful imitation induces a “wealth effect.” By imitating Firm B ’s action, Firm A ’s effective income increases, which then relaxes Manager A ’s risk aversion and this encourages Firm A to initiate experimentation on its own. It is clear, then, that imitation, in all its forms, does not immediately imply convergence in performance. It also implies that, to an outside observer, identifying imitation is a non-trivial task. Yet all of the behaviors identified in Proposition 5 are rightly considered imitative as it is solely the information gleaned from the competitor’s performance that causes the firm to depart from its previous action.

For simplicity we do not explicitly model the behavior of Firm B . Yet, in light of the preceding point, how Firm B is acting is itself of interest. Several properties are immediately apparent. From the requirement that $D_4 > 0$, it is immediate that imitation cannot be mutual: If Firm A is imitating Firm B , even partially, then Firm B is not imitating Firm A . When this is the case, the behavior of Firm B is given simply by Proposition 1 and the single-agent model. For high values of D it is likely that Manager B also departs from his previous action and continues to experiment beyond action d . When both firms experiment, however, it is highly unlikely that they engage in the same experiment. The reason is two-fold. First, Firm A must engage risk in catching up to Firm B and, thus, is less willing than B to experiment further. Second, and more importantly, the firms will differ in performance on the idiosyncractic dimension and this will drive different experimentation levels. In fact, despite Firm A ’s greater baseline risk level, it may very well overtake Firm B and experiment further to the right than Firm B does itself. Counter-intuitively, it is then the imitating firm that is experimenting further to the right than is the imitated firm. This situation, roughly

speaking, approximates the real-world phenomenon of an innovative firm that identifies a valuable breakthrough, only to seemingly be unable to make the breakthrough work, and to have other firms adopt its technology and pass the innovating firm in performance.

A final observation on Proposition 5 is that while imitation is not always chosen in equilibrium, it is always possible. In fact, in the model, firms are able to perfectly observe the strategies of their competitors and replicate them without error. Nevertheless, a trailing firm may choose to not imitate and the reason for this choice is risk. The model captures the idea that whilst replication is possible, imitating a practice that is very different will cause a strain on all of the complementary practices necessary for firm performance. Although performance on the common dimension may be guaranteed, performance on the other dimensions are less sure, possibly improving along with the common dimension but also potentially failing miserably. We show how it is this risk that drives the imitation decision. The centrality of risk to imitation implies a connection between imitation and PPDs. It is risk that explains PPDs, restraining search and experimentation such that performance differences endure and grow, and it is risk that explains why imitation is difficult and why imitation does not lead to convergence in performance across firms.

8 Coordination Failures and Decentralized Search

Decision making in modern firms is by necessity decentralized. The sheer scale of enterprises, and the need to locate decision making where information is held, drive the logic of decentralization. A countervailing force is that decentralization creates the need to coordinate actions across the firm when actions are complementary. In this section we seek to understand these forces and how they interact with the need for managers to experiment and learn about managerial practices.

The literature has been surprisingly quiet on this connection, focusing instead on the role of asymmetric information and hierarchies (see Gibbons, Matouschek, and Roberts (2013) and Garicano and Van Zandt (2013) for surveys). The importance of coordination—and how it impacts managerial decision making—is, however, a recurring theme in Roberts’ (2004) book-length treatment of modern management. He describes the challenges in the following way: “*Search and change must be coordinated. [...] leaving individual managers in charge or particular elements of the organization to find improvements on their own can fail miserable, as can experimentation that is limited in scope. Both can fail to find the better solution and instead leave the firm stuck at an inferior coherent point*” (Roberts 2004, p.60).

Below we extend our model to address these questions formally. We upend Roberts’ intuition and show that coordination problems can lead to more experimentation than is optimal for the

firm as well as less. Most significantly, we show why and how coordination problems emerge endogenously as a function of firm strategy and performance. The multiplicity of equilibria thus produced provides an explanation for how otherwise identical firms can, nevertheless, adopt different managerial practices and depart onto very different growth paths.

Formally, we extend the model as follows. Let the firm now have two managers, A and B , who make decisions $a_t^A \in \mathbb{R}$ and $a_t^B \in \mathbb{R}$ in any period $t = 1, 2, \dots$, and realize incomes m_t^A and m_t^B . In particular, their incomes are given by

$$m_t^A = m(a_t^A) - \frac{1}{2}\delta(a_t^A - a_t^B)^2 \quad \text{and} \quad m_t^B = m(a_t^B) - \frac{1}{2}\delta(a_t^A - a_t^B)^2,$$

where $\delta > 0$ is a parameter that measures the importance of coordination between the managers' actions. The function $m(\cdot)$ is once again the realized path of a Brownian motion with drift μ and variance σ^2 and for which $m_0 = m(a_0)$.

An important feature of our set-up is that we assume the function $m(\cdot)$ is the same for both managers. This assumption is obviously an abstraction, yet it allows us to focus on the pure coordination aspect of the decision problems that managers face. In this setting the agents share the same rankings over actions and are not pulled apart by the different conditions they face. Thus, any failure to coordinate, or failure to coordinate on the optimal pair of actions, is due to the pure difficulty of coordination itself and attributable directly to the decentralization of decision making.

As before, the managers know that the production function is generated by a Brownian motion with drift μ and variance σ^2 . Moreover, in any period t , the managers know the actions that each took in previous periods and the income that these actions generated. Managers A and B therefore have the same information set

$$I_t = \{\mu, \sigma^2, (a_0, m_0), (a_1^A, m_1^A), (a_1^B, m_1^B), \dots, (a_{t-1}^A, m_{t-1}^A), (a_{t-1}^B, m_{t-1}^B)\}.$$

We model decentralized decision making as a simultaneous move game between Managers A and B . The benchmark comparison is to centralized decision making in which the managers are able to explicitly coordinate their actions or, alternatively, in which actions are formally coordinated by a central authority. In the centralized case the problem reduces to the single agent problem analyzed in earlier sections. If learning is decentralized, however, key differences emerge. As before, we begin with the first period and then turn to subsequent periods.

