10.5465/AMBPP.2013.216

RUBIK'S DILEMMA: PARTIAL KNOWLEDGE AND THE EFFICACY OF LEARNING

HART E. POSEN
University of Wisconsin-Madison
Madison WI 53706
hposen@bus.wisc.edu

DIRK MARTIGNONI
University of Zurich, Switzerland

MARKUS LANG
University of Zurich, Switzerland

INTRODUCTION

A core assumption in the strategy literature is the positive cross-sectional relationship between a firm's knowledge and its performance (Nelson and Winter, 1982, Kogut and Zander 1992, Grant, 1996). Does this also imply that superior prior knowledge will engender enhanced learning and superior future performance? While popular managerial wisdom might typically answer in the affirmative, the strategy literature is more circumspect. Based on the assumption that boundedly rational firms engage in search and learning to find better solutions to the challenges they face (Cyert and March, 1963), the strategy literature highlights the benefits of prior knowledge (Gavetti and Levinthal, 2000, Rivkin ,2000, 2001, Agarwal, Echambadi, Franco, and Sarkar 2004, Ethiraj and Zhu, 2008, Dencker, Gruber, and Shah 2009, Csaszar and Siggelkow 2010, Gruber 2010), but also costs such as competency traps and cognitive rigidities (e.g., Levitt and March 1988, Levinthal 1997, Tripsas and Gavetti 2000). We argue that these opposing views on the implications of knowledge derive from the fact that knowledge has not one, but two mechanisms by which it alters future firm performance.

First, prior knowledge has an endowment effect. It provides a better starting position from which to engage in learning (e.g., head-start on the learning curve) and enhances the efficacy of new knowledge accumulation (Dierickx and Cool 1989, Cohen and Levinthal 1990). This is the basis of much popular managerial wisdom. Second, prior knowledge also has a behavioral effect because it alters the search strategy (heuristic) a firm employs in seeking to build upon and supplement its prior knowledge. By examining the interplay between the endowment and behavioral effects of knowledge, we seek to identify the conditions under which prior knowledge negatively (or positively) influences subsequent learning and future performance.

Consider imitative entry (e.g., Csaszar and Siggelkow, 2010 ; Ethiraj and Zhu 2008, Posen, Lee, and Yi 2013, Rivkin 2000, 2001) as an important instantiation of prior knowledge and its behavioral implications. A boundedly rational firm seeks to enter a market that it has not served in the past. The entrant carefully analyzes the current market leader's policy choices. This analysis yields the market leader's choices for some but not all of its policy dimensions (because of tacitness or causal ambiguity). For example, the entrant may successfully identify the leader's marketing strategy and pricing policy while other aspects of the market leader's approach remain unobservable to the entrant (e.g., R&D and production strategy). Thus, the entrant possesses pre-entry knowledge (Agarwal, Franco, Echambadi, and Sarkar 2004, Dencker, Gruber, and Shaw 2009) — a partial understanding of a good configuration of policy choices (Rivkin 2001) derived from its imitation of the observable attributes of a leading incumbent. We refer to this as prior knowledge, which follows a

10.5465/AMBPP.2013.216

continuum from fully incomplete (no prior knowledge) to complete (full prior knowledge). We refer to intermediate levels of prior knowledge as partial knowledge.

The entrant's partial prior knowledge of the market leader's policy choice configuration is the basis from which it engages in subsequent search and learning (Gavetti and Levinthal, 2000 ; Winter 2000, Rivkin 2001). How might an entrant's search behavior change given its prior knowledge endowment? The answer hinges on the observation that, having imitated a subset of the market leader's policy choices that are believed to be good, the entrant knows that it has partial prior knowledge. As such the entrant is likely to focus its search, restricting effort to the sub-problems for which it does not have solutions (e.g., R&D and production strategy in the example above).

