

Mental Representation and the Discovery of New Strategies¹

Felipe A. Csaszar

Ross School of Business, U. of Michigan

Ann Arbor, MI 48109

fcsaszar@umich.edu

Daniel A. Levinthal

The Wharton School, U. of Pennsylvania

Philadelphia, PA 19104

levinthal@wharton.upenn.edu

August 25, 2015

Forthcoming in the Strategic Management Journal

¹The authors thank Rich Bettis and two anonymous referees for their insightful and constructive reviews. The authors also thank the seminar participants at the Academy of Management Meeting (2013) and the Organization Science Winter Conference (2014) for useful comments and suggestions. We also wish to acknowledge some early conversations with Giovanni Gavetti that helped to set the stage for this work.

Mental Representation and the Discovery of New Strategies

Abstract

Managers' mental representations affect the perceived payoffs and alternatives that managers consider. Thus, mental representations affect how managers search for profitable strategies as well as the quality of strategies they discover. To study how mental representation and search interact, we formally model the dual search over possible representations and over policy choices of a strategy 'landscape.' We analyze when it is preferable to emphasize searching for the best policies rather than the best mental representation, and vice versa. We show that in the long run a balance between the two search modes not only results in better expected performance but also reduces the variation in performance. Additionally, the paper describes conditions under which increased accuracy of mental representations can actually worsen firm performance.

Keywords: managerial cognition; mental representation; strategy frameworks; search; NK models

“The voyage of discovery is not in seeking new landscapes but in having new eyes.”

—Marcel Proust

1 Introduction

1.1 Representations in strategy

The strategy literature has offered two disparate research traditions that address strategic decision making by boundedly rational actors. One line of work, which builds on March and Simon (1958/1993), Cyert and March (1963), and the broader ‘Carnegie school’ points to the role of search and trial-and-error learning. A different tradition, one that developed explicitly within the strategy field, highlights the role that mental representations play in strategic decision making (see, e.g., Huff 1990, Prahalad and Bettis 1986, Walsh 1995). However, as suggested by Gavetti and Levinthal (2000), it is important to link the “backward-looking” perspective of the former approach with the “forward-looking” perspective of the later approach.

A *mental representation* is a model of reality held in the mind of an individual, who can use this representation to generate predictions about reality (Craik 1943:61; Holland et al. 1986:12).¹ Mental representations are especially important in strategy because they allow managers to consider alternative strategies in an ‘off-line’ manner (Gavetti and Levinthal 2000). That is, mental representations allow managers to consider the merit of alternative strategies without the need to actually invest in and carry out the various options. Of course, such ‘modeling’ is inevitably imperfect and different representations may offer better or worse characterizations of the likely payoffs in an actual business context. Thus the quality of managers’ mental representations is a basis for performance differences (Barr et al. 1992, Gary and Wood 2011), and a firm’s prior history may have a considerable influence on the mental representations adopted by their managers (Benner and Tripsas 2012).

A fundamental characteristic of mental representations is that they account for only some dimensions of the represented reality (Brunswick 1952, Gärdenfors 2000, Hong and Page 2009). In strategy, the reduced dimensionality of mental representations helps explain why ‘frameworks’ are

¹A basic premise of cognitive science is that thinking can be understood in terms of mental representations and operations on these representations (Thagard 2005:10–12).

pervasive in the field. Frameworks—which provide simple, typically two dimensional representations of the complex underlying strategy context—suggest the dimensions that a mental representation should include. For instance, the Boston Consulting Group (BCG) growth–share matrix suggests that corporate strategy decisions should be based on the dimensions of market share and market growth. A framework is useful to the manager because it reduces the high dimensionality of strategic problems and so makes them tractable within the bounds of the manager’s cognitive capacity. Because frameworks incorporate only a subset of the problem’s actual dimensions, multiple frameworks can apply to a given strategic problem. For instance, if in reality there were 20 strategic dimensions yet frameworks considered just two of them at a time, then there would be $\binom{20}{2} = 190$ possible two-by-two frameworks. This multiplicity of possible representations may explain the profusion of strategy frameworks.²

Thus, an important problem that managers face is to choose the dimensions of the representation—the ‘lens’ or ‘filter’—through which to view their business landscape. That any landscape looks much different under different representations exposes a problem with how the literature has conceptualized the process of searching for a strategy. Models of search (e.g., Lenox et al. 2006, Levinthal 1997, Rivkin 2000) assume that managers search by changing the value of a given set of policies with the goal of increasing the firm’s fitness (e.g., an automobile company’s manager may request changes to engineering specifications or to the firm’s organizational structure so that cars can be produced more profitably). However, these models ignore that the search over policies amounts to only part of devising a new strategy: which policies should be changed depends on the mental representation used by the manager. For example, a manager who viewed autos in terms of safety and mileage dimensions would produce different cars than a manager who viewed them in terms of speed and style. Similarly, a manager who designed a strategy based on the dimensions in the BCG matrix would design a different strategy than one who thought in terms of a price-versus-quality positioning frontier.

Therefore, a characterization of the search for more or less appropriate strategies should incorporate exploring not only the space of alternative policies but also the space of possible representations under which these policies might be evaluated. Or, in terms of frameworks: searching for a strategy requires managers to search for a framework as well as to search for policies given

²For example, without aiming at being comprehensive, Krogerus et al. (2012) describe 50 strategic frameworks.

that framework. One complexity of this dual search process is that there is a trade-off between the two search processes: because managers' time is limited, spending more time searching for representations will reduce the time available to search for policies. A second complexity is that the two search processes are intertwined: not only policies will be seen differently depending on the representation used, but how good a representation appears to be will depend on which policies were considered when that representation was tried.

1.2 Our contribution

Our main objective in this paper is to understand the dual search process over representations and policy choices. Specifically, we examine the conditions under which it is better to emphasize searching for the right dimensions of the mental representation rather than searching for the right policies, and vice versa. Understanding this dual search process provides important practical and theoretical contributions.

From a practical standpoint, our paper sheds light on how managers should deal with the trade-off between searching for policies and searching for representations; that is, whether managers should think creatively about how to represent a strategy problem or whether they should just stick to the current understanding of the problem, and try to find ways to improve performance as suggested by the current representation. While strategy scholars have noted the importance of the role of representations (Benner and Tripsas 2012, Huff 1990, Tripsas and Gavetti 2000), with few exceptions (Gavetti and Levinthal 2000, Martignoni and Siggelkow 2012) the relationship between shifts in representations and search for more or less appropriate strategies has not been developed. These distinct approaches are paralleled in the practitioner literature, which is divided between those emphasizing the search over possible representations—such as the advocates of design thinking in business (Martin 2009) and blue ocean strategy (Kim and Mauborgne 2005)—and those emphasizing the search over policies, such as efforts at continuous improvement (Imai 1986) and specification of activity systems (Porter 1996). Our research sheds light on the interplay between these two modes of search as a function of characteristics of the managers and the environment.

In terms of specific results, modeling the interplay between searching for alternative representations and for specific policy values yields a number of interesting implications. First, the nature of this trade-off depends crucially on the time horizon being considered. In the near term, performance

is maximized by focusing one's energies on the identification of superior policies irrespective of the given representation's merit. Over longer time horizons, however, it is beneficial to balance the search for policies and representations. In the extremely long run, if the competitive context is stable then performance is largely invariant to the mix of search over policies and representations—provided the search process does not focus exclusively on one of those two.

Second, we describe situations in which increasing the accuracy of mental representations is detrimental to performance. These cases run counter to previous work and to the conventional wisdom, by which more accurate representations are *always* preferable (e.g., Barr et al. 1992, Gary and Wood 2011). Third, our analysis offers a novel insight on the relationship between the variability of outcomes and the expected performance of a given search strategy. For a broad range of conditions we show that balancing search across representations and policies improves performance. The reason is that doing so protects managers against assessing policies in terms of a less informative representation but also against assessing representations on a less effective set of policies. As a consequence, search strategies that are more balanced lead to higher average performance as well as less variability in realized outcomes. This offers an alternative mechanism for explaining Bowman's (1980) paradox of the negative relationship between a firm's risk and its returns.