8.1 The First Period

Coordination failures can take two forms: agents may choose mismatched actions or they may match actions but coordinate on an inefficient choice. The assumption that managers face the

same production function implies that coordination failures are limited to the second type. If the agents were to choose different actions then a profitable deviation must exist: The manager with weakly lower expected utility benefits from deviating and imitating his colleague as this strictly lowers his coordination costs without worsening his direct utility from the mapping $m(\cdot)$.

Our interest, therefore, reduces to identifying actions that support coordination in equilibrium. It is immediate that the single-manager optimum supports an equilibrium in the decentralized firm as any deviation from such would lower direct utility and incur a cost of miscoordination. This property holds generally in every period. Therefore, our problem reduces further to the question of whether other, inefficient, actions can also support coordinated equilibrium behavior by the managers. For the first period the answer to this question is quite simply no. Surprisingly, the single-manager equilibrium is unique in the first period, and uniqueness holds regardless of the degree of complementarity between actions.

PROPOSITION 6. *The managers' optimal first period actions are unique for any δ and given by*

$$a_1^{A*} = a_1^{B*} = a_1^*,$$

where a_1^* is the optimal action in the baseline model defined in Proposition 1.

Decentralized decision making, therefore, has no impact on first period behavior, even when coordination costs are arbitrarily large. The managers coordinate on the action that a single manager would choose should he run the firm unilaterally.

To understand the result, consider the managers' problem should both agents locate at their status quo actions. If $m_0 \leq \hat{m}$ we know from the single-agent model that each manager has no incentive to change actions, and the complementarity only reinforces this incentive. If, instead, $m_0 > \hat{m}$, a single manager does wish to deviate as the marginal value of experimenting is positive. The same is true in the decentralized setting as the coordination cost of a marginal deviation is second order. Thus, Manager *A* possesses an incentive to unilaterally deviate and increase the size of his action, if only by a small amount. The resulting mismatch, however, cannot support an equilibrium as Manager *B* would then deviate to again match Manager *A*.

This process iterates at every possible action pair until the managers arrive at the optimal experiment with matching actions, at which point neither manager has an incentive to deviate. Leading the managers inexorably to the single-manager optimum – and driving uniqueness – is the global concavity of expected utility in the first period.

The uniqueness result in Proposition 6 is intuitive, yet it is important to note how dependent it is on the continuity of the action space. If the action space were discrete, with say a gap between

a_0 and a_1^* , then multiplicity of equilibria would exist for sufficiently large coordination costs. The managers would know they were better off at a_1^* than at a_0 , yet with the only option being a discontinuous jump in their action, and no way to explicitly coordinate actions, neither manager would be willing to deviate from a_0 and a coordination failure would result. This discreteness-induced failure is the classic source of coordination failures in the literature. Proposition 6 is notable, therefore, in that it does not fall prey to the same force.

8.2 The Second and Subsequent Periods

If the managers experiment in the first period they learn the outcome of a new action and in later periods know at least two points in the mapping. We show that the possession of a single additional piece of information is sufficient to generate coordination failures and inefficiency.

We begin, however, with a positive result. This result establishes that the accumulation of knowledge does not by itself guarantee a coordination problem and, thus, that organizations do not necessarily succumb to coordination failures. Avoiding this fate, however, is not easy and requires that performance improves, and improves non-trivially.

PROPOSITION 7. *The managers' optimal actions in period $t \geq 2$ are unique for any δ and given by*

$$a_t^{A*} = a_t^{B*} = a_t^* \text{ if } m_t^* \geq m_{t'-1}^* + \mu (a_{t'}^* - a_{t'-1}^*) \text{ for all } t' \leq t,$$

where a_t^* (and $a_{t'}^*$) is the optimal actions in the baseline model defined in Propositions 1 and 3 and $m_t^* = m(a_t^*)$ is the income realization in period t . If $m_t^* < m_{t'-1}^* + \mu (a_{t'}^* - a_{t'-1}^*)$ for any t' then the actions $a_t^{A*} = a_t^{B*} = a_t^*$ are still optimal but not necessarily uniquely so.

A sufficient condition to ensure the absence of coordination failures is for the firm's performance to match or exceed expectations. This is achieved with 50% probability in each period. Thus, uniqueness of equilibrium is not rare in the second period but is increasingly so as time passes. When it is achieved, however, a firm is immune from coordination mistakes.

Proposition 7 provides only a sufficient condition to avoid coordination problems and is silent on when and why they might emerge. To understand the possibilities, consider the second period after performance in the first period falls below expectations. Specifically, suppose that $m(a_1^*) = m_0$ and that first period performance exactly matched status quo performance. Clearly, action a_1^* cannot itself support an equilibrium as the marginal utility of further experimentation is positive and, following the logic of Proposition 6, the managers would iterate rightward along the action space until arriving at the single-manager optimum, a_2^* , and an equilibrium. The second period problem is not identical to the first, however, as the willingness of a manager to deviate from a_0 should the

other manager also locate there has changed. The logic for deviation in the first period rests upon the slope of the drift line being μ . But with knowledge of a second action, the expected gain is no longer μ , and instead it is zero. Middling performance in the first period—even though it matched previous performance—changes manager expectations over the intermediate actions. This change does not make these actions more appealing. In fact, the opposite holds, and the intermediate actions are less appealing. But the newly learned undesirability of these actions is exactly what disrupts the ability of managers to coordinate on the efficient action.

In particular, on the bridge between a_0 and a_1^* that we just described the expected gain in income is zero yet the variance is positive. Thus, if Manager B is located at a_0 , Manager A has no positive incentive to make a marginal deviation. The only viable deviation is for him to jump all the way to the right of a_1^* , yet such a jump creates a large mismatch between his action and that of the other manager. For large enough coordination costs this deviation is not profitable and the managers will find themselves stuck in an equilibrium at the status quo action a_0 .

The logic of this example establishes the possibility for coordination failures. Moreover, the logic does not depend on current performance exactly matching past performance and goes through even when past performance is exceeded. The next proposition describes the possibilities.