On the surface, focus, as a response to partial knowledge, seems to be a reasonable heuristic. Focus is intuitively appealing because focusing economizes on search effort, allowing more exhaustive search in the domain of the remaining sub-problems. More importantly, focus is consistent with theoretical arguments for sequential attention to problems (Cyert and March 1963; Greve 2008; Baumann and Siggelkow, 2013) and organizational structure and process as a means to direct attention (Ocasio 1997, Rivkin and Siggelkow 2003). Recent empirical work finds explicit support for the idea that prior knowledge engenders focus. For example, in a study of technology entrepreneurs, Gruber, MacMillan, and Thompson (2012) show that pre-entry knowledge leads entrepreneurs to focus their subsequent search behavior, constraining the linkages between technologies and market opportunities that the entrepreneurs identify.

While, in the example above, we depict prior knowledge as deriving from imitative entry, our theory is more general. Imitation is one instantiation of prior knowledge. A broad body of research assumes, implicitly or explicitly, that prior knowledge enhances future performance. In these literatures, prior knowledge may accrue from acquisitions and alliances (Puranam, Singh, and Zollo 2006; Puranam, Singh, and Chaudhuri 2009), prior efforts at learning or pre-entry knowledge (Agarwal, Franco, Echambadi, and Sarkar 2004; Dencker, Gruber, and Shaw 2009; Ganco and Agarwal, 2009, Gruber 2010), employee mobility (Corredoira and Rosenkopf 2010), or analogical reasoning (Gavetti, Levinthal, and Rivkin 2005). Our theory is independent of the source of prior knowledge, requiring only that the prior knowledge is acquired, and the firm so endowed recognizes that it possesses such knowledge (so that it can focus).

To assess the implications of partial prior knowledge for subsequent performance, we construct a NK model of learning under complexity and add two features to the standard formulation (e.g., Levinthal 1997; Rivkin 2001). (a) We endow firms with prior knowledge in the sense that a subset of policy choices correctly matches the optimal configuration (Rivkin 2001). (b) We "inform" firms about the subset of correct policy choices such that "they know what they know" and we examine the implications of focused search on the remaining subset of policy choices. In doing so, we hold fixed issues of knowledge and change (inappropriate application of knowledge, inertia) and examine the implications of a focus heuristic.

Given our simple model, we demonstrate that when partial prior knowledge engenders a focus search heuristic, it can generate long-run performance inferior to no knowledge at all. We decompose the performance implications of prior knowledge into the knowledge-endowment effect and the knowledge-focus effect. In the absence of focus, performance is increasing in the completeness of prior knowledge, which is consistent with expectations of the role of knowledge endowments. The effect of focus is more complicated because focus interacts with knowledge endowments in non-obvious ways. Focus generates enhanced performance only when prior knowledge is relatively complete because focus ensures that the firm is not led astray — staying within the region of the best solution — and enhances the likelihood of finding the best solution. When prior knowledge is less complete, focus may lead to what we term a "behavioral impasse." By excluding certain policy configurations in

10.5465/AMBPP.2013.216

one period, other policy configurations are necessarily inaccessible in future periods. As a result, prior knowledge may be detrimental to long-run performance.

Why do we call this "Rubik's dilemma"? Consider a Rubik's Cube on which one face has six (out of nine) yellow tiles. These yellow tiles represent "correct" choices if that face is to be yellow - we call this the prior knowledge endowment. How does one proceed to complete the other faces of the cube and solve the puzzle? One strategy reflects a decision to hold fixed the six yellow tiles in all subsequent moves, only trying the range of alternatives that do not disturb this partial knowledge. Employing this focus strategy, finding a complete solution is (often) impossible. There are only two ways to proceed to solve the puzzle. The first is to abandon focus, and allow changes in the yellow face that reduce the number of yellow tiles temporarily, before reintroducing those correct tiles at a later point. Rubik's dilemma is embodied in the need to (at least temporarily) abandon what is currently known to be correct in order to solve the cube. The alternative strategy, available to firms but unavailable in a Rubik's Cube, is to increase search breadth by making multiple changes simultaneously. The central observation of this paper is that Rubik's dilemma is a common challenge of learning in complex task environments.

