

## ADAPTIVE CAPACITY TO TECHNOLOGICAL CHANGE: A MICROFOUNDATIONAL APPROACH

VIKAS A. AGGARWAL,<sup>1\*</sup> HART E. POSEN,<sup>2</sup> and MACIEJ WORKIEWICZ<sup>3</sup>

<sup>1</sup> INSEAD, Fontainebleau, France

<sup>2</sup> University of Wisconsin-Madison, Madison, Wisconsin, U.S.A.

<sup>3</sup> ESSEC Business School, Cergy-Pontoise, France

**Research summary:** We take a microfoundational approach to understanding the origin of heterogeneity in firms' capacity to adapt to technological change. We develop a computational model of individual-level learning in an organizational setting characterized by interdependence and ambiguity. The model leads to organizational outcomes with the canonical properties of routines: constancy, efficacy, and organizational memory. At the same time, the process generating these outcomes also produces heterogeneity in firms' adaptive capacity to different types of technological change. An implication is that exploration policy in the formative period of routine development can influence a firm's capacity to adapt to change in maturity. This points to a host of strategic trade-offs, not only between performance and adaptive capacity, but also between adaptive capacities to different forms of change.

**Managerial summary:** Why are firms differentially effective at adapting to technological change? We argue that firms differ in the adaptive capacity of the routines that underlie their capabilities. These differences arise well before change occurs, and result because firms build routines that are differentially responsive to signals of performance decline associated with technological change. Thus, early managerial efforts to build superior productive efficiency must be complemented by efforts to build superior adaptive capacity. Our theory suggests that managers can prepare for technological change by implementing policies, in the formative period of organizational development, that promote individuals' exploration of novel actions. However, there are trade-offs because preparation aimed at building adaptive capacity to one type of technological change may limit adaptive capacity to other types of change. Copyright © 2016 John Wiley & Sons, Ltd.

### INTRODUCTION

The development of the strategy literature has been strongly influenced by two key themes: the drivers of inter-firm performance heterogeneity, and the factors underlying heterogeneity in the capacity of firms to adapt to technological change. These themes have evolved somewhat independently. The

former focuses on endowments of unique resources and capabilities (e.g., Barney, 1991; Peteraf, 1993; Winter, 2000), while the latter focuses on factors such as complementary assets, organizational interdependencies, and the nature of the change itself (e.g., Henderson and Clark, 1990; Taylor and Helfat, 2009; Tripsas, 1997; Tushman and Anderson, 1986). In this article, we seek to link these themes theoretically. We propose that the same process underlying the emergence of routines, which ultimately serve as building blocks of firm capabilities, also endows firms with heterogeneity in the capacity to adapt to technological change. We suggest, moreover, that a manager's strategic choices during the formative period of a routine's

Keywords: microfoundations; routines; adaptive capacity; technological change; inter-firm heterogeneity

\*Correspondence to: Vikas A. Aggarwal, INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France. E-mail: vikas.aggarwal@insead.edu

Authors contributed equally and are listed alphabetically.

development (e.g., via setting the organization's exploration policy) may serve as a critical lever, enabling firms to systematically develop routines that are differentially adaptable to alternate forms of technological change.

Organizational routines have become an increasingly central theoretical construct in the strategy literature because of their close link to firm capabilities (Tece, 1982; Winter, 2000, 2003), and consequently, the question of whether and how these routines might change has become important (Feldman and Pentland, 2003).<sup>1</sup> Routines may reflect managerial efforts to purposively coordinate the actions of individuals—for example, through a combination of standard operating procedures and top-down hierarchy (Cyert and March, 1963; March and Simon, 1958). At the same time, a substantial stream of the literature on routines characterizes their *initial development* as stemming from a bottom-up process of learning (e.g., Cohen and Bacdayan, 1994; Greve, 2008; Hodgson, 2008; Hutchins, 1991; Knudsen, 2008; Levitt and March, 1988; Miner, Ciuchta, and Gong, 2008; Nelson and Winter, 1982) that results in an organization-level property that can be defined simply as “a repetitive pattern of interdependent actions, involving multiple actors” (Feldman and Pentland, 2003: 96).<sup>2</sup> These emergent routines then serve as the basis for “continuity in the behavioral patterns” of organizations (Nelson and Winter, 1982: 96), and have been thought to limit adaptation to external change (Hannan and Freeman, 1984; Leonard-Barton, 1992).<sup>3</sup>

<sup>1</sup>The literature has employed a range of terms that broadly relate to the term *routines*. As Cohen, Levinthal, and Warglien (2014) discuss, a range of “recurring organizational patterns” have been studied with labels such as routine, practice, standard operating procedure, etc. They propose using *collective performance* to describe such patterns, with the common denominator being that these all involve “habit-based actions.” This idea is consistent with what we model here; for simplicity (and because a broader discussion of terminology is beyond the present scope), we retain the term *routine*, recognizing that our use of this term is in the spirit of this notion of “collective performance.”

<sup>2</sup>See Becker (2004) for a review of the literature examining the emergence of routines. Pentland and Feldman (2005) note that, while standard operating procedures may indeed reflect top down coordination, they may also reflect the outcome of a process of codifying an emergent routine.

<sup>3</sup>Routines reflect one of three mechanisms of organizational coordination (March and Simon, 1958; Thompson, 1967), which additionally include hierarchical authority and inter-individual communication. While routines may *emerge* from bottom-up individual-learning, the literature has also recognized the important role that purposeful managerial action plays on them once formed (March and Simon, 1958; Teece *et al.*, 1997).

Yet, the literature also recognizes that the mechanisms of routine development do not abate after routine emergence, and routines themselves are able to adapt in dynamic settings (Feldman and Pentland, 2003). Recognizing that routines, as emergent organizational properties, can and do adapt is a first step, but it is not alone sufficient. What matters to strategy scholars is to identify sources of *systematic heterogeneity*. However, the extant literature offers only a very limited understanding of how systematic heterogeneity in the capacity of routines to endogenously adapt to technological change emerges and evolves. Moreover, we have limited insight into the ways in which managers might guide this process through explicitly enacted organizational policies.<sup>4</sup>

To address this gap, we pursue a microfoundational approach (Barney and Felin, 2013; Felin, Foss, and Ployhart, 2015; Foss and Pedersen, 2014; Winter, 2013). We propose that the value of doing so resides not in studying individuals' behaviors per se, but rather from identifying how individual behaviors interact and aggregate to generate macro-level phenomena. Our approach is anchored in individual-level learning and the formation of “habit” (Winter, 2013). To the degree that routines can be conceptualized as “structures of interlocking individual habits” (Hodgson, 2008: 25), the evolving beliefs of individuals are central not only to the emergence of routines (Becker, 2004; Greve, 2008; Knudsen, 2008; Levitt and March, 1988; Miner *et al.*, 2008), but also to their endogenous evolution and adaptation over time (Becker *et al.*, 2005).

We computationally model a process that could lead to the emergence of routines in organizations. We build on the bandit model of individual learning under uncertainty (e.g., March, 1996; Posen and Levinthal, 2012), extending this to an organizational environment in which one individual's actions impact the outcomes of the others in the organization (Levinthal and March, 1993; Thompson, 1967). The individual-level learning process is one in which individuals observe organization-level outcomes, and then update beliefs based on these outcomes. Individuals thus act in the context of substantial ambiguity (March and Olsen, 1975). As a baseline result, our model of individual-level learning in an ambiguous organizational environment produces outcomes with the canonical properties of routines: constancy, efficacy, and organizational

<sup>4</sup>See Foss *et al.* (2012) for a recent discussion that highlights these issues.

memory (e.g., Cohen and Bacdayan, 1994; Grant, 1996; March, 1994; Nelson and Winter, 1982).<sup>5</sup> Moreover, these “routines” exhibit substantial heterogeneity in efficacy, consistent with their role as drivers of inter-firm performance heterogeneity.<sup>6</sup>

Our core contribution is the recognition that the same process which gives rise to routines and to their performance heterogeneity also gives rise to substantial heterogeneity in the ability to self-adapt to technological change, a property we term *adaptive capacity*. We examine this adaptive capacity in response to three different forms of technological change that are all architectural in the sense that change rewires interdependencies (Henderson and Clark, 1990). We observe substantial heterogeneity in adaptive capacity across these forms of change. This suggests a fundamental link between the theory of routines (e.g., Cohen and Bacdayan, 1994; Grant, 1996; Nelson and Winter, 1982; Winter, 2000), the corresponding theory of inter-firm performance heterogeneity (e.g., Barney, 1991; Helfat and Peteraf, 2003; Penrose, 1959), and theories explaining adaptation to change (e.g., Henderson and Clark, 1990; Tushman and Anderson, 1986).