PROPOSITION 8. *The managers' optimal second period actions are not unique if δ is sufficiently large and $m_1^* \in (\widehat{m}, \overline{m}]$, where $\overline{m} = m_0 + \sigma^2 r(m_0)(a_1^* - a_0)/2$, a_1^* is the optimal action in the baseline model defined in Propositions 1, and $m_1^* = m(a_1^*)$ is the income realization in the first period. In particular, in addition to the actions characterized in Proposition 7, actions*

$$a_2^{A*} = a_2^{B*} = a_0$$

are optimal if $m_1^ \in [\widetilde{m}(m_0), \overline{m}]$, where $\widetilde{m}(m_0)$ is defined in (7). And actions*

$$a_2^{A*} = a_2^{B*} = a_1^* + \Delta(m_1^*) > a_1^*$$

are optimal if $m_1^ \in (\widehat{m}, \widetilde{m}(m_0)]$, where $\Delta(m_1^*)$ is the Δ that solves $R(m_1^*, \Delta) = 2\mu/\sigma^2$.*

This result contrasts with Proposition 7. In that result strong performance precluded coordination problems, whereas this result shows that middling performance guarantees them. The coordination failures can be of two types. The first type is intuitive. If first period performance matches or exceeds previous performance by a small margin, the managers can get stuck coordinating at the inefficient status quo action a_0 . The managers would prefer to jump to the single-manager optimum at a_2^* but they cannot coordinate on doing so which can leave them stranded at a_0 . This

equilibrium captures the intuition of Roberts (2004) that coordination problems cause firms to experiment less than efficiency demands.

Given the intuitiveness of this equilibrium, it is striking, therefore, that the inefficiency can also go the other way and that the managers can search too much. The previous result holds for all $m(a_1^*)$ above \hat{m} and up to $\overline{\overline{m}}$. For values just above \hat{m} but below m_0 , the single-manager optimum calls for experimentation to be abandoned and for the manager to return to the status quo action. In this case the equilibrium with experimentation is the *inefficient* equilibrium. Recall that the logic for backsliding is that while experimenting further is preferred to action a_1^* , it is not preferred to the certainty of the status quo action. The status quo action is, however, a long way from a_2^* , and the intermediate actions sufficiently unattractive, that with decentralized decision making the managers can succumb to a coordination failure and experiment where a single-agent wouldn't. This possibility runs counter to the intuition of Roberts as this coordination failure leads the managers to experiment *more* than is jointly optimal for them.

Propositions 7 and 8 provide possibility results for the absence and the existence of coordination problems, respectively. We left several cases open, namely for improved performance that exceeds $\overline{\overline{m}}$ but falls below expectations and for performance that is lower than \hat{m} . The proof of Proposition 7 can be used to show that a firm that nearly meets expectations is also immune from coordination problems, although a precise measure of the necessary performance depends on the particular functional form of utility. However, it is immediate from Proposition 8 that performance must at least exceed $\overline{\overline{m}}$ and, therefore, cannot fall too far below expectations. When performance falls below \hat{m} , multiplicity and coordination problems may result, both among known actions—this time a_0 or a_1^* —and for intermediate actions between these points. Multiplicity of equilibria may persist even for arbitrarily bad outcomes $m(a_1^*)$, implying that the costs of coordination failure can be unbounded. The precise condition necessary for this to hold is not particularly illuminating and for brevity we omit a formal statement of the result.

In combination, these results demonstrate how coordination problems are driven by firm performance. Firms that always match or exceed expectations never succumb to coordination failures. In contrast, firm with more haphazard growth—even if aggregate growth is better—encounter coordination problems. A firm with consistent growth, in effect, leaves little doubt as to the most promising route forward and removes the scope for coordination failures. A firm with uneven performance, in contrast, provides numerous local maxima that can lock-in managers. The potential multiplicity of equilibria is important for what it implies about aggregate growth and also for what it implies about why firms diverge in performance. We showed in Section 6 that, with a single manager,

small differences can persist and grow. The possibility of coordination failures with multiple managers goes further to explain why even otherwise identical firms may choose different actions and, consequently, evolve onto different development paths.

It is interesting to note the important role uncertainty plays in generating—or, more accurately, preventing—coordination failures. It is not a difficult task to specify a production function that produces coordination failures, particularly when the costs of miscoordination are high. Indeed, any production function that fails quasi-concavity, regardless of whether it is smooth or jagged, will exhibit the requisite local peaks to ensnare managers. The modeling challenge is to provide a positive theory as to why firms are sometimes ensnared on these local peaks, when it is expected to occur, and which peak it is that they get stuck at. Surprisingly, uncertainty plays a key role in providing answers to these questions. The realized path of a Brownian motion is assuredly non-monotonic, in fact with an uncountable number of peaks and troughs throughout the landscape. With perfect knowledge of the mapping, any of these peaks could trap decentralized managers and support an equilibrium. It is striking, therefore, that with *less* information about the mapping, coordination problems occur *less* frequently, and potentially not at all. It is here that uncertainty combines elegantly with the Brownian motion to generate expectations that are initially quasi-concave, yet in which local peaks appear endogenously as a function of the actions chosen and performance achieved. Notably, the smoothness and the local peaks that appear are solely in expectations and conceal numerous peaks and troughs underneath in the actual mapping. Yet to rational managers it is, of course, expectations that drive behavior. In that sense the complexity of the problem—and the uncertainty over the production function—is essential to our explanation of coordination failures. The smoothness generated by uncertainty—and the possibility of preserving quasi-concavity when performance matches expectations—provides a positive account of how it is that firms manage to avoid coordination failures and generate growth, yet at the same time explaining when and why coordination problems do emerge, and how the affliction of coordination varies by firm and performance.

9 Applications

The model we analyze is conceptually sparse. It omits many relevant forces so as to focus on the principal mechanisms of search and experimentation. This abstraction does, however, provide the indirect benefit of broad applicability. Our principal application is to the management of firms, where our results on performance differences, imitation, and coordination resonate with practical experience, as described in the introduction. A key factor in the search behavior of managers—

and, thus, in the performance of firms— is the level of risk aversion. Risk aversion in practice provides a real constraint on managers. It manifests as a ‘fear of failure’, a concept that has received considerable practical attention (see, for instance, Harford 2011) and is at the forefront of management at many of the most innovative organizations. The tension between risk aversion and experimentation in practice is neatly captured by Ed Catmull (2014, p.118), the president of Pixar and Disney Animation, who writes:

“While experimentation is scary to many, I would argue that we should be far more terrified of the opposite approach. Being too risk averse causes many companies to stop innovating and to reject new ideas, which is the first step on the path to irrelevance. Probably more companies hit the skids for this reason than because they dared to push boundaries and take risks—and, yes, to fail. To be a truly creative company, you must start things that might fail.”