MODEL AND ANALYSIS

To examine the implications of (partial) prior knowledge, we implement a standard NK model (e.g., Kauffman 1993, Levinthal 1997, Rivkin 2001, Ethiraj and Levinthal 2004a, Siggelkow and Rivkin 2005, Levinthal and Posen 2007, Knudsen and Levinthal 2007, Ganco and Hoetker 2009; Csaszar and Siggelkow, 2010). It has three basic features: (1) a complex performance landscape, (2) a firm that is represented by a position on this performance landscape, and (3) a strategy that guides the search process a firms uses to learn and improve its position on the performance landscape.

The performance landscape maps firm policy choices to performance (fitness) where a firm is associated with a specific policy-choice vector in a given period. Firms seek to improve their positions on the landscape through a process of local search. In the standard NK model, initially firms have no information about the shape of the performance landscape; they start their search process from a random position on the landscape. In contrast, in our study firms are equipped with more- or less-complete prior knowledge about the policy-choice vector associated with the best solution (global peak). This knowledge affects both the starting position of the firm ("endowment effect") and its subsequent search behavior ("focus effect").

In the following, we report results for the case of a landscape with N=15 and K=7. Each experiment involves 10,000 firm replications. We observe firms for 200 periods, which is sufficient to ensure that the model reaches steady state.

In the first experiment, we seek to understand the baseline properties of the model. Figure 1 displays long-run performance (y-axis) over the full range of prior knowledge (x-axis), i.e., from no knowledge (zero policy choices correct), through partial prior knowledge, to complete prior knowledge (15 policy choices correct). The solid line reflects average long-run performance less than that achieved by firms with no prior knowledge. The dashed and dotted lines display the endowment effect and the focus effect, respectively, to which we return later.

Figure 1 about here

The solid line shows that performance is reduced rather than enhanced at low to moderate levels of prior knowledge; firms endowed with no prior knowledge outperform

10.5465/AMBPP.2013.216

firms with partial prior knowledge. Positive effects of prior knowledge materialize only when prior knowledge is relatively complete (i.e., $\gamma>9$). Thus, we observe a U-shaped relationship between the completeness of prior knowledge and long-run performance, with a minimum observed in the case of partial prior knowledge ($\gamma=7$).

Consider this result in the context of imitative entry (example from the introduction). The results of this experiment suggest that if the entrant cannot substantially imitate the market leader's policy choices, it may be better off not imitating at all. That is, attempts to learn, starting with relatively incomplete prior knowledge, may lead to outcomes inferior to entry without the benefit of prior knowledge. In the remainder of this section, we examine the mechanisms underlying the performance implications of prior knowledge.

A central behavioral assumption in this research is that a firm endowed with prior knowledge will focus its search efforts. We decompose (following Posen and Levinthal 2012) the impact of prior knowledge (solid line) into two components: (1) an endowment effect (dashed line) driven by the prior knowledge endowment alone, and (2) a focus effect (dotted line) driven by the assumption that a firm will restrict its subsequent search if it knows what it knows.

Consistent with Rivkin (2001), we find that a more-complete knowledge endowment has an unambiguous positive effect on long-run performance. This positive endowment effect is driven by two mechanisms. First, knowledge endowments (prior knowledge about the position of the best solution) might increase the probability that a firm discovers the best solution. Second, because better solutions (higher local peaks) tend to be co-located on the landscape (Kauffman, 1993; Rivkin and Siggelkow, 2007), knowledge endowments might help a firm discover better solutions.

The impact of focus is less straightforward (and more interesting). The behavioral implication of knowing what you know, focused search, decreases performance if the prior knowledge endowment is less-complete, but increases performance if the prior knowledge endowment is more-complete. This U-shaped relationship is the result of two opposing mechanisms. On the positive side, focus ensures that the firm is not led astray — it stays within the region of the best solution — and enhances the likelihood of finding the best solution. Even if the prior knowledge endowment is relatively complete (e.g., 12 out of 15 policy choices correct), local search in the absence of focus does not ensure that the firm finds the best solution. Unfocused search may (permanently) overturn correct policy choices and lead the firm to the basin of attraction of an average solution (which by definition is inferior to the best solution). When prior knowledge is relatively complete, focus reduces this risk. For example, at a prior knowledge of 12 correct policy choices, focus increases the probability of finding the best solution by over 60 percent.

