A Simple Theory of Growth in a Complicated World

Steven Callander* Stanford University

Niko Matouschek† Northwestern University

PRELIMINARY AND INCOMPLETE

March 2013

Abstract

We examine a setting in which managers learn about the effectiveness of managerial practices by trial-and-error. We show that such a learning process generates persistent performance differences across firms. If learning is centralized, persistent performance differences arise if firms have different initial conditions. And if learning is decentralized and managerial practices are complementary, persistent performance differences can arise even if firms are identical. Moreover, if managerial practices are complementary, performance differences can persist even if managers can imitate each other. Finally, we explore the implications of our model for the growth and development of nations.

Keywords: learning, complementarities, complexity

JEL classifications: D21, D83, L25

*Graduate School of Business, Stanford University, sjc@GSB.Stanford.edu.

†Kellogg School of Management, Northwestern University, n-matouschek@kellogg.northwestern.edu.

1 Introduction

Firms that operate within the same industries, and even plants that operate within the same firms, often exhibit large and persistent performance differences. In recent years, researchers have made significant progress in exploring the implications of this fact for macroeconomics, industrial organization, labor markets, and trade (for a survey of the empirical evidence and implications see Syverson (2011)). An open question, however, is why performance differences exist in the first place. One view that has long been popular with managers and management scholars is that performance differences exist because different managers employ different managerial practices. In line with this view, a variety of recent studies have provided evidence for correlations and, in some cases, causal relationships between managerial practices and performance (see, for instance, Bloom et al. (2007, 2013) and, for a survey, Gibbons and Henderson (2012)).

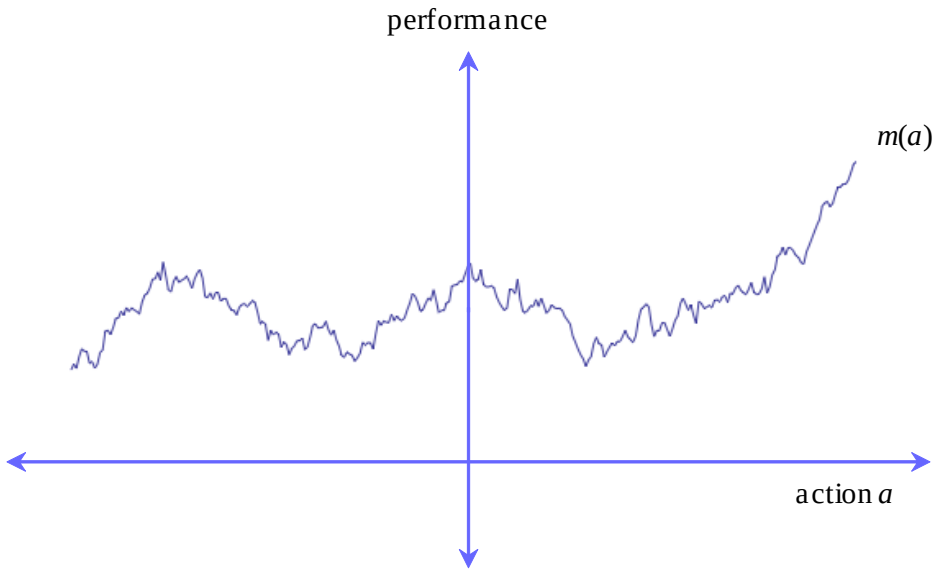
If managerial practices affect performance, however, then why do different managers employ different practices? And why do they continue to employ different practices even when it has become apparent that some work better than others? After all, the managerial practices of successful firms, such as Toyota's just-in-time production system and Lincoln Electric's organizational design, are often well-known and are not protected by patents.

The aim of this paper is to shed light on these questions. Our premise is that the problem of choosing between different managerial practices is too complex to be solved analytically. Instead, managers tackle this problem by trying out different managerial practices and observing how they affect performance. In other words, managers learn about managerial practices by trial-and-error. We show that such a learning process generates persistent performance differences. Moreover, if there are complementarities between managerial practices, these performance differences can persist, even when managers can imitate each other.

Managerial Learning To fix ideas, suppose that there is a single managerial practice that can take any value on the real line. There are infinitely many periods and in every period a manager has decide which specific practice or "action" to take. A production function maps the actions into the firm's performance. This function is increasing on average but also has many local peaks and troughs. The production function therefore looks like a rugged landscape that is tilted upwards, as illustrated in Figure 1.

Specifically, we follow Callander (2011) and model the production function as the realized path of a Brownian motion in which the action replaces time as the independent variable. The path is realized before the first period and then stays the same for all subsequent periods. The production

function therefore does not change over time. We assume that the drift of the Brownian motion is positive, which ensures that the production function is increasing on average. And we interpret the variance of the Brownian motion as a measure of the complexity of the production process. In particular, the more complex the production process is, the larger is the variance; and the larger is the variance, the more rugged is the production function.



The Production Function m(a)

In the first period, the manager knows the performance generated by a status quo action. And he knows the basic characteristics of the production function, as captured by the drift and the variance of the Brownian motion. Before the manager decides what action to take, he uses this information to form beliefs about the performance generated by any action that is different from the status quo.

The assumption that the production function is generated by a Brownian motion ensures that these beliefs take a simple form. In particular, the manager's beliefs are normally distributed with a mean that is increasing in the action and a variance that is increasing in both the distance between the action and the status quo and the complexity of the production process. The manager therefore knows that actions that are further to the right tend to generate better performance. And he knows more about an action, the closer it is to the status quo, or the less complex the production process is.

In this way the model captures the notion that managers know the general direction in which

to search for better managerial practices. They may know, for instance, that they should reduce inventory or tie pay more closely to performance. While managers know the general direction in which to search for better managerial practices, however, they do not know the effect of a specific change on performance. Moreover, they are more uncertain about the effect of a specific change, the larger is the change, and the more complex is the managerial practice.