Risk aversion and experimentation are relevant in other settings as well. As we mentioned in the introduction, one distant but nevertheless immediate and important application is to the development and growth of countries. Our results on inconsistent growth, failed learning, and performance divergence resonate with the historical evidence of economic growth. Several of these points of contact are of particular interest.

Poverty traps & stagnating growth. The optimal search rule we characterize presents two notable characteristics. First, that a performance trap exists, with low-performing firms choosing to not experiment at all. This cut-off point is consistent with the absence of growth in the poorest countries, often described as a poverty trap. A stark feature of our model is that poor and rich countries alike face the same opportunities for growth (in terms of the risk-return trade-off). Yet the fear of the downside is more daunting to a poor country than it is to a wealthier country, and it is this fear—the higher risk aversion among the poor—that deters experimentation and growth.

The second relevant property of the optimal search rule is that growth is not guaranteed even for agents with high income. A bad outcome can derail growth and cause managers to abandon the search for better performance. Stagnation of growth after initial growth was promising resonates with the experience of the many countries that begin to grow, only to encounter roadblocks that knock them off a higher growth trajectory and leave them seemingly unable to recover.

Convergence & imitation. Applied to growth and development, our result on the persistence of performance differences (Proposition 4) translates into an absence of convergence of wealth across countries. This is striking as convergence is a prominent prediction of classic models of growth, and its failure empirically has been long noted. The evidence does support a type of convergence known as “conditional convergence,” which says that countries with similar characteristics do tend

to converge over time. If we think of these similarities as the number of dimensions of choice that are common rather than country-specific, then a straightforward extension of our imitation result in Proposition 5 predicts that more similar countries should converge as they face less risk in imitating each other's successes than would dissimilar countries. As Matsuyama (1996, pp. 148) argues: *“The very diversity of the manners in which different developed economies cope with coordination also imposes a problem for underdeveloped economies if they try to learn from the experiences of more successful economies. They cannot pick and choose different parts from different systems, because of the complementarity inherent in any system. [...] Even if we could decide which system to adopt among all the systems currently known, and then replicate the system completely, it is not at all clear whether this is a desirable thing to do”*.

Coordination. The importance of coordination to growth and development has been a long-standing theme that dates back at least to the work of Rosenstein-Rodan (1943) and Hirschman (1958). Coordination drives Hirschman's theory of unbalanced growth and provides the rationale for a 'Big Push' in development (Murphy, Shleifer, and Vishny 1989). In all of these accounts coordination failure is driven by exogenous discreteness in the space of actions. Our Proposition 8 shows how discreteness is not necessary for coordination failures. Instead we demonstrate how coordination failures emerge endogenously over time as a function of strategy and performance. Coordination failures emerge, most commonly, when performance is poor but they can emerge also when performance is improving yet not as much as expected. This finding resonates with the historical growth experience. It implies that wealthy countries have not necessarily 'solved' their coordination problems, but rather it is their prior good performance that shields them from a multiplicity of equilibria. In contrast, the uneven performance of less developed countries not only damages growth directly, but indirectly it adds the additional burden of coordination problems that make future growth even more problematic.

Our model, quite clearly, does not capture the full richness of the development problem. The model's generality, that enables broad application, simultaneously limits its fit to any particular application. Nevertheless, the resonance to growth and development with such a stark model is suggestive and encouraging of closer examination.

10 Concluding Discussion

In this paper we use the Brownian motion framework to formally model the tension between risk aversion and the search for better performance. We seek to capture the process of experimentation,

why change sometimes succeeds and at other times fails, how managers learn from their experiments and grow their organizations. The Brownian motion possesses several attractive features in this regard, beyond its tractability. Substantively, the model not only captures the decision to experiment or not, it provides a notion of the direction of strategic change and a measure of the boldness, or novelty, of an experiment. The framework also captures the role of experience. In the model all experiments provide valuable information—be they successes, failures, or mediocrities—information that shapes beliefs and guides future choices. Experience accumulates over time leading to development trajectories that are path dependent. A realistic feature of the learning process—that is present neither in the rugged landscapes literature nor the classic bandits formulation—is that the sequence of choices is not blind, rather managers combine their experience with their theoretical understanding of the environment to rationally select the optimal next step in their search for better performance.

Many opportunities to generalize and enrich the model present themselves. That uncertainty is everywhere Gaussian allows us to integrate the Brownian framework with familiar classes of preferences. We chose to work with standard risk aversion (Kimball 1993), although our results extend readily to weaker requirements such as *proper risk aversion* (Pratt and Zeckhauser 1987), which requires only that undesirable independent risks are mutually aggravating.²

It is also possible to enrich the underlying environment, both via the knowledge held by managers and by the generating process. In our model the managers know the drift and variance terms. Thus, their theoretical knowledge of the underlying environment—their understanding of the world—is perfect and unchanging, and they lack only knowledge of the particular setting that they face. It is a straightforward extension to relax this assumption and suppose that managers learn about the underlying environment—about μ and σ^2 —as they learn the mapping. This extension adds realism and brings into play several interesting possibilities of substantive interest, in particular whether momentum effects will emerge in organizational learning and change.

It is also possible to employ stochastic processes other than the Brownian motion. For instance, discontinuities in the mapping via a jump process may more accurately capture organizations in which indivisibilities in production or organizational capacity are important. Other possibilities are to employ a mean-reverting Brownian motion if managerial innovations are more localized in their impact on performance, or a geometric Brownian motion if organizational scale is a component of the choice variable. A further possibility is to retain the Wiener diffusion process but allow the drift term to be non-linear.

²The only difference in results is in the uniqueness claims of Propositions 1 and 3 that may not hold on up to a countable set of points.