On the negative side, focusing search effort on the sub-problems for which the firm does not yet have solutions may prevent a firm from converging to a local or global peak, because it comes to a "behavioral impasse." Focusing excludes certain policy configurations in one period, and by implication, other policy configurations are necessarily inaccessible in future periods. As a consequence, a firm at a behavioral impasse tends to find a solution that is inferior to the local (global) peak associated with its current basin of attraction. A firm at a behavioral impasse finds a "poor solution," with performance approximately 11 percent lower than that of a firm with an "average solution" (local peak).

Consider Microsoft and Dell's failed efforts to replicate Apple's successful iPod/iTunes business model. They obtain partial prior knowledge by imitating the observable features of Apple's policy choices. One interpretation of why Microsoft and Dell failed is that they did not fully imitate Apple's complete set of policy choices (Porter 1996). Given interdependence, incomplete imitation may result in dramatically lower performance. Our model suggests an alternative explanation because we assume that the initial imitation attempt is not the final solution employed by the imitator. Rather, Microsoft and Dell's partial prior

10.5465/AMBPP.2013.216

knowledge obtained via imitation is the starting point for their subsequent effort at search and learning by which they attempt to reconstruct the remainder of Apple's policy choices.

It is commonly understood that Apple's close integration of hardware (iPod) and software (iTunes) was one of the key elements of their success. Microsoft and Dell engaged in search, but were reluctant to abandon this integrated business model, even temporarily (e.g., Dell's Digital Jukebox and Microsoft's Zune). Both Microsoft and Dell get stuck at behavioral impasses, unable to replicate the remainder of Apple's policy choices via local search. Moreover, because of their focus response to partial knowledge, Microsoft and Dell failed to identify other policy configurations, such as Spotify's software-based product, which may have provided a better solution to the one they identified.

Thus, our model suggests that Microsoft and Dell did not fail in digital music because they persisted with outdated or incorrect knowledge, or imperfectly imitated a successful template. Rather, our model suggests that the opposite may be true – Microsoft and Dell performed poorly in digital music because they failed to (at least temporarily) abandon "correct" beliefs. This observation is the heart of "Rubik's dilemma" - sometimes one must forgo partial solutions that are known to be correct, in order to find a path to a more complete and better solution.

In sum, in this experiment, we find a U-shaped relationship between the level of prior knowledge and long-run performance. This effect is driven by two interacting mechanisms, the positive effect of knowledge endowments, and the effect of focus, which is negative when knowledge is less complete and positive when knowledge is more complete.

CONCLUSION

We formally consider how the completeness of a firm's initial endowment of knowledge affects its learning process and subsequent performance. At the core of our analysis rests a simple proposition: A firm's prior knowledge endowment alters its subsequent search strategy and ultimately its prospects for additional learning. We examine one instantiation of a behavioral effect of knowledge, focus. Implicit is the assumption that a firm "knows what it knows," and as such, naturally focuses its search effort on the remaining unknown dimensions of the problem. This focus heuristic is not only intuitively appealing but also consistent with theoretical arguments for sequential attention to problems (Cyert and March 1963; Greve 2008, Baumann and Siggelkow, 2013), and organizational structure as a means to direct attention (Ocasio 1997, Rivkin and Siggelkow 2003).

In a strategic setting, prior knowledge may have many sources. One obvious example is imitation (Ethiraj and Zhu, 2008; Rivkin 2000, 2001). In entering a new market, a firm might seek to imitate the current market leader. In most instances, the entrant can only observe select aspects of the market leader's approach. It seems natural that a firm will focus its subsequent search efforts on those aspects that were not easily imitable. Yet sources of prior knowledge extend well beyond imitation to include: acquisitions and alliances (Puranam, Singh, and Zollo 2006 ; Puranam, Singh, and Chaudhuri 2009), prior efforts at learning or pre-entry knowledge (Agarwal, Franco, Echambadi, and Sarkar 2004; Dencker, Gruber, and Shaw 2009; Ganco and Agarwal 2009; Gruber 2010), employee mobility (Corredoira and Rosenkopf 2010), or analogical reasoning (Gavetti, Levinthal, and Rivkin 2005). Our theory is independent of the source of prior knowledge, requiring only that the prior knowledge is acquired, and the firm so endowed recognizes that it possesses such knowledge (so that it can focus its subsequent search efforts).