After the manager has formed his beliefs in the first period, he either takes the status quo action, in which case he learns no new information. Or he experiments by taking a different action, in which case he learns the performance generated by that action. In the second, and any subsequent, period, the manager then updates his beliefs based on any new information he learned. And after he has updated his beliefs, he decides between taking one of the known actions and engaging in further experimentation. The manager can therefore learn about the production function over time by trying out different actions. The question then is what actions the manager should try out and when, if ever, he should stop and settle for one of the known actions.

We show that for a broad class of utility functions, and under the assumption that the manager maximizes expected utility on a period-by-period basis, the optimal learning rule takes a simple form. In particular, if status quo performance is below a threshold, the manager takes the status quo action in the first period, and he continues to do so in all subsequent periods. If, however, status quo performance is above the threshold, the manager starts to experiment in the first period by taking an action that is strictly to the right of the status quo. In subsequent periods, he then continues to experiment by taking actions that are further and further from the status quo until he reaches a period in which performance falls below another threshold. Once that happens, the manager reverts to the best known action and then continues to take that action in all subsequent periods.

Persistent Performances Differences A key implication of the optimal learning rule is that it generates persistent performance differences. To see this, consider the effect of an increase in status quo performance. We already noted that such an increase makes it more likely that the manager engages in experimentation in the first period. Below we also show that if the manager does engage in experimentation in the first period, an increase in status quo performance increases the size of his experiment and makes it more likely that he will engage in further experimentation in any subsequent period. An increase in status quo performance therefore unambiguously favors experimentation. And because an increase in status quo performance favors experimentation, it generates an even larger increase in expected performance in all subsequent periods. Differences in status quo performance therefore do not only persist, they actually grow larger.

This explanation for persistent performance differences relies on differences in status quo performances. Once we allow for multiple, complementary managerial practices, however, persistent performance differences can arise even if firms are identical. To see this, suppose that there is a second action. For simplicity, the production function for this second action is identical to the first. Each action is taken by a different manager who cares about the performance generated by his action minus the squared difference in the actions. The actions are therefore complementary.

In this setting, we can compare centralized learning–in which case the managers take their actions jointly–to decentralized learning–in which case they take them independently of each other. If learning is centralized, the optimal learning rule is the same as the one described above.

If learning is decentralized, however, there is a key difference. In particular, in any period other than the first one, the manager's actions are now determined by two thresholds. If performance is above the upper threshold, both managers engage experimentation. And if performance is below the lower threshold, both managers revert to the best known action. If performance is between the two thresholds, however, there are multiple equilibria: either both managers engage in experimentation or neither does. For intermediate levels of performance, the managers can therefore get stuck in a coordination failure. Moreover, the more complementary the actions, the more likely it is that coordination failures arise. Indeed, if the actions become very complementary, the only way to rule out coordination failures is for firm performance to improve continuously over time, without a single stumble or fall.

In summary, we show that persistent performances arise naturally if managers learn about managerial practices through trial-and-error. If learning is centralized, or if it is decentralized but managerial practices are independent, they arise if firms face different initial conditions. And if learning is decentralized and managerial practices are complementary, performance differences can arise even if firms are identical.

Barriers to Managerial Imitation So far we have abstracted from imitation. As we observed above, however, the managerial practices of successful firms are often well-known and are typically not protected by patents. The question then is why underperforming firms don't simply imitate the top performing ones. The answer that we develop below is that imitation may simply be too risky.

To see why imitation may be risky, suppose that some managerial practices are industry-specific–in the sense that their production functions are the same for all firms in the industry–while others are firm-specific. The mapping from managerial actions into corporate culture, for example, may be firm specific. This would be the case, for instance, if an American car manufacturer such as

General Motors has to take different actions to create a collaborative culture than a Japanese one such as Toyota.

If managerial practices are independent of each other, one would expect underperformers to imitate the top performers' industry-specific managerial practices. If, however, there are complementarities between industry- and firm-specific managerial practices, then imitating industry-specific managerial practices requires changes to firm-specific ones. Suppose, for instance, that the success of Toyota's just-in-time production system depends on its corporate culture and that, as we argued above, the production function that maps actions into corporate culture is firm-specific. To imitate Toyota's just-in-time production system, General Motors would then also need to change its corporate culture. But if managers learn about the production function for corporate culture by trial-and-error, then trying to change corporate culture is inherently risky.

Our final set of results characterizes conditions under which the risk associated with imitating industry-specific managerial practices is prohibitive. In particular, we show that if the top performers' industry-specific managerial practices are far from those employed by an underperforming firm, and the complementary, firm-specific managerial practices are sufficiently complex, then the underperforming firm is better off if it does not imitate the top performers.

We cast our model in the context of firms and industries. The main ingredients of our model–complexity, complementarities, and learning by trial-and-error–are also relevant in other settings. The vast literature on the growth and development of nations, in particular, touches on many of the issues that arise in our setting. We will discuss to this application of our model in the conclusion.

2 The Model

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

Technology: At the beginning of period $t = 1, 2, ...$ the manager the manager takes an action $a_t \in \mathbb{R}$ that determines his income level $m_t \in \mathbb{R}$. The mapping from actions into income is given by the production function $m(a_t)$, where $m : \mathbb{R} \to \mathbb{R}$. 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 shows one possible realization of the Brownian motion.

Preferences: The manager's utility is given by $u(m)$, where m is his income . The utility function is therefore the same for all agents. We assume that this function is four times continuously differentiable and satisfies

$$u'(m) > 0 \text{ and } u''(m) < 0 \text{ for all } m \in \mathbb{R}.$$

The first condition implies non-satiation and the second risk aversion.

To further characterize the manager's risk preferences, let

$$r(m_t) = -\frac{u''(m_t)}{u'(m_t)}$$

denote the coefficient of absolute risk aversion and let

$$p(m_t) = -\frac{u'''(m_t)}{u''(m_t)}$$

denote the coefficient of absolute prudence. We assume that absolute prudence is decreasing, that is, $p'(m_t) \leq 0$ for all m_t. Prudence was introduced by Kimball (1990) who shows that a risk averse agent engages in precautionary savings if and only if he is prudent, that is, $p(m_t) > 0$. Moreover, he shows that precautionary savings are decreasing in income if and only if absolute prudence is decreasing. Finally, he shows that decreasing absolute prudence implies decreasing absolute risk aversion (on unbounded domain). The manager's utility function therefore satisfies $r'(m_t) \leq 0$ for all m_t, where $r(m_t)$ denotes the coefficient of absolute risk aversion defined above.