Other features of the model give rise to questions of substantive interest as well. Perhaps the most notable is how agents with a longer horizon experiment in our setting. Extending the agent’s planning horizon raises several issues—technical as well as substantive—although the underlying logic of our results appears to be robust. A fundamental insight of the experimentation literature is that the longer the horizon is, the more value there is in experimenting as any information gained will provide benefit longer into the future. This insight surely holds in our model, which then reignites the technical problem of whether an optimal experiment exists in each period. For myopic agents we have shown that a simple crossing condition on the level of absolute risk aversion is sufficient for existence. The equivalent condition for far-sighted agents remains an open and difficult question.³ The difficulty in establishing existence can be considerably eased, however, by varying the stochastic process, several possibilities of which were described above. For instance, if the drift term were convex and, in particular, bounded, an existence result should be attainable.

The substantive question raised by longer horizons is whether the search rule departs significantly from that found here. Much will not change. It is clear that searching to the right remains the optimal direction to search for better performance all else equal. And, for any discount factor less than one, it remains possible for learning to get stuck from a bad realization such that the agent backtracks and stops experimenting to the right. One possible substantive variation in behavior is that when the agent does decide to backtrack, he doesn’t revert to a known action but rather finds it optimal to experiment in the neighborhood of a known action. In any event, such behavior is consistent with the idea of learning stagnating following a bad realization, and it may actually reveal why firms for whom growth stagnates may nevertheless still find it possible to grow, albeit in a constricted region with limited upside.

An altogether different direction is to connect our optimizing model with ideas in the management literature on organizational adaptation, from which the rugged landscapes models evolved. The foundational idea in that area is that managers are boundedly rational. Intriguingly, Simon (1955) argues that the source of these constraints is the difficulty of the problem that managers face. Simon (1990, p.7) went so far as to refer to cognitive constraints and problem difficulty as two blades of scissors, neither of use without the other. Our model provides a formal representation of difficult problems, problems that are non-trivial even to decision makers who think clearly and optimize rationally. Our optimal search rule can be compared to boundedly rational rules-of-thumb that are offered in the management literature and their relative efficiency evaluated. In a setting where formulating beliefs and calculating optimal actions are cognitively taxing, our results can

³An interesting recent contribution by Garfagnini and Strulovici (2014) attacks this problem for agents with a two-period horizon, zero drift, and an exogenous search cost that is increasing in distance.

provide insight into why it is that certain search rules are selected and employed by managers.

The management problem is hard. Managers face a vast array of possible choices yet possess only a tenuous understanding of how these choices map into outcomes. It should come as little surprise, therefore, that firms fail spectacularly almost as often as they succeed, and that the performance of firms follow generally upward yet haphazard trajectories, creating differences in performance that persist over time. Our objective in this paper has been to move the spotlight of research onto the pure difficulty of the management problem. We sought to construct a parsimonious framework to capture this difficulty, abstracting away from many familiar informational, incentive, and selection effects that are obviously important in firms. We hope that with the framework in place, more progress on these important questions will follow.

11 Appendix

This appendix contains the proofs for all lemmas and for Propositions 1-3, which are the results that characterize the manager's optimal behavior and its implications in the baseline model. The proofs for Propositions 4-8, which cover the divergence result and the results in the sections on imitation and decentralized search, are in the online appendix.

Recall that

$$R(m, \Delta) \equiv -\frac{\mathbb{E}\left[u''\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]}{\mathbb{E}\left[u'\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]}$$

and that $R(m, 0)$ is equal to the coefficient of absolute risk aversion $r(m)$. For the proofs below it is convenient to also define

$$P(m, \Delta) \equiv -\frac{\mathbb{E}\left[u'''\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]}{\mathbb{E}\left[u''\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]} \tag{9}$$

and

$$T(m, \Delta) \equiv -\frac{\mathbb{E}\left[u''''\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]}{\mathbb{E}\left[u'''\left(m + \mu\Delta + \sqrt{\Delta}\sigma z\right)\right]} \tag{10}$$

Notice that $P(m, 0)$ is equal to the coefficient absolute prudence $-u'''(m)/u''(m)$ (Kimball 1990) and that $T(m, 0)$ is equal to the coefficient of absolute temperance $-u''''(m)/u'''(m)$ (Gollier and Pratt 1996). We can now prove the following lemma.

LEMMA A1. *For any $m \in \mathbb{R}$ and $\Delta \geq 0$ we have*

$$R(m, \Delta) \leq P(m, \Delta) \leq T(m, \Delta).$$

Moreover, the first inequality is strict if either $\Delta > 0$ or $\Delta = 0$ and $r'(m) < 0$, where $r(m)$ is the coefficient of absolute risk aversion.

Proof of Lemma A1: (i.) Suppose first that $\Delta = 0$. Differentiating the coefficient of absolute risk aversion, we get

$$r'(m) = R(m, 0)(R(m, 0) - P(m, 0)).$$

As we mentioned above, Kimball (1993) shows that standard risk aversion implies decreasing absolute risk aversion, that is, $r'(m) \leq 0$. It then follows from the above expression that $R(m, 0) \leq P(m, 0)$. Moreover, this inequality is strict if absolute risk aversion is strictly decreasing. Similarly, differentiating the coefficient of absolute prudence, we get

$$p'(m) = P(m, 0)(P(m, 0) - T(m, 0)).$$

Kimball (1993) shows that in a setting such as ours, standard risk aversion is equivalent to decreasing absolute prudence, that is, $p'(m) \leq 0$ for all $m \in \mathbb{R}$. It then follows from the above expression that $P(m, 0) \leq T(m, 0)$.