From a theoretical standpoint, our paper illuminates how mental representation influences the process of strategy search. Processes of search have been central in the strategy literature (see, e.g., Levinthal 1997, Lenox et al. 2006). The main focus of this literature has been to identify how search can lead to novel policy configurations (see, e.g., Rivkin and Siggelkow 2007, Sommer and Loch 2004), yet less emphasis has been placed on studying how alternatives are evaluated (Knudsen and Levinthal 2007) and how initial representations guide subsequent search for policy configurations (Gavetti and Levinthal 2000). An important, unstudied question is what is the effect of searching over representations. Our work shows how search over representations is intertwined with the search over policy configurations. Among other implications, our research offers a new perspective on exploration and exploitation. We show that exploration is the outcome of two fundamentally distinct but interrelated search processes: search over representations (i.e., finding how to represent the strategy problem) and search over policies (i.e., finding specific choices given that representation). The trade-off in the allocation of time and energy to one or the other of these search processes

does not map onto a basic exploration/exploitation trade-off. For instance, one can engage in “local” search in the space of representations, but that incremental change in representation may have an extensive effect in the valuation of alternative search paths. Moreover, the payoff to one form of search may depend on the intensity of the other form of search and, thus, there may be complementarities between the two modes of search.

In the next section we establish the theoretical connections between the literatures on search and on mental representation. The subsequent section develops a formal model of search and mental representation. We then present the results that stem from our model, and we build on those results to discuss some broader theoretical and managerial implications of this research.

2 Theoretical motivation

This section has two goals: (1) to show that mental representation plays a central yet often overlooked role in organizational theories of search; and (2) to discuss key determinants of search performance when the role of mental representations is taken into account.

2.1 Mental representation as an overlooked aspect of search

Central to the concept of bounded rationality (Simon 1955) is the now-familiar idea that organizational decision making is better described as a search process than as an optimization process. In most realistic settings, a firm cannot make optimal decisions given the information and cognitive limitations of its members; instead, firms *search* for a satisfactory alternative among the limited set of choices available to them at the time each decision is made.

March and Simon (1958/1993) make this perspective on search central to how they characterize organizational behavior. These authors, and subsequently Cyert and March (1963), develop the concept of search by elaborating on how search is affected by (among other things) political conflict, uncertainty, and aspiration levels. More recent work in the Carnegie tradition has made extensive use of the concept of search in the modeling of organizations (e.g., Csaszar and Siggelkow 2010, Lenox et al. 2006, Levinthal 1997, March 1991, Rivkin 2000).

In the early 1960s, the bulk of research activity involving search processes moved from the field

of organizations to the then-nascent fields of artificial intelligence and cognitive science.³ Within these two fields, mental representation was incorporated as a central element of search models (Newell and Simon 1976). Prior to this time, classic psychological work had assumed that search was performed in the actual problem space—as when a mouse (or human subject) searches for a reward in a T-maze experiment. But in the developing new fields of artificial intelligence and cognitive science, research was based on presuming that search occurred not in the problem space but rather in a mental representation of that space (Newell and Simon 1976:125).

The concept of mental representation is important because it is a key determinant of the outcomes of search. The effect of representation on search was eloquently anticipated by Whitehead (1911:59) when discussing how mathematical notation helps one discover mathematical theorems: ‘[b]y relieving the brain of all unnecessary work, a good notation sets it free to concentrate on more advanced problems, and, in effect, increases the mental power of the race.’ In other words, problems that are difficult under one representation (e.g., multiplying two roman numerals) may well become easy under a different representation (multiplying with Arabic numerals). Research on mental representations has continued in the work on knowledge representation (Brachman and Levesque 2004) and mental models (Gentner and Stevens 1983, Johnson-Laird 1983).

So when considering the problem of search, it is important to distinguish the challenge of identifying more or less appropriate actions from the challenge of identifying more or less appropriate ways to assess the strategic situation. Models of learning typically emphasize the former type of challenge. For instance, research drawing on the idea of fitness landscapes has highlighted the challenge of combinatorial search problems and the need to identify more or less appropriate configurations of policy choices (Levinthal 1997, Rivkin and Siggelkow 2007). Similarly, bandit models provide a canonical characterization of the exploration–exploitation trade-off (Holland 1975, Posen and Levinthal 2012) and examine the updating of beliefs about the quality of alternative policies. However, the fundamental representation of the problem environment—what constitutes the set of possible actions and how performance is measured—is treated as being understood by the actor.

The challenge of searching for representations is often neglected not only in the academic

³To a large extent, these newer fields also descended from Simon’s work on search. In the artificial intelligence field, Newell et al. (1959) were the first to program a computer to solve problems via search; in cognitive science, Simon (1956) was the first to propose search as the mechanism by which individuals make decisions.

literature but also in practical strategic frameworks. For instance, when strategy frameworks are invoked (e.g., the characteristic “two-by-two’s” used in the MBA classroom), a representation is provided with no explicit consideration of whether some alternative representation might yield more insight.

Gavetti and Levinthal (2000) consider the role of cognitive representations in guiding the search process, but their characterization does not include intrinsically better or worse representations (the representations they consider are aggregates of a given NK landscape that do not exhibit systematic differences). Thus, while Gavetti and Levinthal study some of the effects of using a limited representation, the structure of their model precludes them from addressing the question of a search over more or less effective representations.

In sum, the central role that mental representation plays in search has remained mostly unexplored in the strategy literature. Identifying effective strategies requires managers to consider the criteria or dimensions along which strategies should be viewed and also to identify specific strategies. The challenge of selecting evaluative criteria for strategies is no less important and demanding than the problem of identifying those strategies. Because mental representation affects search in fundamental ways, we believe it is important to connect the research on mental representation and on search within the field of strategy.

2.2 Some determinants of search when mental representation is taken into account

The preceding arguments call for incorporating mental representations into models of organizational search. Thus, the logical next step is determining how best to accomplish this goal. Here we consider the main elements that such a model should include. Deciding what does (or does not) belong in the model is not straightforward, since concepts that may be useful in artificial intelligence or cognitive science need not be relevant in the context of strategy making (and vice versa).

Because there are so many elements that can affect a model of mental representation and search, our aim will not be to produce an exhaustive list of relevant elements but instead to describe a minimal set of elements that are important in the context of our research. Because a mental representation depends on elements of the actual decision context faced by the actors involved, we first characterize these contexts and then characterize their mental representations.




Figure 1: Schematic relationship between policies, performance dimensions, and profits.

2.2.1 Characteristics of decision contexts in strategy. We describe the strategy context in terms of the three-element structure described by Adner et al. (2014), which comprises (i) policies, (ii) performance dimensions, and (iii) firm profits. Policies are the levers or choices that managers can change, performance dimensions are the product or strategy characteristics that drive the market's response, and firm profits are an overall measure of that market response. For instance, the manager of a restaurant can choose policies such as staffing levels, the quality of ingredients, and the recipes used (element (i)); these policies affect performance dimensions such as food quality and restaurant ambiance (element (ii)); and how the market responds to these performance dimensions vis-à-vis other alternatives determines the restaurant's profitability (element (iii)).

Figure 1 illustrates the relationship between policies, performance dimensions, and profits. This diagram depicts the sequential relationship between these elements: the performance dimensions of a product are a function of its policies, whereafter the profits from that product are a function of its performance dimensions. We can express this connection more formally, in terms of function composition, as follows: if Performance dimensions = $f(\text{Policies})$ and Profits = $\pi(\text{Performance dimensions})$, then Profits = $\pi(f(\text{Policies}))$.

Models in the fields of economics, marketing, and strategy usually focus on performance dimensions and profits (elements (ii) and (iii) in the setup of Adner et al. 2014). For example, Christensen's (1997) work on disruptive technologies can be understood as highlighting how different customer segments value performance dimensions differently (e.g., how weight and power consumption are extremely important to laptop users but virtually irrelevant for users of desktop computers). A growing literature on product design (e.g., Baldwin and Clark 2000, Sommer and Loch 2004) and strategy (e.g., Adner et al. 2014) has augmented the analysis by incorporating the configuration of

business policies (element (i)), highlighting the interdependence of these policy choices.

The model presented here includes the three elements of policies, performance dimensions, and profits. In the context of our paper, including these three elements is important because the first element captures the traditional idea of search over policies (as in Levinthal 1997), the second allows us to model the idea of search over representations, and the third element corresponds to the outcome variable that we use to measure the quality of different ways of searching over both policies and representations.