Finally, we assume that the coefficient of absolute risk aversion crosses $2\mu/\sigma^2$, where μ and σ^2 are the drift and the variance of the Brownian motion. We will see below that if this were not the case, an equilibrium either does not exist if $r(m) < 2\mu/\sigma^2$ for all $m \in \mathbb{R}$, or it is trivial if $r(m) \geq 2\mu/\sigma^2$ for all $m \in \mathbb{R}$.

Information: The manager knows that the production function is generated by a Brownian motion with drift μ and variance σ^2 and that the status quo action a_0 generates income m_0. And he knows the income generated by any action he took in a previous period. The manager does not, however, know the realization of the Brownian motion. The manager's information set in period t is therefore given by $I_t = \{\mu, \sigma^2, (a_0, m_0), ... (a_{t-1}, m_{t-1})\}$.

Equilibrium: For simplicity we assume that the manager maximizes expected utility on a period-by-period basis. An equilibrium of the game is therefore given by $(a_1^*, a_2^*, ...)$, where

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

We now turn to characterizing the set of equilibria.

3 Managerial Learning

We first examine the optimal actions of the first two periods. We then show that the results about the optimal actions of the first two periods extend easily to the optimal actions in subsequent periods.

3.1 The First Period

In the first period, the manager can either take the status quo action, in which case he is certain to realize status quo income m_0. Or he can take some action $a_1 \neq a_0$, in which case he is uncertain about what income he will realize. Our first lemma characterizes the agent's beliefs in this case.

LEMMA 1. *If the manager takes a first period action $a_1 \in \mathbb{R}$, he believes that his income m_1 will be drawn from a normal distribution with mean*

$$\mathrm{E}\left[m_1\right] = m_0 + \mu(a_1 - a_0)$$

and variance

$$\mathrm{Var}\left(m_1\right) = \left|a_1 - a_0\right|\sigma^2,$$

where μ and σ^2 are the drift and the variance of the Brownian motion.

Notice that if $a_1 < a_0$, expected income is less than the status quo income m_0. Since he is risk averse, the manager will therefore never take an action that is to the left of the status quo. The agent may, however, take an action to the right of the status quo, since the expected income from taking such an action is larger than m_0. We therefore have the next lemma.

LEMMA 2. *If an optimal action a_1^* exists, it satisfies $a_1^* \in [a_0, \infty)$.*

Notice next that if $a_1 \geq a_0$, an increase in a_1 increases both expected income $\mathrm{E}\left[m_1\right]$ and the variance of income $\mathrm{Var}\left(m_1\right)$. The optimal choice of a_1 is therefore determined by a trade-off between risk and return. To characterize the optimal choice of a_1, suppose that $a_1 \geq a_0$ and

let $\Delta_1 = a_1 - a_0 \geq 0$ denote the size of the manager's first period experiment. Since income is normally distributed, we can write expected utility as

$$W(m_0, \Delta_1) = \mathrm{E}\left[u\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right], \tag{1}$$

where $m_0 + \mu\Delta_1$ and $\sqrt{\Delta_1}\sigma$ are the mean and the standard deviation of m_1 and where z is a random variable that is drawn from the standard normal distribution. The manager's problem is then given by

$$\max_{\Delta_1 \geq 0} W(m_0, \Delta_1).$$

Differentiating $W(m_0, \Delta_1)$ we obtain

$$W_{\Delta_1}(m_0, \Delta_1) = \mu\mathrm{E}\left[u'\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right] + \frac{\sigma^2}{2}\mathrm{E}\left[u''\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right], \tag{2}$$

where we used the fact that

$$\mathrm{E}\left[zu'\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right] = \sqrt{\Delta_1}\sigma\mathrm{E}\left[u''\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right].$$

It is useful to rewrite this derivative as

$$W_{\Delta_1}(m_0, \Delta_1) = \mathrm{E}\left[u'\left(m_0 + \mu\Delta_1 + \sqrt{\Delta_1}\sigma z\right)\right]\frac{\sigma^2}{2}\left(\frac{2\mu}{\sigma^2} - R(m_0, \Delta_1)\right), \tag{3}$$

where

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

Notice that $R(m_0, 0)$ is equal to the coefficient of absolute risk aversion $r(m_0)$. It then follows from (3) that the slope of expected utility at $\Delta_1 = 0$ is determined by the relative size of the coefficient of absolute risk aversion and the ratio $\left(2\mu/\sigma^2\right)$. Since the coefficient of absolute risk aversion is itself determined by the level of status quo income m_0, it is useful to introduce the following definition.

DEFINITION. *The income level $\widehat{m}$ denotes the largest income level m for which $r(m) = 2\mu/\sigma^2$.*

Note that the existence of $\widehat{m}$ is guaranteed by the assumption that absolute risk aversion crosses $2\mu/\sigma^2$. The sign of the slope of expected utility at $\Delta_1 = 0$ is therefore given by

$$W_{\Delta_1}(m_0, 0) \begin{cases} > 0 & \text{if } m_0 > \widehat{m} \\ \leq 0 & \text{if } m_0 \leq \widehat{m}. \end{cases} \tag{4}$$

Intuitively, the manager prefers a small experiment to the status quo if and only if he is sufficiently wealthy, in which case his coefficient of absolute risk aversion is sufficiently small.