(ii.) Suppose now that $\Delta > 0$. Property 5 in Meyer (1987) shows that decreasing absolute risk aversion implies $R(m, \Delta) \leq P(m, \Delta)$ and Result 1 in Eichner and Wagener (2003) shows that decreasing absolute prudence implies $P(m, \Delta) \leq T(m, \Delta)$ (see also page 116 in Gollier 2001). All we need to do therefore is to show that our assumption that $r(m)$ crosses $2\mu/\sigma^2$, and thus $r'(m) < 0$ for some m , implies $R(m, \Delta) < P(m, \Delta)$. For this purpose, notice that $R(m, \Delta) < P(m, \Delta)$ is equivalent to

$$\mathbb{E} \left[u' \left(M + \sqrt{V}z \right) \right] \mathbb{E} \left[u''' \left(M + \sqrt{V}z \right) \right] - \mathbb{E} \left[u'' \left(M + \sqrt{V}z \right) \right]^2 > 0, \quad (11)$$

where $M = m + \mu\Delta$ and $V = \Delta\sigma^2$. We can rewrite this inequality as

$$\mathbb{E} \left[u' \left(M + \sqrt{V}z \right) \right] \mathbb{E} \left[u'' \left(M + \sqrt{V}z \right) z \right] - \mathbb{E} \left[u' \left(M + \sqrt{V}z \right) z \right] \mathbb{E} \left[u'' \left(M + \sqrt{V}z \right) \right] > 0, \quad (12)$$

where we use the facts that $\mathbb{E} \left[u' \left(M + \sqrt{V}z \right) z \right] = \sqrt{V} \mathbb{E} \left[u'' \left(M + \sqrt{V}z \right) \right]$ and $\mathbb{E} \left[u'' \left(M + \sqrt{V}z \right) z \right] = \sqrt{V} \mathbb{E} \left[u''' \left(M + \sqrt{V}z \right) \right]$. We can then follow the same argument as in the proof of Property 5 in Meyer (1987) to show that this inequality holds. In particular, let z^* satisfy

$$z^* \int_{-\infty}^{\infty} u' \left(M + \sqrt{V}z \right) dF(z) = \int_{-\infty}^{\infty} u' \left(M + \sqrt{V}z \right) z dF(z),$$

where $F(z)$ is the cumulative density function of the standard normal distribution. We can then rewrite the left-hand side of (12) as

$$\int_{-\infty}^{\infty} u' \left(M + \sqrt{V}z \right) dF(z) \int_{-\infty}^{\infty} r \left(M + \sqrt{V}z \right) u' \left(M + \sqrt{V}z \right) (z^* - z) dF(z), \quad (13)$$

where $r \left(M + \sqrt{V}z \right)$ is the coefficient of absolute risk aversion. The first integral is strictly positive. To sign the second integral, notice that

$$\int_{-\infty}^{\infty} u' \left(M + \sqrt{V}z \right) (z^* - z) dF(z) = 0$$

and that the integrand changes sign from positive to negative once. Since $r \left(M + \sqrt{V}z \right)$ is everywhere decreasing and strictly decreasing for at least some income levels, the second integral in (13) is strictly positive. The overall expression in (13) is therefore strictly positive. ■

Proof of Lemma 1: This lemma is proven in “Hilfsatz” 4.3 in Schneeweiss (1967) and Theorem 1 in Chipman (1973). ■

Proof of Lemma 2: As we mentioned above, Kimball (1993) shows that in a setting such as ours, standard risk aversion is equivalent to decreasing absolute prudence. The lemma then follows from Theorem 3 in Chipman (1973) and Theorem 2 in Lajeri and Nielsen (2000). ■

Proof of Proposition 1: The first-order condition for the manager’s problem is given by

$$\frac{dW(m_0 + \mu\Delta_1, \Delta_1\sigma^2)}{d\Delta_1} = E[u'(\cdot)] \frac{\sigma^2}{2} \left(\frac{2\mu}{\sigma^2} - R(m_0, \Delta_1) \right) \begin{cases} = 0 & \text{if } \Delta_1 > 0 \\ \leq 0 & \text{if } \Delta_1 = 0. \end{cases} \quad (14)$$

With this condition in mind, we now prove the optimal actions for an m_0 such that (i.) $m_0 < \hat{m}_l$, where \hat{m}_l denotes the smallest m such that $r(m) = 2\mu/\sigma^2$, (ii.) $m_0 > \hat{m}$, where \hat{m} denotes the largest m such that $r(m) = 2\mu/\sigma^2$, and (iii.) $m_0 \in [\hat{m}_l, \hat{m}]$. We then conclude the proof by performing the comparative static that is summarized in the proposition.

(i.) Optimal action for $m_0 < \hat{m}_l$: In this case, $dW(m_0, 0)/d\Delta_1 < 0$. Since expected utility $W(m_0 + \mu\Delta_1, \Delta_1\sigma^2)$ is concave in $\Delta_1 \geq 0$ it then follows that the status quo is the uniquely optimal action.

(ii.) Optimal action for $m_0 > \hat{m}_l$: In this case, $dW(m_0, 0)/d\Delta_1 > 0$. If an optimal action exists, it is therefore strictly to the right of the status quo.

To prove that an optimal action does exist, it is sufficient to show that there is a $\Delta_1 > 0$ such that $E[u(m_0 + \Delta_1 + \sigma\sqrt{\Delta_1}z)] < u(m_0)$. For this purpose, consider a status quo income level \underline{m} such that $r(\underline{m}) > 2\mu/\sigma^2$. For such an income level the manager strictly prefers the status quo to any action that is strictly to the right of the status quo. Now let $k(\Delta_1)$ denote the “compensating

premium” that would make the manager indifferent between, on the one hand, taking the status quo action and, on the other hand, receiving $k(\Delta_1)$ and taking an action that is a distance $\Delta_1 \geq 0$ to the right of the status quo. Formally, $k(\Delta_1)$ is given by the k that solves

$$\mathbb{E} \left[u \left(\underline{m} + k + \Delta_1 + \sigma \sqrt{\Delta_1} z \right) \right] = u(\underline{m}) \quad \text{for } \Delta_1 \geq 0,$$

where the Implicit Function Theorem ensures that $k(\Delta_1)$ exists. Implicitly differentiating this expression, we get

$$\frac{dk(\Delta_1)}{d\Delta_1} = \mu \frac{\sigma^2}{2} \left(R(\underline{m} + k, \Delta_1) - \frac{2\mu}{\sigma^2} \right).$$

Notice that this derivative is strictly positive for $\Delta_1 = 0$. Differentiating again we get

$$\begin{aligned} \frac{d^2k(\Delta_1)}{d\Delta_1^2} &= R(\underline{m} + k, \Delta_1) \left(\frac{\sigma^2}{2} \right)^2 * \\ &\quad \left[(R(\underline{m} + k, \Delta_1) - P(\underline{m} + k, \Delta_1))^2 + P(\underline{m} + k, \Delta_1) (T(\underline{m} + k, \Delta_1) - P(\underline{m} + k, \Delta_1)) \right]. \end{aligned}$$

Lemma A1 implies that this expression is positive for all $\Delta_1 \geq 0$.