For the sake of clarity, many of our examples describe product strategy decisions. That being said, the theory we develop is not bound to product strategy and can be applied equally well to other strategy decisions. In the context of a new globalization strategy, for instance, policies might correspond to strategic levers that a manager can change (e.g., organizational structure, hiring policies, acquisition decisions) while performance dimensions correspond to antecedents of profits that the market takes into account, but that are not necessarily direct levers that managers can change (e.g., each of Porter's five forces in different markets, the position of the firm vis-à-vis competitors in a price-versus-quality frontier, or the firm's power in the global supply chain).

2.2.2 *Characteristics of mental representations.* Cognitive science has developed a number of different hypotheses concerning the nature of mental representations. The three main hypotheses are that mental representations consist of images, rules, or connections; Cummins (1989) traces these respective hypotheses back to the work of Aristotle, Hobbes, and Hebb. An emerging consensus is that these hypotheses describe not competing but rather complementary means of representation (Minsky 1986:66; Thagard 2005:20). Thus the mind may use different representation types depending on the nature of what is being represented. For example: the layout of a building is more likely to be represented in terms of images; the conditions for passing a course are more likely to be represented as rules; and tacit knowledge, such as how to recognize a face, is more likely to be represented by connections stemming from associative learning (Hebb 1949).

A key aspect of all these types of mental representations is their use of a simplified version of reality to guide future actions. In the words of Craik, a pioneer in the modern study of mental representation: "If the organism carries a 'small-scale model' of external reality and its own possible actions within its head, [then] it is able to try out various alternatives, conclude which is the best

of them, [and] react to future situations” (1943:61). This simplified or small-scale nature implies that mental representations are not faithful copies of reality, but less accurate depictions of it. The idea that managers operate on the basis of inaccurate information is one of the hallmarks of bounded rationality; according to Simon (1997:17), “bounded rationality [...] assumes that the decision maker [...] has egregiously incomplete and inaccurate knowledge about the consequences of actions.”

Our model, as described in the next section, is faithful to the concept of mental representation in that it characterizes actors’ actions as being guided by an incomplete and thus inaccurate version of their decision context. In particular, managers do not make decisions based on the real effect of policies on profits (i.e., on the true value of π); in our model, they instead make decisions based on their mental representation of that relationship (i.e., on a mere approximation of π , which we denote π').

3 Model

Our model describes the process by which a manager searching on a mental representation can find a new strategy. The model has three basic components: a model of the decision context (i.e., a description of how policies affect performance dimensions and of how performance dimensions affect profits); a model of the mental representation (i.e., a presumably informative but inevitably inaccurate and incomplete representation of the decision context); and a search process by which the manager searches among possible strategies. The search process occurs off-line, so a strategy’s actual performance is not known until after the search has ended (Gavetti and Levinthal 2000); this dynamic is akin to how the performance of a new product or venture cannot be determined until after its launch.

We capture this process by extending the model presented in Levinthal (1997). Namely, we augment that model with multiple performance dimensions and agents whose mental representations need not incorporate all of these dimensions. In Levinthal (1997) and subsequent work drawing on the structure of NK fitness landscapes (see, e.g., Lenox et al. 2006, Rivkin and Siggelkow 2003), a fitness landscape describes how N policy choices determine *one* fitness level. In contrast, our model associates the N policy choices to M fitness levels. For instance, in Levinthal (1997) an $N = 3$

landscape could establish that policies $(1, 0, 1)$ map to fitness 0.9; in contrast, here a multidimensional landscape with $N = 3$ and $M = 2$ could establish that policies $(1, 0, 1)$ map to fitnesses 0.9 and 0.4, where each of these two fitness values corresponds to a different performance dimension. For example, if $(1, 0, 1)$ were the policies describing a automobile (e.g., large chassis, small engine, thick walls) then 0.9 and 0.4 could correspond to two of the car’s performance dimensions—say, high safety and low speed. The mapping from policies to each performance dimension is modeled by M different NK landscapes (all of size N and complexity K). Because each NK landscape connects the value of policies to the value of a performance dimension, these landscapes play the role of a technology. We refer to K as *complexity*.

Profitability depends on the product’s performance dimensions. For instance, the profits of the car mentioned above could be given by $0.5\textit{safety} + 0.25\textit{speed}$. More generally, our model includes a parameter w ($0 < w \leq 1$) that captures the relative relevance of each performance dimension for firm profits. For example, if there are $M = 3$ dimensions and $w = 0.5$, then the first most important performance dimension would have a weight of 0.5, the second most important a weight of 0.25 ($= 0.5^2$), and the third most important a weight of 0.125 ($= 0.5^3$). As w approaches unity, the performance dimensions’ weights become more similar.

Managers do not necessarily have a complete understanding about what performance dimensions drive profits. We therefore assume that managers are aware of only M' of the M performance dimensions (where $M' \leq M$).⁴ For instance, if $M' = 1$ then a manager might believe that the car’s profits depend only on speed. For the sake of succinctness, from now on we normally refer to performance dimensions in the manager’s mental representation simply as ‘dimensions’; this word usage is consistent with the literature on mental representation (e.g., Brunswik 1952, Gärdenfors 2000).

Because managers can search for products only by way of their lower-dimensional mental representation of the landscape, they must choose how to allocate their search efforts between dimensions and policies: at one extreme, they could search for the most valuable dimensions while keeping the policy choices fixed; at the other extreme, they could keep the set of dimensions fixed

⁴It is useful to contrast this notion of a low-dimensional representation with that developed by Gavetti and Levinthal (2000), in which actors are aware of only a subset of the policies ($N' < N$) and there is only one performance dimension. In contrast, we allow for a variety of possible performance dimensions of which only a subset ($M' < M$) are known to the actors.

while searching for the policy choices that maximize the value of those dimensions. The allocation of search effort between these two extremes is the main independent variable of our analysis. We denote this variable e_p (for *effort in policy search*) and examine its full range from 0 to 1. When $e_p = 0$, search focuses entirely on dimensions; when $e_p = 1$, search focuses entirely on policies. Intermediate values interpolate between these extremes.⁵

The parameter e_p is modeled as a probability: at each time step, with probability e_p the manager performs local search among the product policies and with probability $1 - e_p$ the manager performs local search among the dimensions. To keep the dimensions considered within the limit of the manager’s cognitive capacity, when exploring dimensions the manager will add one dimension and remove some other dimension—thereby keeping the number of considered dimensions constant at M' . The manager uses a particular mental representation while searching for a total of T time periods; following this search, the product is launched and its actual performance becomes known. Thus search proceeds not in the real profit landscape (π) but rather in the manager’s mental model of it (π'); however, the product’s actual performance depends on the real profit landscape π . A formal description of the search process is given in the Appendix.

In summary, our model is parameterized by three characteristics of the manager—the search strategy (e_p), the manager’s cognitive limit (M'), and the length of the time horizon allocated for search (T)—and by two characteristics of the environment: its complexity (K) and the relative relevance of the dimensions (w). The main outcome we explore is the performance of different search strategies in different environments. That is, we ask: How effective are different search strategies (e_p ’s) at discovering products that exhibit high performance? We define *performance* in terms of the average profitability of the actual product found by a search strategy in a given environment. We also study the heterogeneity of performance by looking at the cross-sectional variability in performance among firms using the same search strategy. In order to facilitate comparisons across different simulations, performance is scaled so that the global minimum performance is 0 and the global maximum is 1. To ensure that our results are reliable and not merely an artifact of a particular sample, we run 50,000 simulations per scenario (i.e., per combination of parameter levels) and report the average performance per scenario. All reported results are statistically significant at

⁵Our model thus incorporates a trade-off between searching for dimensions and searching for policies given that both searches require the manager’s time.

the 1% level.

4 Results

We use the model to study the performance of different search strategies (e_p) under different time horizons (T), environments (w, K), and cognitive limits (M'). To simplify the exposition, we structure the presentation around a series of plots that are representative of the model's behavior and start by describing the most intuitive effects first.

Among other results, we find that: (i) emphasizing search over representations is detrimental if not much time is available to search; (ii) the value of searching over representations increases with complexity; (iii) using a more accurate mental representation can be detrimental to performance; and (iv) search strategies that lead to high expected performance also lead to low levels of cross-sectional and temporal heterogeneity in performance.