An important implication of our model is that managerial policy in formative periods, during which routines develop and grow, determines the performance and adaptive capacity of routines in maturity. In particular, we examine the impact of a key managerial policy: the degree to which individuals are encouraged by the organization to explore alternative ways of performing their tasks (March, 1991). We observe that such a policy in the formative period of routine development can, in maturity, impact not only performance, but also adaptive capacity, differentially across alternative

types of technological change. This results in a set of managerial trade-offs, not only between performance and adaptive capacity, but also between adaptive capacities to different forms of technological change.

In the next section we briefly review prior explanations for heterogeneity in the ability of firms to adapt to change, focusing on the alternate forms of technological change we model. In our computational model and subsequent analyses we aim to understand the role that routines play in influencing heterogeneous firm adaptation to these forms of change. We conclude by linking our results and insights to the literatures in strategy, focusing in particular on our contributions to understanding the link between heterogeneity in routines and adaptation to technological change.

## THEORETICAL FOUNDATIONS: ADAPTING TO TECHNOLOGICAL CHANGE

A large stream of the strategy literature has sought to address the question of why firms differ in their ability to adapt to technological change. In this section we briefly review this literature, characterizing it across the dimensions of inducements and difficulty. We then discuss how routines might serve as a possible explanation for inter-firm heterogeneity in this context, focusing on their differential adaptability to alternate forms of technological change that vary across these two dimensions.

### Inducements and difficulty of adaptation to change

Why do firms differ systematically in their likelihood of responding to technological change? Explanations in the prior literature on firm adaptation to change can be segmented along two dimensions. First, *inducements* in the form of perceived rewards to organizational adaptation can influence firms’ willingness to undertake adaptive efforts. Second, the *difficulty* of the required adaptive effort shapes the likelihood that a firm’s response to change will be effective.

A firm’s inducement to adapt to change has historically played a central role in the literature on technological change. It is reflected in the relative rewards associated with modifying a firm’s existing activities as a consequence of the balance between

<sup>5</sup>These properties result from our model even in the absence of inter-individual communication and hierarchical intervention. Nelson and Winter (1982) note that “the image of coordination... presented involves no mention of authority figures, backed by a system of incentives and sanctions, who cajole or coerce the required performances from other members.”

<sup>6</sup>With our model, we are able to address *one particular mechanism* through which routines emerge—namely, individual-level learning via performance feedback. We acknowledge, however, that there are numerous other institutional factors that might shape the emergence and development of routines over time—for example, communication, codification, culture, and so on. Our use of the term *routines* abstracts away from these various factors in order to focus on one particular mechanism of routine emergence, together with its implications for the strategy literature. While we refer to the outcomes of our model (in cases where they exhibit constancy, efficacy and organizational memory) throughout this manuscript as “routines,” we do so as shorthand, and with this particular caveat in mind.

the perceived costs and benefits of doing so. For instance, research examines how returns to investment in R&D are shaped by the radicalness of a new technology (e.g., Arrow, 1962; Reinganum, 1983), arguing that an incumbent's inducement to respond is diminished because the new technology may undermine the value of existing product lines. More generally, a firm's own resource base can serve not only to promote, but also to inhibit adaptation to change. Factors such as resource dependence (e.g., Pfeffer and Salancik, 1978), customer demands (e.g., Christensen, 1997; Christensen and Rosenbloom, 1995), and supplier commitments (e.g., Ghemawat, 1991) can influence the incentives for engaging in adaptation. As an example, disruptive change is characterized by current rewards to meeting the needs of existing customers that blunt the short-term impetus to adapt (Christensen and Bower, 1996). Likewise, complementary assets serve to buffer a firm from the need to adapt (Teece, 1986; Tripsas, 1997). Beyond the latent inducement to change, the literature suggests that cognitive blind spots may inhibit managers' recognition of the necessity of, and inducements for, change (e.g., Gilbert, 2005; Tripsas and Gavetti, 2000). In a recent study, Wu, Wan, and Levinthal (2014) argue that the same complementary assets that buffer the firm may also inhibit the recognition of the need for change.

A firm's difficulty in adapting to change has also played a central role in the literature on technological change. To some degree, the difficulty of adapting to change stems from the nature of the change itself. The literature distinguishing between incremental and discontinuous change (e.g., Abernathy and Utterback, 1978; Dosi, 1982; Tushman and Anderson, 1986; Tushman and Romanelli, 1985), for example, suggests that adapting to discontinuous change is more difficult than adapting to incremental change. Disruptive change is, likewise, thought of as reflecting a difficult type of change as it may require the development of a new set of firm capabilities (e.g., Henderson, 2006). Yet even holding constant the nature of the change itself, difficulty of adaptation must be measured relative to firms' ability to respond. Firms may vary in ability as a result of their distinct portfolio of capabilities, as some capabilities may simply be better suited than others for success in a particular environmental setting (Dosi, 1982; Tushman and Anderson, 1986).

### **The role of routines in influencing firm adaptation**

Absent in the discussion above regarding inducements and difficulty is any mention of routines. Yet, the literature suggests that routines are a central component of the explanation as to why firms differ in their ability to adapt in the face of technological change. Henderson and Clark (1990: 27), for example, address how industry-level architectural changes can have profound implications for the ways in which a firm's activities are interdependent. Such changes create particular difficulties because existing routines render a firm's knowledge base "inert and hard to change." Other streams of the technology change literature likewise suggest a similarly important role for routines—for example, Christensen and Bower (1996) argue that the resource allocation process driving investments into new technologies can be stymied by routines geared to existing customer needs. Yet while the technology change literature generally suggests that routines may hinder adaptation, it is relatively silent on how routines might at the same time be a *systematic source of heterogeneity* in the capacity of firms to adapt to change (or likewise, heterogeneity in the extent to which they hinder adaptation).<sup>7</sup>

Examining the process through which routines arise in the first place may elucidate the origins of such heterogeneity. The literature has often conceptualized routines as emergent properties that coordinate and direct action (Becker, 2004; Cohen and Bacdayan, 1994; Nelson and Winter, 1982). Underlying this perspective on the origin of routines is the central role of individuals, and in particular, the "habits" of individuals that result from learning (Hodgson, 2008; Winter, 2013). In an organizational setting, individuals' habits co-evolve in an ambiguous environment. As Cohen *et al.* (1996: 661) note, routines can be thought of as residing "in the memories of the individual actors for their respective roles in the overall pattern." Knudsen (2008: 131–132) echoes this point, noting that "the definition of a routine as an organization-level disposition is firmly rooted in the human individual... individual habits are the microfoundations of organizational routines."

<sup>7</sup>The one exception is that firms may differ in the extent of routinization, and as such, older and larger firms, which tend to be more routinized, are less able to adapt (Hannan and Freeman, 1984).

This view of individuals as the central mechanism in the emergence and evolution of routines also suggests that routines might change via changes in individuals' evolving beliefs (i.e., habits). Knudsen (2008: 131–132) points out that “there is a recursive relation between routines and individual habits,” suggesting that routines not only emerge, but also change, through changes in individual habits. The context in which individuals operate—for example, the structure of interrelationships between individuals—shapes the process of emergence and change: as Hodgson (2008: 25) notes, “routines are more than mere aggregations of habits, because they also depend on the emergent properties of the organization itself, emanating from structured causal relations and interactions between individuals.” Routines, therefore, are likely to change via “forces and actions of agents internal to the routine,” with this change occurring “endogenously because human actors are involved in carrying out routines” (Becker *et al.*, 2005: 776–777).

With this in mind, we propose a behaviorally plausible model of the emergence of routines, and in this context show how the process leading to their emergence might at the same time endow them with heterogeneity in their adaptive capacity to alternate forms of technological change.

### Conceptualizing alternate forms of technological change

In line with Henderson and Clark's (1990) view that technological change can necessitate a rewiring of a firm's interdependencies, a difficult task given the firm's existing routines, we examine firms' adaptive capacity in the context of alternate forms of such architectural change. As we illustrate in Figure 1, we conceptualize (and later model) these three forms—incremental, discontinuous, and entrapping—which differ along the dimensions of inducements and difficulty discussed above. Although these forms of change are not exhaustive in their representation of the broader technology change literature, collectively they offer a backdrop against which we can investigate the sources of systematic heterogeneity in the adaptive properties of routines.