To make further progress, we differentiate (3) to obtain

$$W_{\Delta_1\Delta_1}(m_0,\Delta_1) = E\left[u''(\cdot)\right] + 2\left(\frac{\sigma^2}{2}\right)E\left[u'''(\cdot)\right] + \left(\frac{\sigma^2}{2}\right)^2 E\left[u''''(\cdot)\right],$$

where we used the facts that

$$\mathrm{E}\left[zu''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right] = \sqrt{\Delta_1}\sigma\mathrm{E}\left[u'''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right]$$

and

$$\mathrm{E}\left[zu'''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right] = \sqrt{\Delta_1}\sigma\mathrm{E}\left[u''''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right].$$

We can rewrite the second derivative as

$$\begin{aligned} W_{\Delta_1\Delta_1}(m_0,\Delta_1) &= -\mathrm{E}\left[u'(\cdot)\right]R\left(\frac{\sigma^2}{2}\right)^2 * \\ &\quad * \left[\left(\frac{2}{\sigma^2} - P(m_0,\Delta_1)\right)^2 + P(m_0,\Delta_1)\left(Q(m_0,\Delta_1) - P(m_0,\Delta_1)\right)\right], \end{aligned} \tag{5}$$

where

$$P(m_0,\Delta_1) \equiv -\frac{\mathrm{E}\left[u'''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right]}{\mathrm{E}\left[u''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right]}$$

and

$$Q(m_0,\Delta_1) \equiv -\frac{\mathrm{E}\left[u''''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right]}{\mathrm{E}\left[u'''\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)\right]}.$$

Notice that $P(m_0,0)$ is equal to the coefficient of absolute prudence $p(m_0)$. Moreover, the next lemma shows that the relative size of $P(m_0,\Delta_1)$ and $Q(m_0,\Delta_1)$ is closely related to the notion of decreasing absolute prudence.

LEMMA 3. *The utility function $u\left(m_0+\mu\Delta_1+\sqrt{\Delta_1}\sigma z\right)$ exhibits decreasing absolute prudence if and only if $Q(m_0,\Delta_1) \geq P(m_0,\Delta_1)$ for all $m_0 \in \mathbb{R}$ and $\Delta_1 \geq 0$.*

This lemma follows directly from Chipman (1973) and Lajeri and Nielsen (2000). Together with (5), Lemma 3 implies that $W(m_0,\Delta_1)$ is concave, as stated in the next lemma.

LEMMA 4. *Expected utility is concave in Δ_1, that is, $W_{\Delta_1\Delta_1}(m_0,\Delta_1) \leq 0$ for all $m_0 \in \mathbb{R}$ and $\Delta_1 \geq 0$.*

We can now establish our first proposition which characterizes the manager's optimal action in the first period.

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

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

where $\Delta(m_0)$ *is implicitly defined by*

$$R(m_0, \Delta(m_0)) = \frac{2\mu}{\sigma^2}$$

and satisfies $\Delta(\widehat{m}) = 0$ *and* $\Delta'(m_0) > 0$ *for all* $m_0 \geq \widehat{m}$.

If $m_0 > \widehat{m}$, there therefore exists a unique optimal action that is strictly larger than the status quo. Moreover, the larger is status quo income, the larger is the optimal first period action and thus the optimal first period experiment. In the proof we show that existence and uniqueness of an optimal action follow from the assumptions that absolute prudence is decreasing and that absolute risk aversion crosses $\left(2\mu/\sigma^2\right)$. If, instead, $m_0 \leq \widehat{m}$, it is optimal for the manager to take the status quo action. This result follows immediately from the above observations that expected utility is concave and that for $m_0 \leq \widehat{m}$, the slope of expected utility is negative at $\Delta_1 = 0$.

To conclude this section, we examine the implications of relaxing the assumption that the coefficient of absolute risk aversion crosses $2\mu/\sigma^2$. It must then be the case that either $r(m) \geq 2\mu/\sigma^2$ for all $m \in \mathbb{R}$ or $r(m) < 2\mu/\sigma^2$ for all $m \in \mathbb{R}$. The next lemma characterizes the first agent's action in each case.

LEMMA 5. *If* $r(m) < 2\mu/\sigma^2$ *for all* $m \in \mathbb{R}$ *then an optimal action does not exist. And if* $r(m) \geq 2\mu/\sigma^2$ *for all* $m \in \mathbb{R}$ *then* $a_1^* = a_0$ *is an optimal action. Moreover,* $a_1^* = a_0$ *is the unique optimum unless status quo income* m_0 *satisfies* $r(m_0) = p(m_0) = 2\mu/\sigma^2$.

For the first part of the lemma, consider (2) and notice that $r(m) < 2\mu/\sigma^2$ for all $m \in \mathbb{R}$ implies $W_{\Delta_1}(m_0, \Delta_1) > 0$ for all $m_0 \in \mathbb{R}$ and for all $\Delta_1 \in \mathbb{R}$. In this case the agent therefore always benefits from taking a larger and larger action and, as a result, an optimal action does not exist.

For the second part of the lemma, suppose that $r(m) \geq 2\mu/\sigma^2$ for all $m \in \mathbb{R}$. It then follows from (4) that $W_{\Delta_1}(m_0, 0) \leq 0$ for all $m_0 \in \mathbb{R}$. Since $W(m_0, \Delta_1)$ is concave, it is immediate that $a_1^* = a_0$ is an optimal action. Moreover, taking the status quo action is uniquely optimal unless the manager's income level is given by an m_0 for which $r(m_0) = p(m_0) = 2\mu/\sigma^2$. The assumption that absolute risk aversion crosses $2\mu/\sigma^2$ therefore ensures the existence of a non-trivial solution to the manager's first period problem, as well as to the manager's problems in all subsequent periods.

.

3.2 The Second Period

We just saw that if $m_0 \leq \widehat{m}$, the manager takes the status quo action in the first period. In any subsequent period, the manager then faces the same period as in the first, as takes the status quo

action again. For the remainder of this paper, we therefore assume that

$$m_0 > \widehat{m},$$

in which case the manager's optimal first period action a_1^* is strictly larger than the status quo a_0. In the second period, the manager then knows the income associated with two actions: the status quo action a_0 and his first period action a_1^*. The next lemma characterizes the manager's second period beliefs if he takes an action that is different from either a_0 and a_1^*.