4.1 The best search strategy depends on the available time

We start by analyzing the effect of time spent searching (T). Figure 2 plots real performance (π) for two different search strategies ($e_p = 0.5$ and $e_p = 0.9$) in a representative scenario ($w = 0.8$, $M' = 4$, $K = 3$). This figure can be viewed as depicting the performance a manager would obtain after having spent T time periods searching in her mental representation. Recall that the manager searches in her mental representation of the profit function (π'), but her performance is given by the real profit function (π), akin to how a manager would design a product, to only know its real performance once the product is launched.

A first observation from Figure 2 is that performance increases with the time spent searching. This improvement happens because as long as the mental representation has some resemblance to reality, then searching in the mental representation is no worse than random search, and thus has a better than 50-50 chance of increasing performance. This above-chance probability accumulates over time, which accounts for the observed improvement in performance. Eventually, all the possible benefit of searching in the mental representation will have been realized, whereafter continuing to search in that mental representation has no appreciable effect on performance.

A second observation from Figure 2 is that the preferred search strategy will depend on the time




Figure 2: Comparison of two different search strategies ($e_p = 0.5$ and $e_p = 0.9$) as a function of time.

available for searching. In the short run (i.e., $T < 30$ periods), the search strategy that places more emphasis on policies ($e_p = 0.9$) outperforms the search strategy that puts less emphasis on policies ($e_p = 0.5$); the opposite claim holds in the longer run (i.e., when $T \geq 30$). The reason is that, if there is not much time to improve on the mental representation of a given design, then making too many changes to the mental representation is risky. Suppose, for example, that a manager has only 10 time periods to design a car and that, from $T = 1$ to $T = 8$, he believes that car performance depends on safety; if at $T = 9$ the manager changes his mind and starts thinking that performance depends instead on speed, then he will have just a single time period in which to align product policies to this new mental representation. In that case, the resulting automobile would probably not excel either at safety or speed. Thus it would likely have been better simply to maintain the safety-oriented mental representation even if it turns out that safety was *not* the most relevant dimension, for then the auto would have excelled in that dimension at least.

Conversely, if there is enough time to make adjustments (i.e., $T \geq 30$ in Figure 2), then changing the mental representation more frequently (i.e., using a lower e_p) may have the advantage of allowing the manager to identify the most relevant dimensions and then to choose policies that yield high performance along those dimensions.

Figure 3 sheds additional light on how performance is contingent on time available to search.

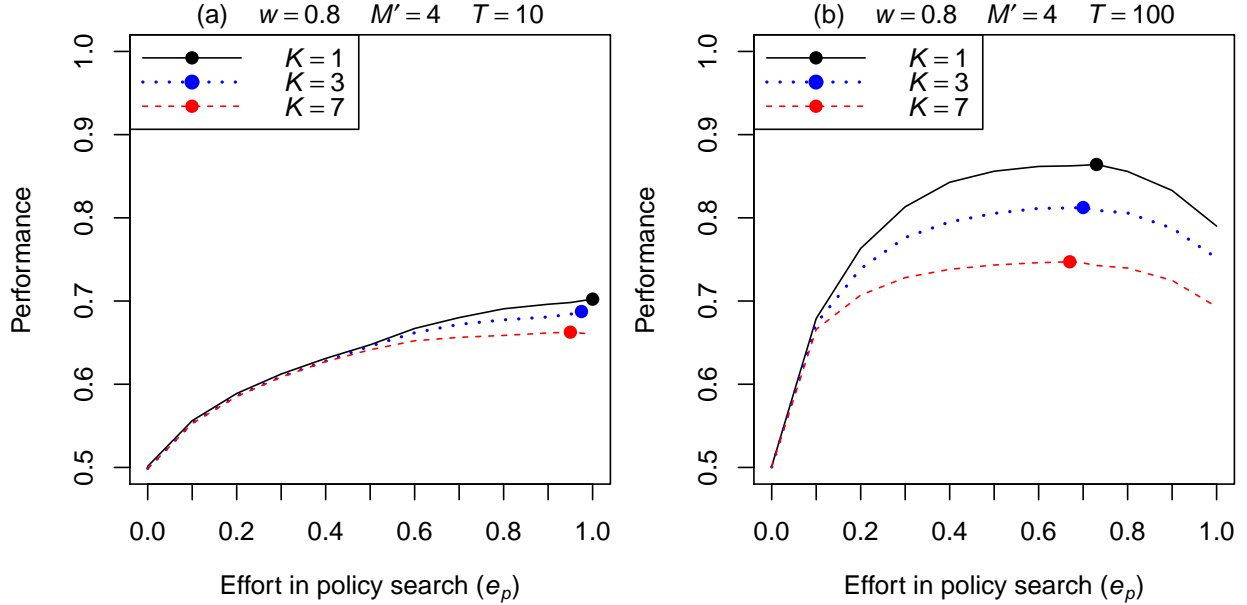


Figure 3: Average performance as a function of effort in policy search (e_p) and complexity (K) for a given scenario ($w = 0.8$, $M' = 4$) in the short and medium run ($T = 10$ and $T = 100$ in panels (a) and (b), respectively). The dots (●) mark the maximum performance (e_p^*) of each curve.

This figure will be analyzed later in more detail; for now, we focus on two aspects of the plots. First, performance increases with e_p in the short run (panel (a), $T = 10$); but second, in the longer term (panel (b), $T = 100$) performance follows an inverted U-shape with respect to e_p . We have already explained the short-run performance illustrated in panel (a): performance increases with e_p because searching over mental representations is risky if the search time is limited.

The inverted U-shape in the longer run (Figure 3(b)) is explained as follows. On the one hand, if the manager devoted search effort exclusively to representations ($e_p = 0$) then the firm would never change the policies affecting performance; all energy would be directed toward seeking the best way to frame a problem, yet no experimentation would take place with respect to alternative solutions or initiatives. On the other hand, if the manager just focuses the search effort on policies (i.e., $e_p = 1$), it is likely that a problem framing would be missed that could identify a more profitable product; in this case, the firm might not be attending to the dimensions that matter most. In short: balancing efforts between search over representations and search over policies is only useful if there is enough time to search; otherwise, it is better for the firm to focus on searching over policies.

Our final analysis of how the time horizon affects search is based on Figure 4, which differs




Figure 4: Percentage of the M' dimensions in the manager’s mental representation that are among the M' most important dimensions in reality.

from Figure 2 in two ways. First, Figure 4 considers a much greater time span (until $T = 2000$, by which time all curves have plateaued). Second, the y -axis signifies not performance but rather a key antecedent of performance: the extent to which different search strategies identify the most important dimensions. We quantify that extent as the percentage of the M' dimensions in the mental representation that are among the M' most important dimensions in reality—in other words, the degree to which the mental representation is as close to reality as the cognitive limit of the manager permits.

There are two striking observations from Figure 4. One is that, in the very long run ($T > 1000$) the performance of all but the two extreme search strategies converge (see, e.g., the plots for $e_p = 0.5$ and $e_p = 0.9$ in Figure 4). This result reflects that, in order for a manager to discover which are the most important dimensions, she needs to collect information about the expected contribution of the different dimensions; and given a sufficiently long time horizon, any search strategy that varies both dimensions and policies (i.e., $0 < e_p < 1$) can yield enough samples to estimate these contributions.⁶ The second notable observation from Figure 4 is that searching only over representations (i.e.,

⁶This approach is akin to estimating the expected payoffs of the arms of a bandit (Holland 1975, Posen and Levinthal 2012) by randomly sampling the payoffs of each arm: given sufficient time, any sampling strategy that samples each arm with at least some probability will deliver correct estimates.

setting $e_p = 0$) is actually less effective at identifying a representation’s most important elements than are mixed search strategies. This result follows because each mental representation is tested on just one, unchanging set of policies, which severely limits how much the manager can learn about the various dimensions’ possible contributions to profits.

Because managers do not have unlimited time to come up with a new strategy, the speed at which different search strategies achieve high performance is paramount. In fact, many strategies are designed under time pressure, as when a firm must respond rapidly to a competitor or when a strategy must be designed during a meeting or company retreat. So despite our theoretical interest in the equifinality of long-term search strategies, they have little relevance in practice. In the real world, search strategies that can make the most of a limited time horizon will be much preferred. Hence the ensuing analyses focus on the short and medium term, which we define as (respectively) $T = 10$ and $T = 100$.⁷

4.2 Complexity calls for emphasizing the search over dimensions

We now revisit Figure 3 so that we can better understand the effect of complexity (K , plotted as different curves) in the short and medium run (panels (a) and (b), respectively). An initial observation from this figure is that as K increases, performance decreases (in each panel, the height of the curves decreases with K). The reason is that, as K increases, the mental representation (π') becomes a more ‘rugged’ landscape and so the search process is more likely to end up at a local peak than at the global peak. This detrimental effect of K accords with previous research on NK landscapes (see, e.g., Levinthal 1997).