The first two types of change are high inducement, varying only in their difficulty. While “incremental” change is low difficulty, “discontinuous” change is high difficulty, requiring substantial organizational effort. The third type of change, which

		Difficulty of Adaptation	
		Low	High
Inducement to Adapt	Low	Status Quo	Entrapping Change
	High	Incremental Change	Discontinuous Change

Figure 1. Adaptation to technological change—inducement and difficulty

we call “entrapping” change, reflects a key dimension of disruptive change (Christensen and Bower, 1996)—it is low inducement, but high difficulty. For example, an established incumbent may be doing well enough serving existing customers, and thus, have reduced inducement to change. At the same time, the change effort, should the incumbent wish to get onto the new technology, would be quite difficult. Of course, the entrapping change we consider abstracts from many of the other features of the Christensen story (e.g., there is no technology evolution).<sup>8</sup>

### MODEL

Our model is anchored in the view that routines emerge from individual performance feedback-based learning (e.g., Cohen and Bacdayan, 1994; Greve, 2008; Hodgson, 2008; Knudsen, 2008; Levitt and March, 1988; Miner *et al.*, 2008; Nelson and Winter, 1982). In so doing, we make three foundational assumptions. First, the basis of individual action is beliefs about the merits of alternatives that result from experiential learning (Bush and Mosteller, 1951; Dewey, 1922; March, 1996; Winter, 2013). Second, individuals are interdependent with one another in an organizational setting (Levinthal and March, 1993; Thompson, 1967). Third, the environment in which learning occurs is one of substantial ambiguity (e.g., Cohen and

<sup>8</sup>We do not intend to coin a new term here; rather, we use *entrapping* as shorthand to capture one centrally important aspect of disruptive change.




Figure 2. Task environment. This figure illustrates eight possible combinations of three agents' choices (patterns of actions). Each of the cells represents a probability of success for the organization when the individuals execute a given pattern of action. When agents select UP, FRONT, and LEFT, for example, the expected probability of success for the organization is 0.85

Bacdayan, 1994; Levitt and March, 1988; March and Olsen, 1975; Nelson and Winter, 1982).<sup>9</sup>

We employ the canonical multi-armed bandit model, which the extant literature suggests accurately represents the psychological processes underlying individual learning in complex settings (Puranam *et al.*, 2015).<sup>10</sup> The bandit model takes its name from a slot machine analogy in which an individual plays a slot machine with multiple arms. The arms can be conceptualized as alternative approaches an individual has to performing a given task (e.g., two ways of connecting pieces of metal, weld or bolt). Each "task approach" is associated with a specific probability of yielding a success/payoff in a given period. These latent payoff probabilities are *ex ante* unknown to the individual, and thus, individuals cannot optimize. Rather, individuals seek to enhance their efficacy through a process of learning about the merits of their task approaches by doing (i.e., trying the approaches). We extend the bandit model, which represents a single individual in isolation, to include three individuals operating with interdependent

payoffs.<sup>11</sup> Individuals take action, observe the resulting *organizational performance*, and then update their individual beliefs regarding the merits of their alternatives. Ambiguity arises for two reasons: first, because individuals make simultaneous choices, leading to causal ambiguity from the perspective of any given individual observing the organization-level performance; and second, because organizational payoffs involve a substantial stochastic component.

### Task environment and performance

The organizational task environment is defined by the set of unique combinations of the task approaches of the three individuals, which we refer to collectively as the organization's "pattern of action." Each of the three individuals undertakes a different task, and has two alternative approaches to performing their task. The task environment contains eight patterns of action, as illustrated in Figure 2. For each of the eight cells, the value indicated is the organization's probability of receiving a payoff of one (success) in a given period.

The organizational payoff values in our task environment are constructed, for simplicity and parsimony, so that two patterns of action yield an

<sup>9</sup>We follow Nelson and Winter (1982) in abstracting from hierarchical authority or purposeful efforts at coordination through direct communication among individuals (we relax this in our robustness analyses). Nelson and Winter (1982: 107) note in their discussion of routines: "the image of coordination... presented involves no mention of authority figures, backed by a system of incentives and sanctions, who cajole or coerce the required performances from other members."

<sup>10</sup>Examples of the use of the single-individual multi-armed bandit model are found in the fields of economics (e.g., Berry and Fristedt, 1985; Gittins, 1979), computer science (e.g., Holland, 1975; Sutton and Barto, 1998), and management (e.g., Denrell and March, 2001; March, 1991, 1996; Posen and Levinthal, 2012). See Puranam *et al.* (2015) for a review of this class of models.

<sup>11</sup>Using a three-individual model suggests that a caveat may be in order regarding how much we are able to say regarding routines. On the other hand, Cohen and Bacdayan's (1994) study uses a two-person experiment, yet is widely recognized as reflecting routine formation. Moreover, as Puranam, Alexy, and Reitzig (2014: 164) suggest, the number of members is not a defining feature of an organization: "the conceptualization of an organization we have adopted precludes neither a dyad nor a corporation from consideration as an organization."

above-average payoff, and the remaining six yield below-average payoffs. Below-average combinations have equal average payoffs, and the average payoff of all combinations is set to 0.5. Consequently, there are two patterns of action we can characterize as “peaks,” with organizational payoff probabilities of 0.85 and 0.6, respectively. The remaining patterns of action have organizational payoff probabilities of 0.425.<sup>12</sup> As we discuss in more detail in the Analysis section (and as we later illustrate in Figure 4), this structure is useful because it enables us to simulate different forms of technological change (incremental, discontinuous, and entrapping) by changing the payoff matrix in such a way that causes firms at the global peak to experience changes that vary along two distinct dimensions: inducements and difficulty.

### Forming beliefs and taking actions

Each individual forms beliefs about the merits of his or her two task approaches, and then takes actions according to these beliefs.<sup>13</sup> We define  $Q_{int} = (q_{i1t}, q_{i2t})$  as the set of beliefs for each individual  $i$  at time  $t$  regarding the payoff probability from a choice of task approach  $n$ . We endow each individual with an uninformed initial prior,  $Q_{in0} = (0.5, 0.5)$ , so that individuals are initially indifferent between the two alternative task approaches. Individuals make simultaneous choices regarding their task approaches in each period, and receive feedback on realized organizational performance per the probabilities in Figure 2. Belief updating follows the fractional updating structure of Bush and Mosteller (1951) and March (1996). In particular, the belief for a given individual regarding the probability of success for a particular choice of task approach at any given point in time  $t$  is simply the average of the individual’s prior experience with that approach.

In any given period, individuals’ choice of task approach is a function of their beliefs about the merits of the alternatives, together with their decision

rule for acting on those beliefs. This decision rule determines, conditional on their beliefs, the extent to which they are willing to explore an alternative task approach. We conceptualize exploration as resulting from an organizational exploration policy that influences the propensity of individuals in the organization to explore (e.g., March, 1994; Siggelkow and Rivkin, 2006), although we recognize that it may also be a function of individual-level dispositions.<sup>14</sup>

We follow the prior literature on bandit models in which a “softmax” rule determines how an exploration policy impacts actions (eg, Luce, 1959; Posen and Levinthal, 2012; Sutton and Barto, 1998). In the softmax rule, the parameter  $\tau$  reflects the organizational exploration policy. It does not define the choices that the individuals make, but rather, how individuals act in response to their own beliefs. As  $\tau \rightarrow 0$ , individuals choose the task approach they believe to be best at that time; and as  $\tau \rightarrow \infty$  individuals become more exploratory, more likely to forgo the alternative currently believed to be superior in hopes of finding a better alternative. For individual  $i$  at time  $t$ , the softmax rule maps beliefs to actions via an equation describing the probability  $m_n$  of an individual selecting a task approach  $n$ , where  $q_n$  represents the corresponding belief from the vector  $Q_{int}$ : 
$$m_n = \left[ \exp(q_n/\tau/10) \right] / \sum_{g=1}^2 \exp(q_g/\tau/10).$$

## ANALYSIS

### The emergence of routines

As a first step, we examine the properties emerging from our model. We hope to observe outcomes with properties that, following Nelson and Winter (1982), are characteristic of routines: constancy, efficacy, and organizational memory.<sup>15</sup> This will

<sup>12</sup>For example, if the pattern of action is such that individual  $i=A$  chooses task approach UP, individual  $i=B$  chooses approach FRONT, and individual  $i=C$  chooses approach LEFT, then the organization will draw from a Bernoulli distribution with a mean of 0.85.

<sup>13</sup>The strength of an individual’s preferences, measured for example as the distance between the beliefs each individual holds over each of the two task approaches, is analogous to the concept of “habit” in the routines literature.

<sup>14</sup>Managerial policy can influence individual actions through multiple levers. In the hiring process, for example, firms often choose individuals based on their propensity to explore (e.g., hiring practices at Google). Incentives and rewards, organizational structures, and culture can also each be powerful organizational drivers of individual-level exploration. As an example, consider Facebook’s motto for many years: “Move fast and break things.” Such an overarching cultural directive likely has implications for how individuals act on and develop their own beliefs and preferences.

<sup>15</sup>Nelson and Winter (1982) note, in chapters 4 and 5, that routines involve “a smooth sequence of coordinated behavior that is ordinarily effective relative to its objectives” with “no mention of authority figures,” representing “the nature and sources of continuity in the behavioral patterns of an individual organization.”