We find that there are conditions under which prior knowledge can generate long-run performance inferior to no knowledge at all. Our result hinges on the implications of focus, when a firm has partial knowledge. In the absence of focus the benefit of prior knowledge unambiguously increases with the extent of knowledge completeness. This result is the basis

10.5465/AMBPP.2013.216

of the endowment effect of prior knowledge. While this result conforms to popular managerial wisdom, it is only part of the story. When prior knowledge engenders a focus heuristic, the future performance implications of prior knowledge are less straightforward. The focus heuristic positively contributes to future performance of prior knowledge when knowledge is relatively more complete, but negatively when prior knowledge is less complete.

Why does the focus heuristic have such a large impact on the future performance implications of prior knowledge? Our study suggests that focus is a double-edged sword. On the positive side, focus ensures that the firm is not led astray — it stays within the region of the best solution — and enhances the likelihood of finding the best solution. On the negative side, focus forecloses pathways for search and learning. This foreclosure occurs because of the path dependent nature of search (Arthur, 1994). Focus excludes certain policy configurations in one period, and as a consequence, this implies that other policy configurations are necessarily inaccessible in future periods. This may prevent a firm from converging to a local or global peak, because it comes to a "behavioral impasse." The former effect dominates at high levels of prior knowledge, while the latter effect dominates at low levels of prior knowledge.

Our results point to the need to enhance our understanding of the implications of current knowledge for future performance. If knowledge is behaviorally neutral, not affecting a firm's subsequent search strategy, then more-complete prior knowledge is certainly better than less-complete prior knowledge (e.g., Nickerson and Zenger, 2004 ; Miller, 2007). Yet the assumption of behavioral neutrality in the face of more-complete knowledge seems unlikely given strong evidence that strategies depend on current knowledge states (Kahneman and Tversky, 1979; Audia and Greve, 2006). Under such circumstances, partial knowledge may reflect what is popularly called "dangerous half truths," suggesting that we need to take a more critical look at the behavioral and performance implications of knowledge.

FIGURE 1 Decomposing the Effect of Prior Knowledge

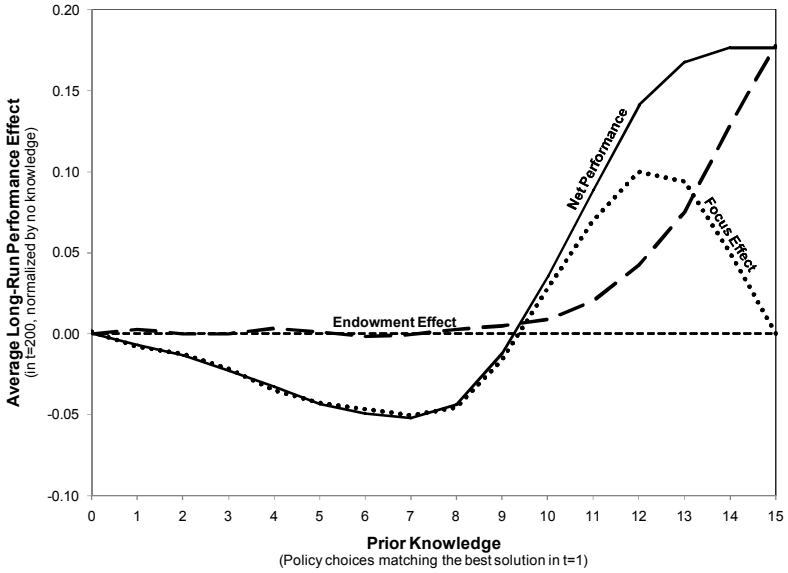

REFERENCES AVAILABLE FROM THE AUTHOR(S)