Because firms typically cannot change the problem’s complexity (K) but can change their search strategy (e_p), a managerially relevant question is what is the optimal search strategy for different values of K . To address this question, in Figure 3 we mark (using heavy dots) the position of the optimal search strategy, which is denoted e_p^* .

Figure 3 shows that increasing complexity calls for less search over policies (i.e., e_p^* moves to the left in the graph as K increases). To understand why this happens, recall that there is a trade-off between searching over policies and searching over dimensions. In particular, reducing the search

⁷The results are not qualitatively sensitive to the exact values used for T . In essence, $T = 10$ characterizes the model dynamics before the crossover point described in Figure 2 whereas $T = 100$ characterizes the model dynamics after that crossover point but before the equifinality that occurs in the very long term (i.e., $T \geq 1000$).

over policies allows the time-constrained manager to increase the search over dimensions. That shift, in turn, makes search less likely to become ‘stuck,’ as a local peak under one representation will probably not be a local peak under another representation. Since the likelihood of getting stuck is increasing in K , the usefulness of searching over dimensions also increases with K . In this sense, searching over the dimensions of a mental representation has an effect that is analogous to that of performing a ‘long jump’ in a traditional NK landscape (Levinthal 1997): both mechanisms can dislodge search from a local peak.

4.3 Incomplete representations are sometimes preferable

Next we explore the effect of varying the cognitive limit (M'). In practice, there are two ways of changing this limit: (1) by hiring managers who are capable of considering more dimensions simultaneously (managers can differ along this characteristic because of differences in working memory or experience; see, e.g., Helfat and Peteraf 2015) and/or (2) by using frameworks designed to consider more dimensions—for instance, a framework like Blue Ocean Strategy (Kim and Mauborgne 2005) considers more dimensions than does a framework like Porter’s (1980) generic strategies.

The main question that we address here is whether one should always use a representation that incorporates as many dimensions as possible. Conventional wisdom holds that considering more dimensions is better, but we shall demonstrate that this generalization does not always hold. Sometimes, incomplete representations are preferable to complete representations.

Figure 5 shows the effect on performance of increasing the cognitive limit (M') in two settings that are identical except for the relative relevance of the dimensions (w). In panel (a) of the figure, dimensions decline in relevance at a rate of 50% ($w = 0.5$); in panel (b), all dimensions are equally relevant ($w = 1$). It is noteworthy that in panel (a) maximal performance happens at an intermediate level of M' , while in panel (b) maximal performance happens when M' is maximal.⁸

The key to understanding this result is to acknowledge that increasing the number of considered dimensions has a double-edged effect. On the one hand, a mental representation that takes more dimensions into account is more accurate and therefore corresponds more closely to the true strategic

⁸Keeping the other parameters in Figure 5 fixed, any value of w below 0.7 produces an inverted U-shape (i.e., akin to panel a), while any value above that threshold produces a monotonically increasing relationship (i.e., akin to panel b).




Figure 5: The effect of varying the cognitive limit M' (x -axes) depends on the relative relevance w (different panels) of the respective performance dimensions.

context. On the other hand, such representations are harder to explore: local search has more opportunities to get stuck in the nooks and crannies of a more elaborate representation. Figure 6 illustrates these counteracting effects. Suppose, that the landscape in panel (a) of this figure is a more accurate representation of reality than the landscape in panel (b). Although the former is more accurate, the global peak of the latter will be easier for local search to discover. In the panel (a) landscape, search is likely to get stuck at one of the many local peaks; however, most points in the panel (b) landscape fall within the basin of attraction that leads to the global peak.

If we bear this imagery in mind, the explanation for Figure 5 becomes more apparent. In the $w = 0.5$ case graphed in Figure 5(a), it is preferable to use an intermediate M' because the least relevant dimensions add very little to the representation's accuracy and so including them would unnecessarily complicate the search for appropriate policies. For instance, the least significant dimension in Figure 5(a) has a weight of 0.002 ($= 0.5^9$), a value that is unlikely to change the policy configuration associated with the global peak; yet accounting for this dimension produces a number of 'wrinkles' (increased ruggedness) in the mental representation, which may leave the search stranded at one of many local peaks. However, in the $w = 1$ case of Figure 5(b) it is preferable to maximize M' . Here all dimensions are equally relevant and, thus, none can be omitted (toward the




Figure 6: Intuition behind the beneficial effect of reducing the cognitive limit M' .

end of streamlining search) without a significant reduction in the representation's accuracy.

We can summarize these results as follows. When determining the optimal M' , one must account for the trade-off between a representation's *accuracy* (higher M' leads to greater accuracy) and its *searchability* (higher M' leads to more ruggedness, which hinders search). Accuracy is less important than searchability if the strategic context contains many dimensions of low relevance (i.e., when w is not high).

4.4 Search strategies that lead to high performance and low variability

The goal of this paper is to examine the dual search process over representations and policy choices. The analyses so far have used expected performance as the main depend variable; however, expected performance does not paint a whole picture of the effect of a given search strategy. Two firms may pursue the same search strategy yet end up performing quite differently because of random events (e.g., different starting positions or different orders in which dimensions and policies are tried). Thus, we now study heterogeneity of performance by measuring the cross-sectional variability in performance among a group of firms that search under the same conditions (i.e., they differ

only in the simulation’s random aspects, not in the value of the model parameters).⁹ In practical terms, it is important to study how performance heterogeneity is affected by increasing search over representations, as management techniques that promote strategic innovation (such as Blue Ocean Strategy) are often justified by pointing to cases where these techniques were successful. But without knowing what happens to a population of similar firms that engaged in a similar exercise—a population that almost certainly includes some unsuccessful firms—such recommendations are unfounded. The simulation method that we employ is well suited to the study of performance heterogeneity because it is limited neither by sampling size nor survivor bias.

Apart from deepening the understanding of the effects of the dual search process, our analysis of performance heterogeneity provides a novel explanation for the inverse risk–return relationship observed empirically (Bowman 1980).

Figure 7 plots the expected performance (on the x -axis) and performance heterogeneity (the y -axis) of different search strategies (different values of e_p , drawn as points on each curve) and different cognitive limits (different values of M' , drawn as different curves). The other parameters are kept fixed at a representative setting ($T = 100$, $K = 3$, and $w = 0.8$), as the effect of varying them does not qualitatively affect the reported results. In the figure, points that are closer to the lower right corner are preferable: they correspond to high performance reliably achieved with low heterogeneity.

A first observation from Figure 7 is that performance heterogeneity decreases as M' increases. This happens because, with increasing M' , a manager’s mental representation becomes closer to reality and so managers are more likely to behave in a similar fashion—that is, in ways that are more consistent with reality than with (various) imagined versions of it. One implication is that performance will be most heterogeneous when the cognitive limit is much smaller than the actual number of dimensions in the strategic context (i.e., when $M' \ll M$).

A second observation from this figure is that performance heterogeneity exhibits a U-shape with respect to the search strategy e_p : regardless of M' , heterogeneity is highest when $e_p = 0$; as e_p

⁹More formally, let $\pi_{i,t}$ denote the performance of firm i after searching for t time periods; we study the standard deviation of $\pi_{i,t}$ across i for a fixed t . Results are derived from analyzing 50,000 simulated firms that are identical in every respect (i.e., of equal e_p , T , M' , w , and K) *except* for the landscape encountered and the search decisions they make, which are randomly drawn (as explained in the Model section). Unlike the concept of risk in finance (which considers variability across t for a given firm i ; Markowitz 1952), our measure captures a population-level construct: the variability among outcomes in a cohort of similar firms.