Figure 3. Emergence of routines. Holding exploration policy fixed across time at  $\tau = 0.5$

form the basis of our subsequent analysis in which we examine heterogeneity in adaptive capacity to different types of change

In Figure 3 we present a set of organization-level outcomes from our model. In these baseline analyses, we set exploration to a moderate level,  $\tau = 0.5$  (per Posen and Levinthal, 2012), reporting averages across 100,000 trials to reduce any statistical artifacts.<sup>16</sup> The measure of *constancy* captures the realized level of stability in the patterns of action across time. We measure it, in any given period, as the probability that the pattern of action will remain unchanged in the subsequent period.<sup>17</sup> The measure of *performance* captures the average realized result of the organizational pattern of action in any given period, determined (stochastically) via the mean probability of organizational success indicated in the corresponding box in Figure 2 (normalized on a 0–1 scale).

The constancy and performance lines in Figure 3 show that the process of individual learning over time leads to both increased constancy and performance. There is substantial inter-firm heterogeneity in performance arising from our model. In period 25, for example, variance in performance is 0.23, which is substantial relative to the average

with routines serving as “organizational memory,” in a way that organizations “remember a routine largely by exercising it.”

<sup>16</sup>This sample is sufficiently large that, unless otherwise noted, any measured differences across alternative model settings are statistically significant.

<sup>17</sup>Specifically, we calculate constancy for a given round by comparing the choices taken by each individual in that round with those taken in the previous one. We assign 0 if the choice is the same, and 1 if they differ. We record constancy as an average across all three individuals.

performance of 0.52. To understand the composition of this performance heterogeneity, we examine the frequency with which firms select particular patterns of action. Approximately half (46%) of the firms execute the best pattern of action, while 16 percent execute the second best pattern of action (these patterns have performance values of 0.85 and 0.6, respectively). The remaining firms execute a pattern of action that is in the “neighborhood” (i.e., on an adjacent pattern) of one of the two favorable patterns, with the associated patterns having average performance values of 0.425.<sup>18</sup>

We also construct a measure of *organizational memory* (Cohen and Bacdayan, 1994; Nelson and Winter, 1982). To do so, we analyze how the replacement of one of the three individuals (a turnover event) impacts the observed pattern of action. In the turnover event, a new individual enters with no preference for either of her two task approaches and does not communicate with the replaced individual. Consequently, the replaced individual’s entire stock of knowledge is lost. We consider a routine to exhibit memory if the new individual is more likely than chance to behave like the replaced individual, and as such, the organization executes the same pattern of action it would have executed had the turnover event not occurred.<sup>19</sup> The reported values of the line labeled *organizational memory* in Figure 3 represent this probability, showing an increase in organizational memory over time.<sup>20</sup>

Ambiguity is central to the mechanisms by which these various properties of routines—constancy, performance, and organizational memory—result from our model. Given the presence of interdependence across individuals, not only do a given individual’s preferences across task approaches shape the individual’s own actions, but one individual’s preferences shape other individuals’ actions

<sup>18</sup>We also examined how heterogeneity in performance changes with the exploration policy,  $\tau$ . In the results presented here, we set  $\tau = 0.5$ . We have examined the range of  $\tau$  from 0 through 1. Over this range, performance heterogeneity follows a U-shaped pattern. These results are available on request.

<sup>19</sup>We calculate the probability that within 100 rounds subsequent to a turnover event the organization returns to the same pattern of action it would have attained had the turnover event not occurred.

<sup>20</sup>The patterns in Figure 3 remain qualitatively the same with changes in  $\tau$ . Although for some high levels of  $\tau$  (e.g., greater than 5), constancy may drop to the point at which we might doubt whether the action we observe can still be called a routine, for the ranges of  $\tau$  we consider here a routine will emerge at some point in time during the simulation.

as well. This property is critical to the emergence of routines because it gives rise to local search, which we also plot in Figure 3. Local search arises because constancy of the actions of individuals—which ultimately leads to the “repetitive, recognizable pattern of interdependent action, involving multiple actors” per the Feldman and Pentland (2003) definition of a routine—is necessary for any individual to learn effectively in an ambiguous environment.<sup>21</sup>

Consider the pattern of organizational search over time. In early periods, individual habitual behavior is yet uninformed, and thus, any given pattern of action at the organization-level is equally likely to be selected. This is an ineffective and slow means of learning because it engenders causal ambiguity due to the difficulty in mapping performance feedback to a particular individual’s action (Gavetti and Levinthal, 2000; Zollo and Winter, 2002). As time progresses, given interdependence between individuals, some level of constancy is necessary for efficacious learning that enhances performance because learning is more effective when only one individual searches at a time. As beliefs are formed and reinforced, an individual will begin to form increasingly strong preferences across task-approaches, which then sets in motion a self-reinforcing process: as the extent of simultaneous exploration diminishes, the efficacy of learning increases, performance improves, individuals’ preferences across task approaches strengthen and become increasingly habitual, actions become more constant, and search becomes increasingly local.

The resulting organizational behavior has the definitional characteristics of a routine: increasing constancy and performance over time, and organizational memory. This does not imply that performance is homogeneous. Even though all firms and individuals are identical in terms of the task environment they face and their initial beliefs, search may stall at a suboptimal pattern of action. Organizational memory exists

because incumbent individuals’ preferences for task approaches impact the performance feedback received by a new individual on the choice of a particular task approach. This, in turn, shapes the new individual’s preferences and subsequent action in a manner that makes the new individual more likely than chance to behave like the replaced individual.

### Inter-firm heterogeneity in adaptive capacity

Our central argument is that routines are innately heterogeneous in their capacity to adapt to technological change, just as they are heterogeneous in their efficacy. In the technology change literature, as we noted earlier, routines are typically conceptualized as inhibiting organizational adaptation. Our claim is not simply that routines may exhibit a capacity to adapt (for a nice discussion see Feldman and Pentland [2003])—such a claim alone is insufficient to be strategically relevant. Rather, from a strategy perspective that focuses on inter-firm heterogeneity, we explore the possibility that routines may differ *systematically* in their capacity to adapt to technological change, and that this capacity to adapt is managerially tunable.

We examine heterogeneity in routines’ adaptive capacity in relation to three forms of technological change to which our model is amenable: incremental, discontinuous, and entrapping change. These forms of change can be distinguished from one another along two dimensions: *inducement to adapt* and *difficulty of adapting* to change (per the earlier Figure 1). We illustrate these three forms of change in Figure 4. Each of the nonstatus quo panels in this figure reflects a shift in the task environment, and Figure 4 illustrates the implications for each shift from the perspective of a market leader (an organization previously on Peak 1).

In the case of incremental change, a leader (on Peak 1) before the shock needs just one individual to change his or her preferred task approach in order for the organization to return to Peak 1 after the change (“low difficulty”). At the same time, expected organizational performance to the current pattern of action post-shock drops from 0.85 to 0.425 (“high inducement”). With discontinuous change, a leader (on Peak 1) before the shock needs two individuals to change their preferred task approach (“high difficulty”), while expected organizational performance again drops from 0.85 to 0.425 (“high inducement”). Finally, with entrapping

<sup>21</sup>Our definition of local search is at the organization level, and not at the individual level. With three individuals, each of which is a two-armed bandit, there are eight possible organizational outcomes. Local search is defined as a situation in which one individual changes her preferred task approach, and thus, the organization moves “locally” from one adjacent cell in the  $2 \times 2 \times 2$  cube to the other. In the extant literature, local search is an assumption (e.g., Ahuja and Katila, 2004; Cohen and Levinthal, 1989; Cyert and March, 1963; Nelson and Winter, 1982). By contrast, in our model, local search is an outcome of individual-level learning: we make no *ex ante* assumption that the organization must search locally; rather, local search emerges endogenously over time.




Figure 4. Three types of technological change. This figure illustrates the three forms of technological change. The Status Quo organizational pattern involves the three individuals’ selected task approaches being (UP, FRONT, LEFT); these choices result in Peak 1, which has a mean organizational payoff of 0.85. The three forms of change differ with respect to “difficulty”—how many individuals must change their pattern of action to get back to Peak 1; and also with respect to “inducement”—the magnitude of the drop in expected probability of organizational success. Example: Going from *Status Quo* to *Discontinuous Change* means that the focal organizational pattern, which is at (UP, FRONT, LEFT), needs to change to (UP, BACK, RIGHT) in order for the organization to stay at Peak 1. This is “high difficulty” since two individuals need to change their preferred task approach. At the same time, the organization choosing (UP, FRONT, LEFT) goes from a mean organizational payoff of 0.85 before the change with the *Status Quo*, to 0.425 after the *Discontinuous Change*. This is “high inducement”

change, a leader (on Peak 1) before the shock also needs two individuals to change (“high difficulty”), but the post-shock expected organizational performance drops only to 0.59, instead of 0.425 as in the other cases.<sup>22</sup> This case thus represents a “low inducement” form of change.