LEMMA 6. *If the manager takes a second period action $a_2 \notin \{a_0, a_1^*\}$, he believes that his income m_2 will be drawn from a normal distribution. If $a_2 \in (a_0, a_1^*)$, the mean and the variance are given by*

$$\mathrm{E}\left[m_2\right] = \frac{a_2 - a_0}{a_1^* - a_0} m_0 + \frac{a_1^* - a_2}{a_1^* - a_0} m_1^*$$

and

$$\mathrm{Var}\left(m_2\right) = \frac{\left(a_2 - a_0\right)\left(a_1^* - a_2\right)}{a_1^* - a_0} \sigma^2.$$

where $m_1^* \equiv m(a_1^*)$ denotes the income level associated with the optimal first period action a_1^* and where σ^2 is the variance of the Brownian motion. If, instead, $a_2 \in (-\infty, a_0) \cup (a_1^*, \infty)$, the mean and the variance are given by

$$\mathrm{E}\left[m_2\right] = \begin{cases} m_1^* + \mu(a_2 - a_1^*) & \text{if } a_2 \in (a_1^*, \infty) \\ m_0 + \mu(a_2 - a_0) & \text{if } a_2 \in (-\infty, a_0) \end{cases}$$

and

$$\mathrm{Var}\left(m_2\right) = \begin{cases} \left(a_2 - a_1^*\right)\sigma^2 & \text{if } a_2 \in (a_1^*, \infty) \\ \left(a_0 - a_2\right)\sigma^2 & \text{if } a_2 \in (-\infty, a_0)\,, \end{cases}$$

where μ is the drift of the Brownian motion.

Notice that for any action strictly between a_0 and a_1^*, expected income is a convex combination of the two known income levels m_0 and m_1^*. The manager therefore always prefers one of the two actions with known income to any action between them. Moreover, just as in the first period, the manager never finds it optimal to take an action to the left of the status quo. We therefore have the next lemma.

LEMMA 7. *If an optimal action a_2^* exists, it satisfies $a_1^* \in a_0 \cup [a_1^*, \infty)$.*

To characterize the optimal action, we proceed in two steps. First, we assume that the manager is constrained to take an action that is larger than the his first period action, that is, we assume $a_2 \geq a_1^*$. Second, once we know the constrained optimal action, we examine under what conditions the manager prefers that action to the status quo.

Suppose therefore that $a_2 \geq a_1^*$ and let $\Delta_2 = a_2 - a_1^* \geq 0$ denote the size of the manager's second period experiment. The manager's constrained problem is then given by

$$\max_{\Delta_2 \geq 0} W\left(m_1^*, \Delta_2\right),$$

where expected utility $W\left(m_1^*, \Delta_2\right)$ is defined in (1). Notice that this is the same problem the manager faced in the first period, except that the income level associated with not experimenting is now given by m_1^* rather than m_0. It then follows from Proposition 1 that the solution to the manager's constrained problem depends on whether m_1^* is smaller than $\widehat{m}$. In particular, if $m_1^* \leq \widehat{m}$, the solution is given by a_1^*. And if, instead, $m_1^* > \widehat{m}$, the solution is given by $a_1^* + \Delta^*\left(m_1^*\right)$, where $\Delta^*\left(m_1^*\right) > 0$ is implicitly defined by

$$R\left(m_1^*,\, \Delta^*\left(m_1^*\right)\right) = \frac{2\mu}{\sigma^2}. \tag{6}$$

Next we need to examine when the manager prefers the constrained optimal action to the status quo. For this purpose, consider the next lemma.

LEMMA 8. *There exists a unique income level $\widetilde{m}\left(m_0\right) \in \left(\widehat{m},\, m_0\right)$ such that*

$$u\left(m_0\right) = W\left(\widetilde{m}\left(m_0\right), \Delta^*\left(\widetilde{m}\left(m_0\right)\right)\right), \tag{7}$$

where $\Delta^*\left(\widetilde{m}\left(m_0\right)\right) > 0$. The derivative of this income level satisfies

$$0 < \frac{\mathrm{d}\widetilde{m}\left(m_0\right)}{\mathrm{d}m_0} < 1.$$

If first period income m_1^* is equal to $\widetilde{m}\left(m_0\right)$, the manager is indifferent between the status quo and the constrained optimal action. It then follows that the manager strictly prefers the status quo to the constrained optimal action if $m_1^* < \widetilde{m}\left(m_0\right)$ and that he strictly prefers the constrained optimal action to the status quo if $m_1^* > \widetilde{m}\left(m_0\right)$. We therefore have the next proposition.

PROPOSITION 2. *The manager's optimal second period action is unique and given by*

$$a_2^* = \begin{cases} a_1^* + \Delta^*\left(m_1^*\right) & \text{if } m_1^* > \widetilde{m}\left(m_0\right) \\ a_0 & \text{otherwise,} \end{cases}$$

where $\Delta^*(m_1^*) > 0$ *is defined in* (6).

Notice that if $m_1^* \in (\widehat{m},\ \widetilde{m}(m_0))$, the manager reverts to the status quo even though the marginal return from engaging in further experimentation is positive. The reason is that while in this case experimentation is better than m_1^*, it is not better than the status quo $m_0 > m_1^*$.

3.3 Subsequent Periods

The problem the manager faces in periods $t > 2$ is very similar to the one he faces in period 2. To characterize his optimal actions, we denote the largest income level known in period t by $m_t^{\max}$, that is, we define

$$m_t^{\max} = \max\{m_0, m_1^*, m_2^*, ..., m_{t-1}^*\}.$$

And we denote the action that generated $m_t^{\max}$ by $a(m_t^{\max})$, that is, we define

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

And finally, we denote the largest action the manager took before period t by $a_t^{\max}$, that is, we define

$$a_t^{\max} = \max\left\{a_0, a_1^*, a_2^*, ..., a_{t-1}^*\right\}.$$

We can then characterize the optimal actions in periods $t \geq 2$ as follows.