Figure 7: Performance (x -axis) and performance heterogeneity (y -axis) for a given environment ($K = 3$, $w = 0.8$, $T = 100$) as a function of cognitive limit M' and search strategy e_p (which is varied along each curve from 0 to 1 in increments of 0.1).

increases, heterogeneity first declines and then increases. The reason is that, when $e_p = 0$, search consists only of changing dimensions; thus firm performance is given by whatever policies the firm had in place when it began searching (i.e., performance will have as much heterogeneity as there is heterogeneity in initial policies). At the other extreme, when $e_p = 1$, search consists only of changing policies and so performance is fundamentally constrained by whatever the firm's representation was when it began searching. Both extreme settings result in considerable heterogeneity in performance owing to the legacy effect of representations ($e_p = 0$) or policies ($e_p = 1$). However, $e_p = 0$ leads to even greater heterogeneity as the performance of such firms depends entirely on (random) initial conditions whereas the $e_p = 1$ firms can at least experiment with alternative policies (though the experiments will always be assessed using the same criteria). It is only the intermediate levels of e_p that give firms leeway to adapt both their policies and the criteria by which those policies are evaluated.

A final observation from Figure 7 is the negative relationship plotted between the performance of each search strategy and the heterogeneity it produces: all points in the graph fall roughly along a -45 -degree line. In practice, this means that higher-performing firms will tend to exhibit the least performance heterogeneity.

This last result is relevant because—given that firms periodically search for new strategies—low heterogeneity translates into low risk.¹⁰ This negative relationship between risk and performance is contrary to the risk–return trade-off commonly postulated in finance (e.g., Sharpe 1964); however, it is consistent with the empirical finding of Bowman (1980), which has come to be called the ‘Bowman paradox.’ Several mechanisms have been proposed to explain this paradox.¹¹ Our model proposes an additional one: the Bowman paradox may emerge from the dynamics of search in mental representations because the mechanism that reduces heterogeneity (i.e., choosing a search strategy closer to e_p^*) is the same mechanism that increases expected performance. In other words, high performance and low risk tend to co-occur because they are both driven by the same mechanism: the capability of managers to search effectively within their mental representations.

5 Discussion

Here we discuss the broader implications of our research on mental representation in strategy. In particular, we show that acknowledging the central role that mental representation plays in strategy affects our understanding of strategy search as well as other key theories used in strategy. We conclude by proposing ways in which the research on mental representation in strategy can be furthered.

5.1 How mental representation affects strategy search

Strategy making inevitably requires dealing with an abstract representation of the strategic context. Strategists must not only identify the key success factors underlying current competitive positions, they must also consider hypothetical and cognitively distant competitive positions if the goal is to capture superior opportunities (Gavetti 2012). Thus strategists must engage in a dual search process: a search over alternative representations and a search over possible policies for a given

¹⁰That is, reiterated searches by one firm using a low-heterogeneity search strategy will produce less longitudinal performance heterogeneity (i.e., risk) than will reiterated searches by a firm using a high-heterogeneity search strategy.

¹¹One explanation uses prospect theory (Kahneman and Tversky 1979) and the behavioral theory of the firm (Cyert and March 1963, March and Simon 1958/1993) to argue that poorly performing firms might engage in riskier search strategies (Fiengenbaum and Thomas 1988). Another explanation, developed by Bettis (1981) and Bettis and Mahajan (1985), suggests that the risk–return relationship reflects synergies resulting from the focal firm’s degree of diversification, with firms engaging in closely related diversification exhibiting a negative relationship between risk and return. Other explanations have shown that the negative risk–return relationship is consistent with simple stochastic models of firm adaptation (Andersen et al. 2007, Andersen and Bettis 2015).

representation. The literature on search focuses mainly on the latter as it examines the search for alternative policies under a fixed representation.

Our modeling of the interplay between searching for policies and for representations produces a number of insights. First, which search type should be emphasized depends on how much time is available for search. If not much time is available, then performance is maximized by focusing one’s energies on the identification of superior policies (i.e., regardless of the possible merit of other representations); in contrast, if more time is available then balancing the search between policies and representations is beneficial. The explanation for this finding is that a change in representation helps the firm dislodge itself from a local peak—but that doing so is worthwhile only if enough time remains to find a better peak. One managerial implication of this result is that firms should not engage in search processes that alter the criteria for appraisal *unless* there is plenty of search time available. A corollary is that using frameworks that account for more dimensions (or hiring managers capable of considering more dimensions) should be accompanied by ensuring that managers have enough time to search for new strategies.

A second insight derived from our analysis is that, contrary to conventional wisdom and previous research (Barr et al. 1992, Gary and Wood 2011), there are situations in which less accurate mental representations are preferable to more accurate ones. Our explanation for this result is that accurate mental representations are more difficult to search. Hence a representation’s optimal level of accuracy depends on a trade-off between accuracy and ease of search. In particular, extremely accurate representations can be detrimental when there is considerable dispersion in the different dimensions’ relevance (i.e., when w is low). This result implies that there may be negative returns to experience and to highly nuanced insights about the strategic context—even in stable environments. This result offers an alternative rationale for why ‘simple rules’ (Bingham and Eisenhardt 2011) and heuristics (Gigerenzer and Goldstein 1996) are common in many business settings. In practical terms, this calls for managers to avoid incorporating into their mental representations dimensions that matter little. This result highlights the fact that representations yield power for the elements that they exclude as well those that they include.

Finally, our analysis points to an interesting relationship between the variability of outcomes and the expected performance under a given search strategy. We show that, if there is enough time to search for strategies, then a balanced search strategy not only results in a higher expected

performance but also—by mitigating the legacy effect of initial representations and policies—reduces the variation in performance associated with a given search strategy. This offers an alternative mechanism to explain Bowman’s (1980) paradox.

5.2 How our research illuminates other strategy theories

A relevant question in strategy is how to control the effectiveness with which a firm explores new opportunities (March 1991). Answering this question requires understanding the process underlying exploration. The initial literature on strategy search equated *exploitation* with performing local search in the business landscape and *exploration* with performing random jumps (usually called “long jumps”; Levinthal 1997:938). According to this view, a firm can increase its degree of exploration by increasing the extent to which it introduces random variation in its search process. However, these views on exploration do not shed much light on the behavioral process underlying exploration. Following Adner and Levinthal (2008), we propose that exploration may be better understood as not simply a particular draw from a distribution characterizing how “noisy” the search process is, but as a consequence of search along different dimensions.

The current paper points to search over mental representations as an additional process that can underlie exploration. What an outside observer may see as a long jump may actually be the outcome of small changes to the mental representation used by the manager. Moreover, the effectiveness of this process can be controlled by managers (by matching the extent of the search over representations with characteristics of the environment and the manager, as shown in the Results section). Thus, this paper advances our understanding of exploration by describing how effective exploration can result from a process that managers can control. At a more general level, this paper suggests that one way to advance toward answering questions such as “where does exploration come from?” and “what does determine the direction and effectiveness of exploration?” is to further study how managers search over mental representations.

Our conceptualization of search over mental representations also sheds light on the problem-solving perspective (Nickerson and Zenger 2004), which studies how different governance forms affect the ability of firms to solve problems. A later incarnation of this theory (Nickerson et al. 2007) highlights that strategy problems are not given—they are found—and points out that not much is known about the problem-finding process. Our paper contributes to the problem-solving perspective

by showing how formulating a strategy problem (i.e., choosing a representation) is intertwined with solving a problem (i.e., choosing which policies to implement). Moreover, our paper describes the situations under which problem-finding will be more or less relevant vis-à-vis problem-solving.

Our research also complements previous work on how the modular structure of task design impacts search processes. Previous research has studied the degree to which a modular structure is a more or less accurate reflection of the true interdependencies among tasks (Ethiraj and Levinthal 2004, Rivkin and Siggelkow 2003, Dosi et al. 2003). However, while a given modular structure may reflect a belief about interdependencies among tasks, previous research on modular structures has not considered whether some performance dimensions are more or less important and how do managers represent the value of these dimensions. Our research suggests that the process that leads to a modular structure should be as much determined by the manager's understanding of the interdependence among tasks (i.e., a perceived interdependence matrix), as by the manager's representation of how value is created (i.e., a perceived profit function).

Finally, our conceptualization of search over mental representations provides a more general understanding of why search strategies that rely on temporally ignoring dimensions can be successful. Levinthal and Posen (2007) show that ignoring for a period of time some interdependencies in the business landscape may be beneficial to performance; similarly, Ethiraj and Levinthal (2009) show that focusing on a few performance metrics at any given time leads to better performance than trying to maximize all performance metrics at once. In terms of our paper, ignoring parts of the problem and changing which parts to consider is akin to searching in a lower-dimensional representation of the problem. As explained in the context of Figure 6, lower-dimensional representations are helpful when they make search less likely to get stuck as long as they guide the search toward a high-performing policy configuration.