We measure adaptive capacity to technological change as the likelihood of the firm having the same expected performance after the change as it

<sup>22</sup>This specific value is chosen to satisfy several constraints as we seek to minimize the difference between discontinuous

and entrapping as much as possible to ensure comparability. In particular, the value must be below that of Peak 2 (0.6) and higher than the floor (0.425). Within these constraints the lowest inducement is the maximum value in this range.

would have had if the change had not happened.<sup>23</sup> To construct the measure, we first simulate the organization through period 125 in the absence of a change event, calculating the modal pattern of action in the latter 100 periods. We then repeat the calculation with a change that occurs at period 25, comparing it to what would have happened had the change not occurred.<sup>24</sup> We simulate 1,000 organizations over 1,000 trials each to estimate the value of adaptive capacity.

In examining heterogeneity in adaptive capacity, we consider both across and within change-type heterogeneity. There are substantial differences across types of technological change. The average value of adaptive capacity is greatest for incremental change (0.354), moderate for discontinuous change (0.131), and lowest for entrapping change (0.085). There is also substantial heterogeneity in adaptive capacity within a particular type of change. The variance in adaptive capacity within incremental change is highest, 0.064, while for discontinuous and entrapping change the variance is, respectively, 0.025 and 0.017. One implication of this set of results is that it makes little sense to talk about the extent to which routines are adaptive (or inert) in general. Rather, the extent of adaptive capacity is specific to a particular type of change. This observation is important in a strategy sense because, to the extent that firms can shape the adaptive properties of routines, there will likely be substantial trade-offs involved (we examine this possibility in the next experiment).

### Mechanisms underlying heterogeneity in adaptive capacity

What explains heterogeneity in adaptive capacity in our model, both *across* and *within* the three

forms of technological change? The core mechanism underpinning our results involves individual beliefs regarding the merits of alternative task approaches, and in particular, the degree to which these beliefs are strongly ingrained.<sup>25</sup> When a technological change occurs, individuals must readjust their beliefs in response to the new environment. For this to occur successfully, individuals must be “dislodged” from their current beliefs about the merits of the alternative task approaches. The possibility of this occurring is a function of not only any individual’s beliefs, but also of the properties of others’ beliefs. In this sense, a routine is characterized by its “strength of preferences,” which reflects the degree to which beliefs of individuals in the aggregate are strongly ingrained. Strength of preferences underlies both the ability to search distantly as well as the willingness to search, and via these joint effects, the adaptive capacity to different types of technological change.

Average adaptive capacity decreases as we proceed from incremental to discontinuous, and then to entrapping change. When a technological change occurs, individuals must readjust their beliefs in response to the new environment. For this to occur successfully, individuals must be “dislodged” from their current beliefs. At any given point in time, the strength of preferences associated with a routine will determine the degree to which the routine will be capable of adapting to a given form of technological change. As the level of difficulty associated with technological change increases, for example, from incremental to discontinuous change, more distant search is required to adapt to the change (i.e., two individuals must change their beliefs versus just one). Thus, for a given strength of preferences, a routine will exhibit lower adaptive capacity to discontinuous (versus incremental) change. More interestingly, holding difficulty fixed, inducement associated with technological change is equally important. As the level of inducement decreases, for example, from discontinuous to entrapping change, the challenge is not one of a need for more distant search, but rather, the willingness to search at all, which requires foregoing the current pattern of action that is only mildly inferior. In this case, for a given strength

<sup>23</sup>Note that our definition of *adaptive capacity* involves moving to the pattern with the same expected level of performance following the shock. We examined an alternate set of measures based on moving to the same or better pattern of action as well as based on the time needed to return the original level of performance. While there are some quantitative differences when employing these alternate definitions, our results are qualitatively similar.

<sup>24</sup>We select period 25 in which to introduce different forms of technological change as this is the period in which constancy passes 90 percent, and 95 percent of its long-run value. (We also replicate the results with the change in period 100 and find that the results are qualitatively similar.) In this period there is substantial inter-firm performance heterogeneity, with this heterogeneity resulting from the process of routine emergence discussed previously.

<sup>25</sup>This mechanism is consistent with the claim that the emergence of routines is closely linked to learned individual “habit” (Winter, 2013).

of preferences, a routine will exhibit lower adaptive capacity as the inducement associated with the change diminishes. Thus, difficulty and inducement act together, via their effects on changing individual beliefs and routine-level strength of preferences, to determine variation in adaptive capacity across forms of change.

Individual beliefs, and the strength of preferences they embody, can also explain why adaptive capacity differs across firms *within* a particular type of change. In this case, beliefs and strength of preferences are influenced by two key drivers: position and organizational path. The first driver, position, relates to the particular pattern of action on which the firm is coordinating at the time of the shock—that is, whether this is Peak 1, Peak 2, or the other patterns in the neighborhoods of these peaks. Figure 5 shows that the performance quartile the organization is on has implications for the level of adaptive capacity. To understand why this is the case, consider what happens to industry leaders (Peak 1). A firm at Peak 1 receives positive performance feedback 85 percent of the time (per Figure 2). In each period, individuals update their beliefs based on the new (organization-level) feedback. The reliably positive feedback tends to rapidly solidify individuals' beliefs, leading to habitual behavior that is more firmly ingrained (i.e., success affirms the beliefs of the individuals, making them more certain of their current beliefs). Consequently, the organization will be much more strongly bound to the particular pattern of action. This is, in part, why top performers (e.g., first quartile firms) tend to be less adaptable to discontinuous or entrapping change.

The second driver relates to the specific path the organization took to get to that particular position (pattern of action). Figure 5 shows that even *within* performance quartiles there is substantial heterogeneity. We plot (along the vertical axis for each quartile) the mean and distribution of individual values of adaptive capacity for the firms in our simulation. Heterogeneity occurs because firms may have arrived at a particular position via very different paths. This is important because, conditional on a firm's position, the specific path the organization took to get to that position shapes the extent to which individuals' habits are strongly ingrained. In turn, this impacts adaptive capacity.<sup>26</sup>

<sup>26</sup>Figure 5 also shows that there are conditions under which our model behaves as expected in the extant literature in the sense that

We illustrate the implications of heterogeneity in strength of preferences for a firm's adaptive capacity in Figure 6. This figure presents a scatter plot of organizations' strength of preferences versus adaptive capacity in  $t=24$  for firms that are choosing the best pattern of action (Peak 1) immediately pre-shock. An organization's strength of preferences at any given point in time is measured as the distance between the two values of each individual's belief vector,  $Q_{int}$ . We decompose the results across the three types of change, with the fitted lines in the figures showing that, conditional on choosing the best pattern of action pre-shock, firms with individuals holding stronger preferences for their task approach exhibit lower adaptive capacity to technological change. The negative slope is greatest for incremental change, and gets smaller (in absolute value) across discontinuous and entrapping change respectively. This implies that for incremental change, heterogeneity in adaptive capacity is more strongly driven by differences across organizations in the strength of preferences of their members (than it is for discontinuous and entrapping change). Strength of preferences impacts adaptive capacity because, with behavioral habits more strongly ingrained, search is more local, and consequently, adaptation to change is harder to achieve.

### Managerial policy and heterogeneity in adaptive capacity

Our model suggests that routines are heterogeneous in their capacity to adapt to technological change, both across different types of change and within a particular type of change. This, in turn, suggests that differences across firms in the adaptive capacity of routines may help us to understand why firms are differentially effective at adapting to change. From a strategy perspective, the relevance and importance of our adaptive capacity results rest on the existence of a managerial policy that can tune these adaptive properties of routines. Here, we examine one such policy—a firm's exploration policy—that affects adaptive capacity, and discuss the strategic trade-offs from adjusting that policy.

\_\_\_\_\_

routines exhibit reduced adaptive capacity over time as feedback reinforces current beliefs. This shows that our model reproduces key features of the existing understanding in the literature. As the figure shows, however, there is substantial heterogeneity in adaptive capacity, even for entrapping change, and even among firms in the top performance quartile.