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

$$a_t^* = \begin{cases} a_t^{\max} + \Delta^*(m(a_t^{\max})) & \text{if } m(a_t^{\max}) > \widetilde{m}(m_t^{\max}) \\ a(m_t^{\max}) & \text{otherwise,} \end{cases}$$

where $\Delta^*(m(a_t^{\max})) > 0$ *is defined in* (6) *and* $\widetilde{m}(m_t^{\max})$ *is defined in* (7).

In any period $t \geq 2$, the manager therefore engages in experimentation if and only if the income level $m(a_t^{\max})$ associated with the largest previously taken action $a_t^{\max}$ is above a threshold $\widetilde{m}(m_t^{\max})$, where the threshold is increasing in the largest known income level $m_t^{\max}$.

3.4 Discussion

As discussed in the Introduction, the optimal learning rule therefore takes a simple form. In particular, if status quo income is below a threshold, the manager takes the status quo action in the first period, and he continues to do so in all subsequent periods. If, however, status quo income is above the threshold, the manager starts to experiment in the first period by taking an action that is strictly larger than the status quo. In subsequent periods, he then continues to experiment

by taking larger and larger actions until he reaches a period in which income falls below another threshold. Once that happens, the manager reverts to the best known action and then continues to take that action in all subsequent periods.

Figuratively, the manager's exploration of the rugged landscape depends on the height of his starting point. If that point is too low, he just stays put. If his starting point is sufficiently high, however, he starts exploring the rugged landscape by taking discrete steps towards the right. The manager continues his rightward march until he falls off a sufficiently large cliff. At that point, continuing his march is too risky for the manager, who instead returns to the highest peak he discovered during his exploration.

Notice that even though the manager eventually settles for an action that guarantees him some level of income, he is not satisficing. In particular, the manager does not stop searching for better actions once he found one that is "good enough." Instead, he stops searching for better actions after having taken a sufficiently bad one. The manager then reverts to the best known action because further experimentation is too risky. Notice also that when the manager does settle for an action, that action is a local peak *given the manager's information*. Since the manager's information is coarse, however, his action will not, in general, be an actual local peak.

A key implication of the model is that it generates persistent performance differences. We discuss this implication in the next section.

4 Persistent Performance Differences

In this section we first show that persistent performance differences arise if firms have different status quo income levels. We then show that if there are multiple, complementary managerial practices and learning is decentralized, persistent performance differences can arise even if firms are identical.

4.1 Due to Differences in Status Quo Income Levels

To understand the effect of an increase in status quo income, recall that the size of the optimal first period experiment is increasing in status quo income. An increase in status quo income therefore generates higher expected income in the first period, both because the starting point from which the manager experiments is higher, and because his experiment is larger. A larger experiment, however, also generates more uncertain income. And higher status quo income raises the threshold below which the manager finds it optimal to revert to the status quo in the second period. It is therefore not immediately obvious whether an increase in status quo income makes it more or less

likely that the manager will engage in further experimentation in the second period. Specifically, it follows from Proposition 2 that the probability with which the manager engages in further experimentation in the second period is given by

$$prob\left(m_1^* > \widetilde{m}\left(m_0\right)\right) = F\left(\frac{m_0 + \mu\Delta^*\left(m_0\right) - \widetilde{m}\left(m_0\right)}{\sqrt{\Delta^*\left(m_0\right)}\sigma}\right),$$

where $F\left(\cdot\right)$ is the cumulative density function of the standard normal distribution. Notice that an increase in m_0 increases both the numerator–the difference between expected first period income $m_0 + \Delta^*\left(m_0\right)$ and the threshold level of income $\widetilde{m}\left(m_0\right)$ above which the second agent engages in further experimentation–and the denominator–the standard deviation of first period income$\sqrt{\Delta^*\left(m_0\right)}\sigma$. The next lemma shows that the effect on the numerator dominates the effect on the denominator. And it shows that this is true for all subsequent periods.

LEMMA 9. *The probability that the manager engages in further experimentation in period $t \geq 2$ is strictly increasing in status quo income, that is,*

$$\frac{\mathrm{d}prob\left(m_t^* > \widetilde{m}\left(m_t^{\max}\right)\right)}{\mathrm{d}m_0} > 0.$$

An increase in status quo income therefore favors experimentation unambiguously: it makes it more likely that the manager engages in experimentation in the first period, it increases the size of the first period experiment if the manager does engage in experimentation, and it makes it more likely that the manager will engage in further experimentation in any subsequent period. As a result, an increase in status quo income generates an even larger increase in expected income in any subsequent period, as shown in the next proposition.

PROPOSITION 4. *Expected income in any period $t \geq 2$ is strictly increasing in status quo income, that is,*

$$\frac{\mathrm{dE}\left[m_t^*(m_0)\right]}{\mathrm{d}m_0} > 0.$$

To interpret the lemma, suppose there are two firms that differ in their status quo income. The two firms use the same production function. Their managers, however, cannot observe each other and thus learn about the production function independently. The lemma implies that the difference in expected income in any period $t \geq 1$ is strictly larger than the difference in status quo income. Differences in status quo income therefore do not only persist, they actually grow larger.

4.2 Due to Complementarities and Decentralized Learning

The explanation for persistent performance differences that we just discussed relies on differences in status quo income. In this section we show that if there are multiple, complementary managerial

practices, persistent performance differences can arise even if firms are identical.

For this purpose, suppose there are two managers, A and B. In any period $t = 1, 2, ...$ the managers make decisions $a_t^A \in \mathbb{R}$ and $a_t^B \in \mathbb{R}$ and then realize their incomes m_t^A and m_t^B. In particular, their incomes are given by

$$m_t^A = m\left(a_t^A\right) - \frac{1}{2}\delta\left(a_t^A - a_t^B\right)^2$$

and

$$m_t^B = m\left(b_t\right) - \frac{1}{2}\delta\left(a_t^A - a_t^B\right)^2,$$

where $\delta \geq 0$ is a parameter that measures the importance of coordination between the managers' decisions. The function $m(\cdot)$ is once again the realized path of a Brownian motion with drift μ and variance σ^2 and for which $m_0 = m(0)$. Notice that this function is the same for both managers. The managers are therefore learning about the same production function.