5.3 Further steps toward integrating mental representation into strategy

This paper presents a simple model of how firms search both for representations and for policies that fit those representations, and our findings constitute a first step in uncovering the central role that mental representation plays in the search for new strategies. We believe that this important topic offers a number of opportunities for future research. For instance, future research could explore how a strategy's quality depends on the fit among the mental representations that managers

use, the environments they face, and the aggregation process employed to combine their opinions (Csaszar and Eggers 2013). Further work could also study the circumstances that trigger managers to change representations as well as whether some managers are systematically better than others at picking representations. Furthermore, closer attention could be paid to the role of frameworks in strategy. Although frameworks are used extensively by consultants and managers, not much is known about how various characteristics of frameworks affect the strategic decision-making process. Understanding frameworks in terms of their effect on mental representation and search provides a new way to address this concern.

The quest for profitable competitive positions is a core problem of business strategy. To describe these competitive positions, practitioner-oriented contributions have introduced myriad frameworks. Yet these frameworks largely ignore the vastness of the design problem that strategists face: in reality, there are an arbitrarily large number of dimensions to which a strategist could attend (e.g., activities to include in the activity system or dimensions to consider regarding a new product or strategy). In other words, there has been insufficient acknowledgement of the degree to which frameworks are simplified, small-world representations of an actual problem environment (Levinthal 2011). Therefore, devising a strategy is a challenge that consists largely of searching the vast space of representations to find a ‘small world’ that has analytical purchase on real-world performance.

Separately from these practitioner-oriented contributions, a long-standing line of work (dating from March and Simon 1958/1993) has given prominence to the role of search processes; however, that line of work has not devoted much attention to the actual representations used by managers. We have drawn inspiration from how frameworks are used in practice and also from the search and representation ideas of the Carnegie tradition in order to establish a firmer bridge between these practitioner- and theory-oriented traditions within the strategy field.

The space of possible representations is huge, so more or less insightful representations may have important effects on the firm’s capacity to identify a more profitable competitive position. At the same time, the firm must establish a profitable set of policy choices. Those tasks compete for top management time and attention, and they also affect each other. Couched in these terms, the tools of strategic analysis are the tools of cognitive representations and heuristic search. By recognizing this connection, researchers can more closely link practitioner and theoretical approaches, thereby advancing the field in terms of both rigor and relevance.

References

- Adner, R., F. A. Csaszar, P. B. Zemsky. 2014. Positioning on a multi-attribute landscape.
- Adner, R., D. A. Levinthal. 2008. Doing versus seeing: Acts of exploitation and perceptions of exploration. *Strategic Entrepreneurship Journal* **2**(1) 43–52.
- Andersen, T. J., R. A. Bettis. 2015. Exploring longitudinal risk-return relationships. *Strategic Management Journal* **36**(8) 1135–1145.
- Andersen, T. J., J. Denrell, R. A. Bettis. 2007. Strategic responsiveness and Bowman’s risk–return paradox. *Strategic Management Journal* **28**(4) 407–429.
- Baldwin, C. Y., K. B. Clark. 2000. *Design Rules, Volume 1: The Power of Modularity*. MIT Press, Cambridge, MA.
- Barr, P. S., J. L. Stimpert, A. S. Huff. 1992. Cognitive change, strategic action, and organizational renewal. *Strategic Management Journal* **13** 15–36. Special Issue.
- Benner, M. J., M. Tripsas. 2012. The influence of prior industry affiliation on framing in nascent industries: The evolution of digital cameras. *Strategic Management Journal* **33**(3) 277–302.
- Bettis, R. A. 1981. Performance differences in related and unrelated diversified firms. *Strategic Management Journal* **2**(4) 379–393.
- Bettis, R. A., V. Mahajan. 1985. Risk/return performance of diversified firms. *Management Science* **31**(7) 785–799.
- Bingham, C., M. Eisenhardt. 2011. Rational heuristics: The ‘simple rules’ that strategists learn from process experience. *Strategic Management Journal* **32**(13) 1437–1464.
- Bowman, E. H. 1980. A risk-return paradox for strategic management. *Sloan Management Review* **21**(3) 17–31.
- Brachman, R. J., H. J. Levesque. 2004. *Knowledge Representation and Reasoning*. Morgan Kaufmann, Boston, MA.
- Brunswik, E. 1952. *The Conceptual Framework of Psychology*. University of Chicago Press, Chicago, IL.
- Christensen, C. M. 1997. *Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston, MA.
- Craik, K. J. W. 1943. *The Nature of Explanation*. Cambridge University Press, Cambridge, UK.
- Csaszar, F. A., J. P. Eggers. 2013. Organizational decision making: An information aggregation view. *Management Science* **59**(10) 2257–2277.
- Csaszar, F. A., N. Siggelkow. 2010. How much to copy? Determinants of effective imitation breadth. *Organization Science* **21**(3) 661–676.
- Cummins, R. 1989. *Meaning and Mental Representation*. MIT Press, Cambridge, MA.
- Cyert, R. M., J. G. March. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.
- Dosi, G., D. A. Levinthal, L. Marengo. 2003. Bridging contested terrain: Linking incentive-based and learning perspectives on organizational evolution. *Industrial and Corporate Change* **12**(2) 413–436.
- Ethiraj, S. K., D. A. Levinthal. 2004. Bounded rationality and the search for organizational architecture: An evolutionary perspective on the design of organizations and their evolvability. *Administrative Science Quarterly* **49**(3) 404–437.
- Ethiraj, S. K., D. A. Levinthal. 2009. Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science* **20**(1) 4–21.
- Fiegenbaum, A., H. Thomas. 1988. Attitudes toward risk and the risk-return paradox: Prospect-theory explanations. *Academy of Management Journal* **31**(1) 85–106.
- Gärdenfors, P. 2000. *Conceptual Spaces the Geometry of Thought*. MIT Press, Cambridge, MA.
- Gary, M. S., R. E. Wood. 2011. Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal* **32**(6) 569–594.
- Gavetti, G. 2012. Toward a behavioral theory of strategy. *Organization Science* **23**(1) 267–285.

- Gavetti, G., D. A. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* **45**(1) 113–137.
- Gentner, D., A. L. Stevens. 1983. *Mental Models*. Cognitive science, Erlbaum, Hillsdale, NJ.
- Gigerenzer, G., D. G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* **103**(4) 650–669.
- Hebb, D. O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York.
- Helfat, C. E., M. A. Peteraf. 2015. Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal* **36**(6) 831–850.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Holland, J. H., K. J. Holyoak, R. E. Nisbett, P. Thagard. 1986. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA.
- Hong, L., S. E. Page. 2009. Interpreted and generated signals. *Journal of Economic Theory* **144**(5) 2174–2196.
- Huff, A. S. 1990. *Mapping Strategic Thought*. Wiley, Chichester, NY.
- Imai, M. 1986. *Kaizen: The Key To Japan's Competitive Success*. Random House Business Division, New York.
- Johnson-Laird, P. N. 1983. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. No. 6 in Cognitive science series, Harvard University Press, Cambridge, MA.
- Kahneman, D., A. Tversky. 1979. Prospect theory: Analysis of decision under risk. *Econometrica* **47**(2) 263–291.
- Kim, W. C., R. Mauborgne. 2005. *Blue Ocean Strategy*. Harvard Business School Press, Boston, MA.
- Knudsen, T., D. A. Levinthal. 2007. Two faces of search: Alternative generation and alternative evaluation. *Organization Science* **18**(1) 39–54.
- Krogerus, M., R. Tschäppeler, J. Piening. 2012. *The Decision Book: Fifty Models for Strategic Thinking*. W.W. Norton & Co., New York.
- Lenox, M. J., S. F. Rockart, A. Y. Lewin. 2006. Interdependency, competition, and the distribution of firm and industry profits. *Management Science* **52**(5) 757–772.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7) 934–950.
- Levinthal, D. A. 2011. A behavioral approach to strategy: What's the alternative? *Strategic Management Journal* **32**(13) 1517–1523.
- Levinthal, D. A., H. E. Posen. 2007. Myopia of selection: Does organizational adaptation limit the efficacy of population selection? *Administrative Science Quarterly* **52**(4) 586–620.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2** 71–87.
- March, J. G., H. A. Simon. 1958/1993. *Organizations*. John Wiley & Co., New York.
- Markowitz, H. 1952. Portfolio selection. *Journal of Finance* **7**(1) 77–91.
- Martignoni, D., N. Siggelkow. 2012. Consequences of misperceiving external and internal interdependencies: The effects of neuroticism and blind spots. Working paper, Wharton School.
- Martin, R. L. 2009. *The Design of Business: Why Design Thinking is the Next Competitive Advantage*. Harvard Business Press, Boston, MA.
- Minsky, M. L. 1986. *The Society of Mind*. Simon and Schuster, New York.
- Newell, A., J. C. Shaw, H. A. Simon. 1959. Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing*. UNESCO, Paris, 256–264.
- Newell, A., H. A. Simon. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* **19**(3) 113–126.
- Nickerson, J. A., B. S. Silverman, T. R. Zenger. 2007. The ‘problem’ of creating and capturing value. *Strategic Organization* **5**(3) 211–225.
- Nickerson, J. A., T. R. Zenger. 2004. A knowledge-based theory of the firm: The problem-solving perspective. *Organization Science* **15**(6) 617–632.