Figure 5. Heterogeneity in adaptive capacity across types of technological change. We keep  $\tau$  fixed at 0.5. We simulate 1,000 organizations and divide this group into quartiles based on their performance in the five periods prior to the shock (cumulative performance from the 20th to 24th round). For each organization, we then estimate adaptive capacity by simulating 1,000 post-shock alternative futures. The figure shows the mean adaptive capacity for each performance quartile, as well as individual values of adaptive capacity for all simulated organizations

Our earlier results point to the critically important formative period during which routines emerge (Helfat and Peteraf, 2003). During the formative period, routine-like behavior results as individuals’ preferences for particular task approaches strengthen, local search emerges, and this in turn leads to increasingly effective learning and enhanced performance. Thus, we examine the implications of changing a firm’s exploration policy,  $\tau$ , during the formative period of routine development. In particular, we examine  $\tau$  over the range 0.1–1.0 during the first 25 periods; we refer to this as the *formative*  $\tau$ . In period 25, we introduce the change, and thereafter, for the remainder of the simulation, we set  $\tau$  to 0.25 (regardless of the value of  $\tau$  during the formative period). Consequently, differences in the observed

adaptive capacity of routines arise only from heterogeneity in managerial exploration policy in the formative period.

Figure 7 displays the results, where adaptive capacity is normalized by the average level of adaptive capacity to a particular type of change. At a very high level, we observe that this managerial policy during the formative period of routine emergence has a substantial impact on the long-run properties of those routines (even though, in this experiment, managerial policy post-shock is identical across firms). Increasing formative  $\tau$  does not necessarily increase adaptive capacity to change. Underlying this result is that formative  $\tau$  does not linearly impact the strength of preferences associated with a routine. Low and high formative  $\tau$  lead to weaker preferences than does a moderate level of




Figure 6. Strength of preferences versus adaptive capacity for firms at Peak 1 pre-shock. We simulate 1,000 organizations and record the strength of preferences for each organization in the period preceding the environmental shock. We also estimate the value of adaptive capacity to each type of change by running 1,000 iterations for each organization and recording the number of runs when the given organization managed to adapt to the change versus those when it did not. We then produce a scatterplot to record the relationship between the strength of preferences just before the shock and the resulting value of adaptive capacity, and superimpose a fitted line

formative  $\tau$ . With low formative  $\tau$  search is insufficient, resulting in a poorly performing pattern of action, and with high formative  $\tau$  excess exploration by all members of the firm increases ambiguity, making learning very hard. In both cases, preferences are weak, and poor patterns of action are easily abandoned in the face of a shock. Moderate formative  $\tau$  leads to the strongest preferences for the pattern of action.

The relationship between formative  $\tau$  and adaptive capacity (i.e., shape of the curve in Figure 7) differs across the three types of technological change. Incremental change tends to benefit uniformly from an increase in formative  $\tau$ . By contrast, the two more challenging forms of technological change, discontinuous and entrapping, exhibit strong U-shaped properties. The marginal returns to increasing adaptive capacity (slope of the line) differ across the three types of change

(e.g., at higher levels of formative  $\tau$ , the slope is lowest for incremental change). Formative  $\tau$  is important not because of its impact on the level of exploration, but rather, because it shapes the beliefs that individuals hold about their alternative task approaches, and thus, the extent to which habitual behavior emerges and co-evolves into routine-like recurring organizational patterns of action.

Moreover, formative  $\tau$  conditions the returns to exploring post-shock. In results available in the Appendix, we show that the “optimal” level of post-shock  $\tau$  varies quite strongly with formative  $\tau$ , with substantial differences across the alternative forms of technological change. For example, in the case of entrapping change, the concave shape of the surface suggests that increasing formative exploration tends to reduce the level of post-shock exploration needed to adapt to such change. For incremental change, increasing formative exploration




Figure 7. Impact of formative exploration policy on adaptive capacity. To estimate the effects of different pre-shock exploration policies (formative  $\tau$ ) on the subsequent adaptive capacity, we fix the post-shock exploration policy ( $\tau = 0.25$ ) and vary exploration policy in the initial 25 periods. We then compare the resulting adaptive capacity against a scenario where the initial exploration policy is chosen at random

reduces sensitivity to the level of post-shock exploration. That is, exploring in advance of the shock shapes the properties of the resulting routine such that under-exploring post-shock is much less damaging. More generally, these results suggest that the appropriate choice of post-shock exploration is a function of how routines were conditioned by formative exploration. An interesting implication of this observation is that two firms having settled on the same pattern of action may respond differently to the same shock because their exploration in the formative period was different.

One way of interpreting these results is that a higher exploration policy in the formative periods leads to greater adaptive capacity to all types of technological change. Yet, this figure does not imply that the “optimal” formative  $\tau$  should be set very high, because the figure does not take into account the opportunity costs associated with alternative settings of  $\tau$  during the formative period. In setting formative  $\tau$ , there are substantial strategic trade-offs of two sorts. First, higher  $\tau$  during the formative period may be suboptimal in terms of performance during the formative period. While increasing formative  $\tau$  may increase adaptive capacity, it may at the same time reduce a firm’s performance in the formative period.<sup>27</sup>

<sup>27</sup>We provide a more detailed example of this point in the online appendix, which illustrates the performance decline (though with increased adaptive capacity) that stems from increasing levels of exploration.

A second strategic trade-off stems from the observation that a firm does not know what type of change will occur in the future. This is important because the appropriate level of formative  $\tau$  differs across types of change. Consequently, firms may implement an organizational exploration policy that is, *ex post*, inappropriate for the realized form of change. To examine this possibility, we consider the implications of setting low versus high organizational exploration policy ( $\tau = 0.2$  and 5) in the formative period. We then examine the implications for post-shock (period 25–50) average performance, assuming the firm sets the exploration policy that maximizes performance in the post-shock period. We find that exploration policy in the formative period has a significant and lasting effect. As Figure 8 shows, under conditions of incremental change, the best achievable performance is 44 percent higher for firms choosing a high (relative to low) exploration policy in the formative period. Under discontinuous change, this performance advantage is 26 percent. High formative exploration policy, however, does not always put firms at an advantage. With entrapping change, the best possible performance is 3.3 percent lower for firms choosing a high (relative to low) exploration policy.<sup>28</sup>

### Robustness analyses

We conducted a number of robustness analyses to examine the implications of embedded assumptions in our model. In the Appendix, we include results showing alternative information flow regimes, alternative task environments, skill-based learning, routine development under different levels of  $\tau$ , and the joint effects of formative and post-shock  $\tau$ . We also conducted analyses, available on request, to examine alternative timings of the shock, alternative measures of adaptive capacity based on the time needed to return to the original level of performance, and using different periods post-shock to measure performance. While the results vary quantitatively, the conclusions and insights with

<sup>28</sup>There will be a variety of additional costs that firms would experience in preparing for different kinds of change. These costs could be included exogenously in our model. However, we conceptualize the cost in terms of lost performance as the maximum additional change costs the firm would be willing to incur.




Figure 8. Post-shock performance advantage of high formative exploration policy across different forms of technological change. We calculate the post-shock performance advantage of high (relative to low) exploration policy in the formative period. The high (low) exploration policy regimes are characterized by  $\tau = 5$  ( $\tau = 0.2$ ) for the first 25 periods. The technological change occurs in the 25th period, and thereafter, the firm employs the exploration policy that produces the best possible performance post-shock (period 25 through 50). Thus, any observed difference in post-shock performance is due only to exploration policy in the formative period. Positive values imply that high formative  $\tau$  produces better outcomes than low formative  $\tau$ , while negative values suggest the opposite

respect to routine emergence and adaptive capacity are qualitatively robust to these alternative assumptions.

## DISCUSSION

We examine the proposition that inter-firm heterogeneity in the capacity to adapt to technological change, as well as inter-firm heterogeneity in performance, stem from a common underlying process through which routines emerge. We develop a microfoundational model of routine emergence, anchored in behaviorally plausible assumptions about individual learning in an environment characterized by interdependence and ambiguity (e.g., Hodgson, 2008; Hutchins, 1991; Levitt and March, 1988; Nelson and Winter, 1982).<sup>29</sup> Our results exhibit the canonical properties of routines: efficacy, constancy, and organizational memory (e.g.,

<sup>29</sup>We show that the outcomes we observe arise under conditions where our modeling assumptions map to those described by the literature as leading to routine emergence. Although it is possible that other processes—for example, rational choice game theoretic models such as the Prisoner's Dilemma (Dawes, 1980; Rapoport and Chammah, 1965)—may result in similar outcomes, such processes would not be consistent with the behavioral assumptions made in the routines literature.

Cohen and Bacdayan, 1994; Grant, 1996; March, 1994; Nelson and Winter, 1982).

Our primary contribution is the claim that routines may be a source of systematic heterogeneity in the ability of firms to adapt to technological change. While the literature has classically emphasized the stability of routines (e.g., Cyert and March, 1963; Nelson and Winter, 1982) and implications for organizational inertia (e.g., Gersick and Hackman, 1990; Hannan and Freeman, 1984), over the last two decades the literature has come to renew the claim that routines can and do change (e.g., Feldman and Pentland, 2003; Knudsen, 2008). Yet, from a strategy perspective, the relevant questions are not whether routines can change, but rather: (1) Is there systematic heterogeneity in the adaptive capacity of routines to technological change; and (2) What are the managerial levers that can be used to alter the nature of this adaptive capacity? We suggest a candidate process through which such systematic heterogeneity in adaptive capacity might emerge, linking the theory of routines, the corresponding theory of inter-firm performance heterogeneity, and theories explaining adaptation in the context of technological change.