The managers know that the production function is generated by a Brownian motion with drift μ and variance σ^2 and that $m(0) = 0$. Moreover, in any period t, the managers know the actions that they took in previous periods and the income that these actions generated. Managers A and B therefore have the same information set

$$I_t = \left\{\mu, \sigma^2, (0, m_0), \left(a_1^A, m_1^A\right), \left(a_1^B, m_1^B\right), ..., \left(a_{t-1}^A, m_{t-1}^A\right), \left(a_{t-1}^B, m_{t-1}^B\right)\right\}.$$

And they have the same utility function $u(\cdot)$ that satisfies the same properties described above.

We distinguish between centralized learning–in which case the managers coordinate their actions–and decentralized learning–in which case the managers take their actions independently and simultaneously. If learning is centralized, the optimal learning rule is the same as the one we derived above. If learning is decentralized, however, there is a key difference, which we derive below. Since the formal analysis of this version of the model is similar to the single-action version we described above, we relegate the formal analysis to the appendix.

4.3 The First Period

The managers never find it optimal to take actions to the left of the status quo, just as in the single-action model. Suppose therefore that $a_1 \geq a_0^A$ and $b_1 \geq a_0^B$ and let $\Delta_1^A = a_1^A - a_0^A$ and $\Delta_1^B = a_1^B - a_0^B$ denote the size of each manager's experiment. We can then write manager A's expected utilities as

$$W^A\left(m_0, \Delta_1^A, \Delta_1^B\right) = E\left[u\left(m_0 + \mu\Delta_1^A + \sigma\sqrt{\Delta_1^A}z - \frac{1}{2}\delta\left(\Delta_1^A - \Delta_1^B\right)^2\right)\right],$$

where z is again a random variable drawn from the standard normal distribution. The definition of manager B's expected utility is analogous. The managers' first period problem is then given by

$$\max_{\Delta_1^A \geq 0} W^A\left(m_0, \Delta_1^A, \Delta_1^B\right)$$

and

$$\max_{\Delta_1^B \geq 0} W^B\left(m_0, \Delta_1^B, \Delta_1^A\right).$$

The solution to this problem gives the managers' optimal first period actions, which we characterize in the next proposition.

PROPOSITION 4. *The managers' optimal first period actions are unique and given by*

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

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

The optimal first period actions are therefore the same as in the single-action model. Notice that this implies that there cannot be any coordination failures in the first period, even if the actions are very complementary, that is, even if δ is very large.

To understand why coordination failures cannot arise in the first period, suppose that δ is very large and that $\Delta_1^A = \Delta_1^B = 0$. Notice that if $\Delta_1^A = \Delta_1^B$, the effect of a marginal increase in the size of a manager's experiment on his or her utility is second order. As long as $m_0 > \widehat{m}$, manager A has an incentive to unilaterally increase the size of his investment by a small amount. But once A has increased his experiment, manager B benefits from increasing her investment by the same amount. And once B has increased her investment by the same amount, A benefits by unilaterally increasing her experiment by a further small amount. This process continues until each manager's experiment is equal to the optimal experiment in the single-action model, at which point neither manager has an incentive to unilaterally change his or her experiment.

4.4 The Second And Subsequent Periods

Consider first the second period. As in the single-action model, it is never optimal to take an action that is either smaller than the status quo or lies strictly between the status quo and the optimal first period action a_1^*. And as in the first period of the two-action model, the two managers will always take the same action. In contrast to either setting, however, the managers may now get stuck in a coordination failure.

To understand why coordination failures can arise in the second period, suppose that m_1^* is just marginally larger than the threshold income level $\widetilde{m}(m_0)$. We saw above that if $m_1^* > \widetilde{m}(m_0)$,

then in the single-action model the optimal second period action is given by $a_1^* + \Delta^*(m_1^*) > a_1^*$, where $\Delta^*(m_1^*)$ is defined in (6). Suppose now that in the second period manager A takes action $a_1^* + \Delta^*(m_1^*)$. Since the managers benefit from coordination, it is then optimal for manager B to take the same action. It is therefore an equilibrium for each manager to take action $a_1^* + \Delta^*(m_1^*)$. Suppose now, however, that manager A takes the status quo action a_0. Since the managers benefit from coordination, it is then optimal for manager B to take the same action. In particular, this will be the case if δ is large enough. It is therefore also an equilibrium for each manager to take the status quo action, in which case the managers are stuck in a coordination failure.

To understand the conditions under which coordination failures can arise in the second period, suppose that $m_1^* > \widetilde{m}(m_0)$ and let $\Delta^{k*}(m_1^*, 0)$ denote the solution to

$$\max_{\Delta_2^k \geq 0} W^k\left(m_1^*, \Delta_2^k, 0\right) \text{ for } k = A, B.$$

In the appendix we show that there then exists a unique income level $\widetilde{m}(m_0, \delta)$ such that

$$u(m_0) = W^k(\widetilde{m}(m_0, \delta), \Delta^{k*}(\widetilde{m}(m_0, \delta), 0), 0). \tag{8}$$

In words, there exists a unique income level such that if one manager takes the status quo action, the other is indifferent between also taking the status quo action and engaging in experimentation. Notice that $\widetilde{m}(m_0, 0) = \widetilde{m}(m_0)$. Moreover, we show in the appendix that while $\widetilde{m}(m_0, \delta)$ is strictly increasing in δ, it is always strictly smaller than m_0. For any $\delta > 0$ there therefore exists a region $(\widetilde{m}(m_0, \delta), m_0)$ such that if m_1^* is within that region, there are two equilibria, one in which both managers take the status quo action and another in which both managers take action $a_1^* + \Delta^*(m_1^*)$, which is the optimal second period action in the single-action model.