- Porter, M. E. 1980. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. Free Press, New York.
- Porter, M. E. 1996. What is strategy? *Harvard Business Review* **74**(6) 61–78.
- Posen, H. E., D. A. Levinthal. 2012. Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science* **58**(3) 587–601.
- Prahalad, C. K., R. A. Bettis. 1986. The dominant logic: A new linkage between diversity and performance. *Strategic Management Journal* **7**(6) 485–501.
- Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science* **46** 824–844.
- Rivkin, J. W., N. Siggelkow. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* **49**(3) 290–311.
- Rivkin, J. W., N. Siggelkow. 2007. Patterned interactions in complex systems: Implications for exploration. *Management Science* **53**(7) 1068–1085.
- Sharpe, W. F. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* **19**(3) 425–442.
- Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**(1) 99–118.
- Simon, H. A. 1956. Rational choice and the structure of the environment. *Psychological Review* **63**(2) 129–138.
- Simon, H. A. 1997. *An Empirically Based Microeconomics*. Cambridge University Press, Cambridge, UK.
- Sommer, S. C., C. H. Loch. 2004. Selectionism and learning in projects with complexity and unforeseeable uncertainty. *Management Science* **50**(10) 1334–1347.
- Thagard, P. 2005. *Mind: Introduction to Cognitive Science*. 2nd ed. MIT Press, Cambridge, MA.
- Tripsas, M., G. Gavetti. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal* **21**(10-11) 1147–1161.
- Walsh, J. P. 1995. Managerial and organizational cognition: Notes from a trip down memory lane. *Organization Science* **6**(3) 280–321.
- Whitehead, A. N. 1911. *An Introduction to Mathematics*. H. Holt and Co., New York.

A Appendix: Formal description of the search process

A product \mathbf{p} is described by N binary policy choices ($\mathbf{p} \in \{0, 1\}^N$). Each of the M performance dimensions of product \mathbf{p} is defined as an independent NK landscape (Levinthal 1997); that is,

$$f_i(\mathbf{p}) = \frac{1}{N} \sum_{j=1}^N c_i^j(\mathbf{p}_j; K \text{ other elements of } \mathbf{p}) \text{ for } i \text{ in } \{1, \dots, M\}.$$

Here $f_i(\cdot)$ is the value of performance dimension i and $c_i^j(\cdot)$, referred to as a ‘contribution function,’ determines the contribution of policy j (when interacted with K other policies) toward the value of performance dimension i . Because K controls how many policies interact to determine the value of each contribution, we say that K controls the landscape’s *complexity*. The higher is K , the more rugged the landscape (i.e., the greater the number of local peaks).

The profits of the product \mathbf{p} are defined as a weighted average of its performance dimensions; thus,

$$\pi(\mathbf{p}) = \sum_{i=1}^M w^i \times f_i(\mathbf{p}).$$

Here w , $0 < w \leq 1$, parameterizes the relative importance of the performance dimensions.¹²

The manager’s mental representation is an incomplete version of the previous equation; that is, her representation accounts for only M' out of the M performance dimensions of the product:

$$\pi'(\mathbf{p}, \mathbf{d}) = \sum_{\{i|\mathbf{d}_i=1\}} w^i \times f_i(\mathbf{p}).$$

Here \mathbf{d} is a binary vector of size M , with 1s marking the M' performance dimensions considered by the manager (and with 0s elsewhere).¹³

¹²Although the relevance of dimensions is ordered according to index i , this ordering is not known to the actors and so does not reduce the search task’s difficulty.

¹³In other models (not reported here but available from the authors upon request), we considered mental representations that—besides not accounting for all product dimensions—use inaccurate weights. Adding this extra source of inaccuracy to the mental representation does not qualitatively change the reported results.

A.1 Manager’s search heuristic

We are interested in determining the relative advantages of different search strategies that managers can use. In particular, we focus on identifying the optimal balance between searching for the right set of dimensions (\mathbf{d}) and searching for the right set of policy choices (\mathbf{p}).

We assume that the search for new products occurs off-line (Gavetti and Levinthal 2000). In other words, because trying new products is expensive and time consuming, managers must first simply make their best attempt at designing one; only then can a real product be manufactured and its real performance (i.e., $\pi(\mathbf{p})$) be known. Thus a manager’s only resource when designing a new product is an (imperfect) understanding of the actual performance associated with any given design—namely, the mental representation $\pi'(\mathbf{p}, \mathbf{d})$. Hence the manager faces a twofold problem: (i) finding the right representation (i.e., which M' elements in \mathbf{d} should be set to 1); and (ii) choosing the business policies that define the product (i.e., choosing the N values in the binary vector \mathbf{p}).

The manager’s search process is modeled as follows. At each time step, the manager can either make a change to the policy choices (\mathbf{p}) or to the considered dimensions (\mathbf{d}). In either case, the manager explores a neighboring position. If exploring policy choices, the manager will ‘flip’ a single bit (i.e., from 0 to 1 or vice versa) of \mathbf{p} ; if exploring dimensions, the manager will flip *on* one bit of \mathbf{d} and flip *off* some other bit of \mathbf{d} while keeping the number of considered dimensions constant at M' . If the perceived fitness of the position being explored is greater than the perceived fitness of the current position, then the manager ‘moves’ to this new position. So at each time step, the manager compares the current position to a neighbor in either the space of policies or the space of performance dimensions. As presented by Rivkin and Siggelkow (2003), the search algorithm used in both cases is described by a search radius and number of alternatives considered equal to 1.

More precisely, at each time step the manager changes one policy with probability e_p or flips a dimension with probability $e_d = 1 - e_p$. Thus e_p represents how much of the search effort is devoted to searching in policy space. We model search effort as a limited resource, so more search among policies leads to less search among performance dimensions. If $e_p = 1$ then the manager searches only among policies, and if $e_p = 0$ (i.e., if $e_d = 1$) then the manager searches only among performance dimensions. Intermediate levels of e_p interpolate between these two extremes, so that, for example, if $e_p = 0.5$ then the search is equally balanced between the two search types. The

Parameter	Description
Task environment parameters	
N	Number of policy choices
M	Number of performance dimensions
K	Technological complexity (i.e., ruggedness of the landscapes representing the M performance dimensions)
w	Relative relevance of the performance dimensions
Manager parameters	
M'	Cognitive limit (i.e., dimensionality of the mental representation)
e_p	Search effort in policy space
Choices of the manager	
\mathbf{p}	Vector of policy choices (size N , binary)
\mathbf{d}	Vector of performance dimensions to consider (size M , binary)
Product performance (real and perceived)	
$\pi(\mathbf{p})$	Profits of the product described by policy choices \mathbf{p}
$\pi'(\mathbf{p}, \mathbf{d})$	Mental representation (i.e., perceived profits of product described by policy choices \mathbf{p} when the dimensions in \mathbf{d} that are set to 1 are considered)

Table A.1: Summary of the notation.

search process continues for T periods or time steps, where we interpret T as the extent of the search effort.

In Table A.1 we summarize the notation used to describe the model.