Our contribution to understanding firm response to technological change is based on taking a forward-looking perspective, with the idea that firms may take strategic steps to *prepare* for technological change. One approach to such preparation may be structural—for example, ensuring organizational ambidexterity (e.g., Stettner and Lavie, 2014) by physically isolating agile subunits (Tushman and O'Reilly, 1996) or maintaining flexibility in current operations (Gibson and Birkinshaw, 2004). Our results, which highlight how managerial policy well before a change occurs is central to building routines with the capacity to adapt, suggest a complementary approach. By utilizing alternative exploration policies in the formative stages of routine development, managers can build routines that are differentially adaptable to alternative types of technological change. The managerial cognitive challenge of such preparation is then not one of recognizing change after it occurs and responding appropriately (e.g., Eggers and Kaplan, 2009; Tripsas and Gavetti, 2000), but rather, one of forming forward-looking assessments of future change in order to grow routines *ex ante* that exhibit the appropriate adaptive characteristics.

Our work on routines and their adaptive capacity in the face of technological change takes a step forward in furthering our understanding of how firm capabilities might respond to change. Routines are characterized in the literature as underpinning capabilities (Winter, 2003), and capabilities are, more generally, intertwined with firms' ability to respond to different forms of technological change. This has implications for the broader theme in the strategy literature addressing how firms renew and reconfigure in response to discrete threats, replenishing capabilities that have been devalued by change (e.g., Aggarwal and Helfat, 2009; Knott and Posen, 2009; Simons, 1994; Teece, Pisano, and Shuen, 1997; Winter, 2003). In particular, it points to the importance of understanding the properties of the routines on which capabilities are built, and the path-dependent process through which they are formed.

Our theory also links work in the knowledge-based view (e.g., Grant, 1996; Kogut and Zander, 1992) with work on strategic human capital (e.g., Campbell, Coff, and Krzyscynski, 2012; Coff, 1997). The literature has long held that routines are "repositories and carriers of knowledge" (Hodgson, 2008: 25). We show that organizational memory is a natural outcome of the process through which routines emerge, and knowledge is thus not only in the minds of (and embodied in the habits of) individuals, but also in the connections between individuals. Consider employee turnover. We find that the exit of one individual does not undermine the performance of the routine. Even if the departing individual does not explicitly transfer his or her knowledge to the replacement person, the routinized behavior of the remaining individuals will lead him or her to select a task approach that resembles his or her predecessor. Thus, the knowledge embodied in the connections between individuals has the properties of tacit knowledge. Once formed, such knowledge is not subject to expropriation by individuals—it is inherently the property of the organization.<sup>30</sup>

Several aspects of our model offer opportunities for future extensions. In particular, we purposefully

<sup>30</sup>This raises the interesting question of what share of individuals must depart for the tacit knowledge to be eroded, and how the properties of the replacement impact this tacit knowledge. In work not reported here, but available on request, we have examined these questions, finding that as the share of departing individuals increases, the extent of organizational memory decreases. We also find that the stronger the preferences of the replacement (the extent to which his or her habits are strongly ingrained), the weaker is organizational memory.

abstract away from various organizational issues, including individual motivation, emotions, incentives, information asymmetry, heterogeneous utility functions, information transfer, hierarchy, and varying forms of interdependence across individuals. This suggests that some caveats may be in order regarding the inferences we can make about organizational routines with our model of three interdependent individuals. As with any modeling effort, however, our goal is not to offer a rich representation of organizations, but rather to illustrate a core mechanism that could underlie learning and behavior in real organizations. Future research might explore the degree to which enriching the stylized view of organizations we put forth in this article can result in a deeper understanding of the mechanisms we examine here. Additionally, a key boundary condition of our model is that we focus on architectural forms of technological change. Future work might examine the role of adaptive capacity in the context of nonarchitectural (e.g., component-related) technological change.

Beyond these modeling extensions, there are several additional avenues for future research. In the empirical domain, our notion of adaptive capacity as a property of routines and capabilities suggests that it may be fruitful for future empirical work to verify the existence, and examine the causes and consequences, of this property. There may also be opportunities to conduct experimental work. For example, our results suggest that learning by interdependent individuals leads to the emergence of local search. Would such an effect occur in a controlled lab setting? To the degree that learning is distributed among individuals, could we recover a reliable pattern of search in such a setting? These questions suggest a number of important avenues for future research that stem from our model and from the notion of adaptive capacity.

In sum, the microfoundational approach we pursue in this article furthers our understanding of the role of routines in the context of firm adaptation to technological change. It raises the interesting possibility that managerial policy during routines' formative period may greatly impact their adaptive properties in maturity. This, in turn, has implications for how managers prepare for the possibility of technological change. The challenge of adapting to such change is a central issue for firms and remains a fertile domain of inquiry for scholars of strategy.

## ACKNOWLEDGEMENTS

We thank Connie Helfat and two anonymous reviewers, as well as Adam Brandenburger, Russ Coff, Sendil Ethiraj, Henrich Greve, Henning Piezunka, Phanish Puranam, Maurizio Zollo; seminar participants at INSEAD, LMU-Munich, University College London, University of Texas-Dallas, and University of Wisconsin-Madison; and participants at the Academy of Management, Strategic Management Society, and Vienna Organization Design conferences, for valuable comments. We gratefully acknowledge support from the INSEAD Alumni Fund.

## REFERENCES

- Abernathy WJ, Utterback JM. 1978. Patterns of industrial innovation. *Technology Review* **80**(7): 41–47.
- Agarwal R, Helfat CE. 2009. Strategic renewal of organizations. *Organization Science* **20**(2): 281–293.
- Ahuja G, Katila R. 2004. Where do resources come from? The role of idiosyncratic situations. *Strategic Management Journal* **25**(8–9): 887–907.
- Arrow KJ. 1962. The economic implications of learning by doing. *Review of Economic Studies* **29**(3): 155–173.
- Barney JB. 1991. Firm resources and sustained competitive advantage. *Journal of Management* **17**(1): 99–120.
- Barney JB, Felin T. 2013. What are microfoundations? *Academy of Management Perspectives* **27**(2): 138–155.
- Becker MC. 2004. Organizational routines: a review of the literature. *Industrial and Corporate Change* **13**(4): 643–678.
- Becker MC, Lazaric N, Nelson RR, Winter SG. 2005. Applying organizational routines in understanding organizational change. *Industrial and Corporate Change* **14**(5): 775–791.
- Berry DA, Fristedt B. 1985. *Bandit Problems: Sequential Allocation of Experiments*, Monographs on Statistics and Applied Probability. Springer: New York.
- Bush RR, Mosteller F. 1951. A mathematical model for simple learning. *Psychological Review* **58**(5): 313–323.
- Campbell BA, Coff RW, Kryscynski D. 2012. Rethinking sustained competitive advantage from human capital. *Academy of Management Review* **37**(3): 376–395.
- Christensen CM. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Press: Boston, MA.
- Christensen CM, Bower JL. 1996. Customer power, strategic investment, and the failure of leading firms. *Strategic Management Journal* **17**(3): 197–218.
- Christensen CM, Rosenbloom RS. 1995. Explaining the attacker's advantage: technological paradigms, organizational dynamics, and the value network. *Research Policy* **24**(2): 233–257.
- Coff RW. 1997. Human assets and management dilemmas: coping with hazards on the road to resource-based theory. *Academy of Management Review* **22**(2): 374–402.
- Cohen MD, Bacdayan P. 1994. Organizational routines are stored as procedural memory: evidence from a laboratory study. *Organization Science* **5**(4): 554–568.
- Cohen WM, Levinthal DA. 1989. Innovation and learning: the two faces of R&D. *Economic Journal* **99**(397): 569–596.
- Cohen MD, Levinthal DA, Warglien M. 2014. Collective performance: modeling the interaction of habit-based actions. *Industrial and Corporate Change* **23**(2): 329–360.
- Cohen MD *et al.* 1996. Routines and other recurring action patterns of organizations: contemporary research issues. *Industrial and Corporate Change* **5**(3): 653–698.
- Cyert RM, March JG. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall: Englewood Cliffs, NJ.
- Dawes RM. 1980. Social dilemmas. *Annual Review of Psychology* **31**(1): 169–193.
- Denrell J, March JG. 2001. Adaptation as information restriction: the hot stove effect. *Organization Science* **12**(5): 523–538.
- Dewey J. 1922. *Human Nature and Conduct*. Henry Holt and Co.: New York.
- Dosi G. 1982. Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research Policy* **11**(3): 147–162.
- Eggers JP, Kaplan S. 2009. Cognition and renewal: comparing CEO and organizational effects on incumbent adaptation to technical change. *Organization Science* **20**(2): 461–477.
- Feldman MS, Pentland BT. 2003. Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly* **48**(1): 94–118.
- Felin T, Foss NJ, Ployhart RE. 2015. The microfoundations movement in strategy and organization theory. *Academy of Management Annals* **9**(1): 575–632.
- Foss NJ, Heimeriks KH, Winter SG, Zollo M. 2012. A Hegelian dialogue on the micro-foundations of organizational routines and capabilities. *European Management Review* **9**(4): 173–197.
- Foss NJ, Pedersen T. 2014. Microfoundations in strategy research. *Strategic Management Journal* Virtual Special Issue.
- Gavetti G, Levinthal DA. 2000. Looking forward and looking backward: cognitive and experiential search. *Administrative Science Quarterly* **45**(1): 113–137.
- Gersick CJ, Hackman JR. 1990. Habitual routines in task-performing groups. *Organizational Behavior and Human Decision Processes* **47**(1): 65–97.
- Ghemawat P. 1991. *Commitment*. The Free Press: New York.
- Gibson CB, Birkinshaw J. 2004. The antecedents, consequences, and mediating role of organizational ambidexterity. *Academy of Management Journal* **47**(2): 209–226.
- Gilbert CG. 2005. Unbundling the structure of inertia: resource versus routine rigidity. *Academy of Management Journal* **48**(5): 741–763.
- Gittins JC. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B Methodological* **41**(2): 148–177.