Consider now any $m_0 > \hat{m}$. Since $k(\Delta_1)$ is strictly increasing and convex, there exists a $\Delta_1 > 0$ such that $m_0 = \underline{m} + k(\Delta_1)$. We then have

$$\mathbb{E} \left[u \left(m_0 + \Delta_1 + \sigma \sqrt{\Delta_1} z \right) \right] = u(\underline{m}) < u(m_0),$$

where the equality follows from the definition of $k(\Delta_1)$ and the inequality from the fact that $k(\Delta_1) > 0$. This implies that an optimal action exists.

Finally, we need to show that the optimal action is unique. For this purpose, we differentiate (14) to obtain

$$\frac{d^2W(\cdot, \cdot)}{d\Delta_1^2} = \mu^2 \left[\mathbb{E} [u''(\cdot)] + 2 \left(\frac{\sigma^2}{2\mu} \right) \mathbb{E} [u'''(\cdot)] + \left(\frac{\sigma^2}{2\mu} \right)^2 \mathbb{E} [u''''(\cdot)] \right], \quad (15)$$

where we used the facts that $\mathbb{E} [u'(\cdot) z] = \sigma \sqrt{\Delta_1} \mathbb{E} [u''(\cdot)]$ and $\mathbb{E} [u''(\cdot) z] = \sigma \sqrt{\Delta_1} \mathbb{E} [u'''(\cdot)]$. We can then use the definitions of $P(m_0, \Delta_1)$ and $T(m_0, \Delta_1)$ in (9) and (10) to rewrite this expression as

$$\begin{aligned} \frac{d^2W(\cdot, \cdot)}{d\Delta_1^2} &= -\mu^2 \mathbb{E} [u'(\cdot)] R(m_0, \Delta_1) \left(\frac{\sigma^2}{2\mu} \right)^2 * \\ &\quad \left[\left(\frac{2\mu}{\sigma^2} - P(m_0, \Delta_1) \right)^2 + P(m_0, \Delta_1) (T(m_0, \Delta_1) - P(m_0, \Delta_1)) \right]. \end{aligned} \quad (16)$$

Lemma A1 implies that this expression is negative for any $\Delta_1 \geq 0$. This confirms that expected utility is concave. Moreover, Lemma A1 implies that the above expression is strictly negative for any $\Delta_1 > 0$ that satisfies the first-order condition (14). This, in turn, implies that the optimal action is unique.

(iii.) Optimal action for $m_0 \in [\widehat{m}_l, \widehat{m}]$: In this case, $dW(m_0, 0)/d\Delta_1 = 0$. The status quo is therefore an optimal action. Moreover, it follows from Lemma A1 and (16) that the status quo is the unique optimum.

(iv.) Comparative statics for any $m_0 \geq \widehat{m}$: Above we showed that for any $m_0 \geq \widehat{m}$, the uniquely optimal action is given by $a_1^* = a_0 + \Delta(m_0)$, where $\Delta(m_0)$ is the $\Delta_1 \geq 0$ that solves

$$\frac{dW(m_0 + \mu\Delta_1, \Delta_1\sigma^2)}{d\Delta_1} = E[u'(\cdot)] \frac{\sigma^2}{2} \left(\frac{2\mu}{\sigma^2} - R(m_0, \Delta_1) \right) = 0. \quad (17)$$

Implicitly differentiating this expression we get

$$\frac{d\Delta(m_0)}{dm_0} = \frac{2\mu}{\sigma^2} \frac{\frac{2\mu}{\sigma^2} - P(m_0, \Delta_1)}{\left(\frac{2\mu}{\sigma^2} - P(m_0, \Delta_1)\right)^2 + P(m_0, \Delta_1)(T(m_0, \Delta_1) - P(m_0, \Delta_1))} > 0,$$

where the inequality follows from the first-order condition (17) and Lemma A1. ■

Proof of Lemma 3: We first show that there exists a unique $\tilde{m}(m_0) \in (\widehat{m}, m_0)$ such that

$$u(m_0) = E \left[u \left(\tilde{m}(m_0) + \mu\Delta(\tilde{m}(m_0)) + \sqrt{\Delta(\tilde{m}(m_0))}\sigma z \right) \right]. \quad (18)$$

For this purpose, notice that

$$E \left[u \left(\widehat{m} + \mu\Delta(\widehat{m}) + \sqrt{\Delta(\widehat{m})}\sigma z \right) \right] = u(\widehat{m}) < u(m_0)$$

and

$$E \left[u \left(m_0 + \mu\Delta(m_0) + \sqrt{\Delta(m_0)}\sigma z \right) \right] > u(m_0).$$

Expected utility $E \left[u \left(m + \mu\Delta(m) + \sqrt{\Delta(m)}\sigma z \right) \right]$ is therefore strictly less than $u(m_0)$ for $m = \widehat{m}$ and strictly larger than $u(m_0)$ for $m = m_0$. To show the existence of a unique $\tilde{m}(m_0)$ it is therefore sufficient to show that expected utility $E \left[u \left(m + \mu\Delta(m) + \sqrt{\Delta(m)}\sigma z \right) \right]$ is strictly increasing in $m \in [\widehat{m}, m_0]$. Applying the Envelope Theorem we obtain

$$\frac{dE \left[u \left(m + \mu\Delta(m) + \sqrt{\Delta(m)}\sigma z \right) \right]}{dm} = E \left[u' \left(m + \mu\Delta(m) + \sqrt{\Delta(m)}\sigma z \right) \right] > 0,$$

which completes the proof of the existence of a unique $\tilde{m}(m_0) \in (\widehat{m}, m_0)$.

To prove the comparative statics, we implicitly differentiate (18). Once again applying the Envelope Theorem, we have

$$\frac{d\tilde{m}(m_0)}{dm_0} = \frac{u'(m_0)}{\mathbb{E} \left[u' \left(\tilde{m} + \mu\Delta(\tilde{m}) + \sqrt{\Delta(\tilde{m})}\sigma z \right) \right]} > 0,$$

where the inequality follows from non-satiation.