As in the single-action model, the problem faced by the managers in periods $t \geq 2$ is very similar to the first period problem. To characterize the optimal actions in periods $t \geq 2$, notice first that it is never optimal for managers A and B to take different actions. We can therefore once again use $m_t^{\max}$ to denote the largest income level known in period t. Similarly, we can once again let $a(m_t^{\max})$ denote the action associated with $m_t^{\max}$. And finally, we can once again let $a_t^{\max}$ denote the largest action that was taken before period t. We then have the following proposition.

PROPOSITION 5. *The managers' optimal actions in period $t \geq 2$ are given by*

$$a_t^{A*} = a_t^{B*} = a_t^* = \begin{cases} a_t^{\max} + \Delta^*(m(a_t^{\max})) & \text{if } m(a_t^{\max}) > \widetilde{m}(m_t^{\max}, 0) \\ a(m_t^{\max}) & \text{if } m(a_t^{\max}) \leq \widetilde{m}(m_t^{\max}, \delta) \end{cases}$$

where $\Delta^*(m(a_t^{\max})) > 0$ and $\widetilde{m}(m_t^{\max}, \delta)$ are defined in (6) and (8) and where $\widetilde{m}(m_t^{\max}, \delta)$ is strictly increasing in δ and satisfies $\widetilde{m}(m_t^{\max}, \delta) < m_t^{\max}$ for any $\delta \geq 0$.

In any period t, multiple equilibria therefore arise if the income level associated with the largest action $m\left(a_t^{\max}\right)$ lies strictly between $\widetilde{m}\left(m_t^{\max}, 0\right)$ and $\widetilde{m}\left(m_t^{\max}, \delta\right)$. We already noted that the critical income level $\widetilde{m}\left(m_t^{\max}, \delta\right)$ is strictly increasing in δ. In the appendix we further show that as $\delta \to \infty$, $\widetilde{m}\left(m_t^{\max}, \delta\right) \to m_t^{\max}$. This implies that as coordination becomes very important, coordination failures become pervasive. In particular, the only way to avoid coordination failures for sure is for each period's income level to be larger than the previous period's income level.

5 Barriers to Imitation

So far we have abstracted from imitation. As we observed in the Introduction, however, the managerial practices of successful firms are often observable and are typically not protected by patents. This raises the question of why underperformers do not imitate the managerial practices of top performers.

To address this question, suppose now that there is a single manager who, in every period t, has to take two actions a_t^A and a_t^B. These actions are highly complementary. In fact, they are so complementary that the manager always needs to set them equal to each other, that is, $a_t^A = a_t^B = a_t$. The manager's income is then given by

$$m_t = m^A\left(a_t\right) + m^B\left(a_t\right),$$

where $m^A\left(a_t^A\right)$ and $m^B\left(a_t^B\right)$ are the production functions for the two actions. Each production function is an independently realized path of a Brownian motion with drift μ and variance σ^2. The status quo actions are given by $a_0^A = a_0^B = a_0 = 0$ and each status quo action generates status quo income $m^A(0) = m^B(0) = m_0/2$, where we divide by two for notational convenience. The rest of the model is as in the single-action model.

In the first period, the manager's expected utility is then given by

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

where $\Delta_1 = a_1 - a_0$ denotes the size of the manager's experiment for each action. Notice that this expression is equivalent to the manager's first period expected utility in the single-action model if the drift were given by 2μ and the variance were given by $2\sigma^2$. We can therefore use Proposition 1 to determine the manager's optimal decisions in the first period.

To focus on imitation, suppose that in the absence of any other information, the manager finds it optimal not to engage in experimentation in the first period, that is, suppose that $m_0 < \widehat{m}$, where $\widehat{m}$ is defined in Section 3.1. Moreover, suppose that there is a second firm that shares

the same production function for action A. The production function for this action is therefore industry-specific. In contrast, the production function for action B is firm-specific, in the sense that for each firm the production function is an independent realization of the Brownian motion with the same drift, variance, and status quo action described above. The manager of the first firm knows that the actions of the second firm are given by $b^A = b^B = d > 0$, where d stands for "distance" from the first firm. Moreover, the manager knows that $m^A(d) = M + m_0/2$ for $M > 0$.

We are interested in the conditions under which the manager prefers the status quo to imitation, that is, the conditions under which he prefers action $a_1 = 0$ to $a_1 = d$. To determine these conditions, notice that the expected utility from imitation is given by

$$W(m_0 + M, d) = E\left[u\left(m_0 + M + \mu d + \sigma\sqrt{d}z\right)\right].$$

The manager therefore prefers the status quo to imitation if and only if $u(m_0) \geq W(m_0 + M, d)$.

We can now use the results from the single-action model to characterize $W(m_0 + M, d)$. In particular, it follows from (4) that $W_d(m_0 + M, d) > 0$ if and only if $m_0 + M > \widehat{m}$. Moreover, it follows from Lemma 4 that $W(m_0 + M, d)$ is concave in d. And finally, it follows from Proposition 1 that if $m_0 + M > \widehat{m}$ there exists a unique $d > 0$ at which $W(m_0 + M, d)$ is maximized. Together these properties imply that there exists a critical distance $\widehat{d} > 0$ such that the manager prefers the status quo to imitation if and only if $d \geq \widehat{d}$. The manager therefore prefers the status quo to imitation if the other firm's managerial practices are too far its own, in which case imitation is simply too risky. The following proposition summarizes this result and shows how the critical distance varies with the economic environment.

PROPOSITION 6. *The manager prefers the status quo to imitation if and only if*

$$d \geq \widehat{d},$$

where $\widehat{d} > 0$ *is increasing in* M *and decreasing in* σ^2.

The manager therefore prefers the status quo to imitation if the other firm is too far away in terms of its actions, not in terms of its performance. Indeed, the further ahead the other firm is in terms of performance, that is, the larger M is, the more likely it is that the manager prefers imitation to the status quo. Notice also that the barrier to imitation is the complexity of the firm specific managerial practice. The less complex this managerial practice is, that is, the smaller σ^2 is, the more likely it is that the manager prefers imitation to the status quo. And if $\sigma^2 = 0$, the manager will always find it optimal to imitate that other firm.