- Grant RM. 1996. Toward a knowledge-based theory of the firm. *Strategic Management Journal* **17**(S2): 109–122.
- Greve HR. 2008. Organizational routines and performance feedback. In *Handbook of Organizational Routines*, Becker MC (ed). Edward Elgar Publishing Limited: Cheltenham, UK; 187–204.
- Hannan MT, Freeman J. 1984. Structural inertia and organizational change. *American Sociological Review* **49**(2): 149–164.
- Helfat CE, Peteraf MA. 2003. The dynamic resource-based view: capability lifecycles. *Strategic Management Journal* **24**(10): 997–1010.
- Henderson RM. 2006. The innovator's dilemma as a problem of organizational competence. *Journal of Product Innovation Management* **23**(1): 5–11.
- Henderson RM, Clark KB. 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly* **35**(1): 9–30.
- Hodgson GM. 2008. The concept of a routine. In *Handbook of Organizational Routines*, Becker MC (ed). Edward Elgar Publishing Limited: Cheltenham, UK; 15–30.
- Holland JH. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan: Ann Arbor, MI.
- Hutchins E. 1991. Organizing work by adaptation. *Organization Science* **2**(1): 14–39.
- Knott AM, Posen HE. 2009. Firm R&D behavior and evolving technology in established industries. *Organization Science* **20**(2): 352–367.
- Knudsen T. 2008. Organizational routines in evolutionary theory. In *Handbook of Organizational Routines*, Becker MC (ed). Edward Elgar Publishing Limited: Cheltenham, UK; 125–151.
- Kogut B, Zander U. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science* **3**(3): 383–397.
- Leonard-Barton D. 1992. Core capabilities and core rigidities: a paradox in managing new product development. *Strategic Management Journal* **13**(2): 111–125.
- Levinthal DA, March JG. 1993. The myopia of learning. *Strategic Management Journal* **14**(S2): 95–112.
- Levitt B, March JG. 1988. Organizational learning. *Annual Review of Sociology* **14**: 319–340.
- Luce RD. 1959. *Individual Choice Behavior a Theoretical Analysis*. John Wiley and Sons: New York.
- March JG. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2**(1): 71–87.
- March JG. 1994. *Primer on Decision Making: How Decisions Happen*. Simon and Schuster: New York.
- March JG. 1996. Learning to be risk averse. *Psychological Review* **103**(2): 309–319.
- March JG, Olsen JP. 1975. The uncertainty of the past: organizational learning under ambiguity. *European Journal of Political Research* **3**(2): 147–171.
- March JG, Simon HA. 1958. *Organizations*. Wiley: New York.
- Miner AS, Ciuchta MP, Gong Y. 2008. Organizational routines and organizational learning. In *Handbook of Organizational Routines*, Becker MC (ed). Edward Elgar Publishing Limited: Cheltenham, UK; 152–186.
- Nelson RR, Winter SG. 1982. *An Evolutionary Theory of Economic Change*. Belknap Press: Cambridge, MA.
- Penrose ET. 1959. *The Theory of the Growth of the Firm*. Oxford University Press: New York.
- Pentland BT, Feldman MS. 2005. Organizational routines as a unit of analysis. *Industrial and Corporate Change* **14**(5): 793–815.
- Peteraf MA. 1993. The cornerstones of competitive advantage: a resource-based view. *Strategic Management Journal* **14**(3): 179–191.
- Pfeffer J, Salancik GR. 1978. *The External Control of Organizations: A Resource Dependence Perspective*. Harper & Row: New York.
- Posen HE, Levinthal DA. 2012. Chasing a moving target: exploitation and exploration in dynamic environments. *Management Science* **58**(3): 587–601.
- Puranam P, Alexy O, Reitzig M. 2014. What's "new" about new forms of organizing? *Academy of Management Review* **39**(2): 162–180.
- Puranam P, Stieglitz N, Osman M, Pillutla MM. 2015. Modelling bounded rationality in organizations: progress and prospects. *Academy of Management Annals* **9**(1): 337–392.
- Rapoport A, Chammah AM. 1965. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. University of Michigan Press: Ann Arbor, MI.
- Reinganum JF. 1983. Uncertain innovation and the persistence of monopoly. *American Economic Review* **73**(4): 741–748.
- Siggelkow N, Rivkin JW. 2006. When exploration backfires: unintended consequences of multilevel organizational search. *Academy of Management Journal* **49**(4): 779–795.
- Simons R. 1994. How new top managers use control systems as levers of strategic renewal. *Strategic Management Journal* **15**(3): 169–189.
- Stettner U, Lavie D. 2014. Ambidexterity under scrutiny: exploration and exploitation via internal organization, alliances, and acquisitions. *Strategic Management Journal* **35**(13): 1903–1929.
- Sutton RS, Barto AG. 1998. *Introduction to Reinforcement Learning*. MIT Press: Cambridge, MA.
- Taylor A, Helfat CE. 2009. Organizational linkages for surviving technological change: complementary assets, middle management, and ambidexterity. *Organization Science* **20**(4): 718–739.
- Teece DJ. 1982. Towards an economic theory of the multiproduct firm. *Journal of Economic Behavior and Organization* **3**(1): 39–63.
- Teece DJ. 1986. Profiting from technological innovation: implications for integration, collaboration, licensing and public policy. *Research Policy* **15**(6): 285–305.
- Teece DJ, Pisano G, Shuen A. 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* **18**(7): 509–533.
- Thompson JD. 1967. *Organizations in Action*. McGraw-Hill: New York.
- Tripsas M. 1997. Unraveling the process of creative destruction: complementary assets and incumbent survival in the typesetter industry. *Strategic Management Journal* **18**(S1): 119–142.

- Tripsas M, Gavetti G. 2000. Capabilities, cognition, and inertia: evidence from digital imaging. *Strategic Management Journal* **21**(10–11): 1147–1161.
- Tushman ML, Anderson P. 1986. Technological discontinuities and organizational environments. *Administrative Science Quarterly* **31**(3): 439–465.
- Tushman ML, O'Reilly CA III. 1996. Managing evolutionary and revolutionary change. *California Management Review* **38**(4): 8–28.
- Tushman ML, Romanelli E. 1985. Organizational evolution: a metamorphosis model of convergence and reorientation. In *Research in Organizational Behavior*, Cummings LL, Staw BM (eds). JAI Press: Greenwich, CT; 171–222.
- Winter SG. 2000. The satisficing principle in capability learning. *Strategic Management Journal* **21**(10–11): 981–996.
- Winter SG. 2003. Understanding dynamic capabilities. *Strategic Management Journal* **24**(10): 991–995.
- Winter SG. 2013. Habit, deliberation, and action: strengthening the microfoundations of routines and capabilities. *Academy of Management Perspectives* **27**(2): 120–137.
- Wu B, Wan Z, Levinthal DA. 2014. Complementary assets as pipes and prisms: innovation incentives and trajectory choices. *Strategic Management Journal* **35**(9): 1257–1278.
- Zollo M, Winter SG. 2002. Deliberate learning and the evolution of dynamic capabilities. *Organization Science* **13**(3): 339–351.

## SUPPORTING INFORMATION

**Additional supporting information may be found in the online version of this article:**

**Appendix.** Additional robustness analyses.