To show that $d\tilde{m}/dm_0 \leq 1$ we need to establish that (18) implies

$$u'(m_0) \leq \mathbb{E} \left[u' \left(\tilde{m} + \mu\Delta(\tilde{m}) + \sqrt{\Delta(\tilde{m})}\sigma z \right) \right]$$

or, equivalently,

$$v(m_0) \geq \mathbb{E} \left[v \left(\tilde{m} + \mu\Delta(\tilde{m}) + \sqrt{\Delta(\tilde{m})}\sigma z \right) \right], \quad (19)$$

where we define $v(m_0)$ as the utility function $v(m_0) = -u'(m_0)$. Notice that (18) implies (19) if a manager with utility $v(\cdot)$ is more risk averse than a manager with utility function $u(\cdot)$. It is therefore sufficient to show that

$$-\frac{u'''(m)}{u''(m)} \geq -\frac{u''(m)}{u'(m)} \text{ for all } m \in \mathbb{R},$$

where the LHS is the coefficient of absolute risk aversion associated with $v(\cdot)$ and the RHS is the one associated with $u(\cdot)$. This inequality is satisfied since the utility function $u(\cdot)$ satisfies decreasing absolute risk aversion, which completes the proof. ■

Proof of Proposition 2: Follows immediately from the discussion in the text. ■

Proof of Proposition 3: Follows immediately from the discussion in the text. ■

References

- [1] BOLTON, PATRICK, AND CHRISTOPHER HARRIS. 1999. Strategic Experimentation. *Econometrica*, 67(2): 349-374.
- [2] BRYNJOLFSSON, ERIC AND PAUL MILGROM. 2013. Complementarity in Organizations. In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [3] CALLANDER, STEVEN. 2011. Searching and Learning by Trial and Error. *American Economic Review*, 101(6): 2277-2308.

- [4] CATMULL, ED. 2014. *Creativity, Inc.: Overcoming the Unseen Forces That Stand in the Way of True Inspiration*. Random House.
- [5] CHIPMAN, JOHN. 1973. The Ordering of Portfolios in Terms of Mean and Variance. *The Review of Economic Studies*, 40(2): 167-190.
- [6] CYERT, RICHARD, AND JAMES MARCH. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs.
- [7] EICHNER, THOMAS, AND ANDREAS WAGENER. 2003. More on Parametric Characterizations of Risk Aversion and Prudence. *Economic Theory*, 21(4): 895-900.
- [8] GARFAGNINI, UMBERTO, AND BRUNO STRULOVICI. 2014. Social Experimentation with Interdependent and Expanding Technologies. Working paper.
- [9] GARICANO, LUIS. 2000. Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, 108(5): 874-904.
- [10] ——— AND TIMOTHY VAN ZANDT. 2013. Hierarchies and the Division of Labor. In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [11] GIBBONS, ROBERT AND REBECCA HENDERSON. 2013. What Do Managers Do? In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [12] ———, NIKO MATOUSCHEK, AND JOHN ROBERTS. 2013. Decisions in Organizations? In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [13] GOLLIER, CHRISTIAN. 2001. *The Economics of Risk and Time*. MIT Press.
- [14] ——— AND JOHN PRATT. 1996. Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica*, 64(5): 1109-1123.
- [15] HANNAN, MICHAEL AND JOHN FREEMAN. 1977. The Population Ecology of Organizations. *American Journal of Sociology*, 82(5): 929-964
- [16] HARFORD, TIM. 2011. *Adapt: Why Success Always Starts with Failure*. Macmillan.

- [17] HIRSCHMAN, ALBERT. 1958. *The Strategy of Economic Development*. Yale University Press.
- [18] KIMBALL, MILES. 1990. Precautionary Saving in the Small and in the Large. *Econometrica*, 58(1): 53-73.
- [19] ———. 1993. Standard Risk Aversion. *Econometrica*, 61(3): 589-611.
- [20] LAJERI, FATMA AND LARS NIELSEN. 2000. Parametric Characterizations of Risk Aversion and Prudence. *Economic Theory*, 15(2): 469-476.
- [21] LEVINTHAL, DANIEL. 1997. Adaptation on Rugged landscapes. *Management Science*, 43(7): 934-950.
- [22] MATSUYAMA, KIMINORI. 1996. Economic Development as Coordination Problems. In *The Role of Government in East Asian Development: Comparative Institutional Analysis*, eds. Masahiko Aoki, Hyung-Ki Kim, and Masahiro Okuno-Fujiwara, Oxford University Press.
- [23] MEYER, JACK. 1987. Two-Moment Decision Models and Expected Utility Maximization. *American Economic Review*, 77(3): 421-430.
- [24] MILGROM, PAUL, AND JOHN ROBERTS. 1995. Complementarities and Fit Strategy, Structure, and Organizational Change in Manufacturing. *Journal of Accounting and Economics*, 19(2): 179-208.
- [25] MURPHY, KEVIN, ANDREI SHLEIFER, AND ROBERT VISHNY. 1989. Industrialization and the Big Push. *The Journal of Political Economy*, 97(5): 1003-1026.
- [26] PRATT, JOHN AND RICHARD ZECKHAUSER. 1987. Proper Risk Aversion. *Econometrica*, 55(1): 143-154.
- [27] RIVKIN, JAN. 2000. Imitation of Complex Strategies. *Management Science*, 46(6): 824-844.
- [28] ROBERTS, JOHN. 2004. *The Modern Firm*. Oxford University Press.
- [29] ROBERTS, JOHN AND GARTH SALONER. 2013. Strategy and Organization. In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [30] ROSENSTEIN-RODAN, PAUL. 1943. Problems of Industrialisation of Eastern and South-Eastern Europe. *The Economic Journal*, 53(210/211): 202-211.

- [31] SCHNEEWEISS, HANS. 1967. *Entscheidungskriterien bei Risiko*. Springer Verlag.
- [32] SIMON, HERBERT. 1955. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69(1): 99-118.
- [33] SIMON, HERBERT. 1990. Invariants of Human Behavior. *Annual Review of Psychology*, 41(1): 1-19.
- [34] SYVERSON, CHAD. 2011. What Determines Productivity?, *Journal of Economic Literature*, 49(2): 326-65.