Strategic Management Journal
Strat. Mgmt. J., **26**: 691–712 (2005)
Published online 7 June 2005 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/smj.475

STRATEGY MAKING IN NOVEL AND COMPLEX WORLDS: THE POWER OF ANALOGY

GIOVANNI GAVETTI,[1] DANIEL A. LEVINTHAL[2]* and JAN W. RIVKIN[1]
[1] *Harvard Business School, Boston, Massachusetts, U.S.A.*
[2] *The Wharton School, Philadelphia, Pennsylvania, U.S.A.*

We examine how firms discover effective competitive positions in worlds that are both novel and complex. In such settings, neither rational deduction nor local search is likely to lead a firm to a successful array of choices. Analogical reasoning, however, may be helpful, allowing managers to transfer useful wisdom from similar settings they have experienced in the past. From a long list of observable industry characteristics, analogizing managers choose a subset they believe distinguishes similar industries from different ones. Faced with a novel industry, they seek a familiar industry which matches the novel one along that subset of characteristics. They transfer from the matching industry high-level policies that guide search in the novel industry.

We embody this conceptualization of analogy in an agent-based simulation model. The model allows us to examine the impact of managerial and structural characteristics on the effectiveness of analogical reasoning. With respect to managerial characteristics, we find, not surprisingly, that analogical reasoning is especially powerful when managers pay attention to characteristics that truly distinguish similar industries from different ones. More surprisingly, we find that the marginal returns to depth of experience diminish rapidly while greater breadth of experience steadily improves performance. Both depth and breadth of experience are useful only when one accurately understands what distinguishes similar industries from different ones. We also discover that following an analogy in too orthodox a manner—strictly constraining search efforts to what the analogy suggests—can be dysfunctional. With regard to structural characteristics, we find that a well-informed analogy is particularly powerful when interactions among decisions cross policy boundaries so that the underlying decision problem is not easily decomposed. Overall, the results shed light on a form of managerial reasoning that we believe is prevalent among practicing strategists yet is largely absent from scholarly analysis of strategy. Copyright © 2005 John Wiley & Sons, Ltd.

INTRODUCTION

Strategy making is most critical in times of change and in unfamiliar environments. It is in such a context that the strategy makers of a firm must identify a viable new strategic position or face the potential demise of their enterprise. Yet how are managers to make intelligent choices in these novel contexts? Popular accounts suggest that managers in such settings fall back on past learning and their experience in a variety of business settings. Managers with such experience may be capable of seeing through surface features of the current context to deeper truths that underlie it. This paper aims to shed light on the experiential wisdom that enables strategy makers to cope with novel environments. We argue that the basis of this sort of experiential wisdom lies in analogical reasoning. For past

Keywords: analogy; cognition; complex systems; strategic decision making; modularity
* Correspondence to: Daniel A. Levinthal, The Wharton School, University of Pennsylvania, 2028 Steinberg-Dietrich Hall, Philadelphia, PA 19104, U.S.A.
E-mail: levinthal@wharton.upenn.edu

Copyright © 2005 John Wiley & Sons, Ltd.
Received 25 March 2003
Final revision received 24 December 2004

experience to be of value in a world of novelty, actors must be able to generalize from prior settings to the current environment. This process of mapping from a source context of prior experience to the current, 'target' context is precisely what constitutes analogical reasoning (Gick and Holyoak, 1980). To shed light on strategy making in the face of novelty, we tackle two related tasks. First, we describe the anatomy of analogical reasoning, discussing in detail the components that make such logic work in the context of business strategy. Second, we build a simulation model of firms that use analogies to seek good strategic positions in unfamiliar settings. An analysis of the model allows us to identify situations in which analogical reasoning works well or poorly.

Our focus on analogical reasoning contrasts with the two dominant perspectives on strategic choice in prior literature. In one corner of the scholarly ring, the positioning school within strategic management has emphasized the importance of deductive reasoning and rational choice in the strategy-making process. The imagery is that of a general and staff in a situation room, coolly surveying topographical maps of the business landscape (Ghemawat, 1999), picking out a peak, and directing the troops to take the summit. This image is hard to sustain, however, in the world that positioning scholars have increasingly painted of firms' strategies embodied in intricate arrays of interlocking choices (Porter, 1996; Ghemawat, 1999; Rivkin, 2000; Siggelkow, 2002). With this recent emphasis on nuanced interactions among strategic choices, the positioning school has created a bit of a logical predicament for itself: it is precisely in complex worlds of highly interactive systems that deductive reasoning and rational choice seem least able to pinpoint effective positions. If the choices involved in a strategy are numerous and each affects the pay-offs associated with many others, the computational load created by a deductive process can quickly outstrip the bounded processing power (Simon, 1955) of any management team.[1] We explore below whether analogical reasoning can cope not only with novel settings, but also with novel settings that are complex.

In the opposite corner of the ring are those who argue that firms discover effective positions through local, boundedly rational search and luck. In any given industry, a broad set of firms begin their lives with a wide variety of strategic choices. Each firm's decisions and its routines are gradually perturbed in ways that enhance immediate performance, in a process that behavioral theorists of the firm (Cyert and March, 1963) and evolutionary scholars (Nelson and Winter, 1982) depict as largely automatic, experiential, and emergent (Mintzberg, 1978). A few fortunate firms happen upon highly effective sets of choices and survive an ensuing shakeout. The imagery here depicts blindfolded individuals spread out widely on a rugged landscape, each engaged in local hill climbing (Levinthal, 1997). The water level rises over time, and a few lucky survivors discover effective positions not by virtue of deductive power but simply because they started their search in a fortunate place and scrambled up the right inclines. This imagery doubtlessly paints an accurate picture of some successes. Pascale's (1984) well-known description of how Honda discovered its effective position in the U.S. motorcycle industry, for instance, is marked by happenstance and incremental response to unforeseen problems and opportunities.

We hesitate, however, to ascribe all effective positions to luck and local search. Human rationality may be limited, but intendedly rational action surely remains possible (March and Simon, 1958). By simplifying the problems they face, managers can bring problems within the bounds of their processing power and possibly come up with effective solutions (Simon, 1991). This is a central role of managerial cognition. Cognition can be especially powerful when coupled with local search (Gavetti and Levinthal, 2000); cognition can raise the odds that a firm begins its search near a high, promising mountain rather than in the valley beneath a mere foothill, and local search can complete the journey uphill. We see cognition, then, as a means to cope with complexity.

Novelty, however, remains a challenge. What cognitive processes are available to managers when they face unfamiliar problems? In novel situations, where deduction is likely difficult, past lessons from prior settings can be a tremendously powerful source of wisdom (Neustadt and May, 1986). One process for transferring such wisdom is analogical reasoning. As Thagard (1996: 80) points out, 'analogies can be computationally powerful in situations when conceptual and

[1] Indeed, Rivkin (2000) shows that strategic problems can become NP-complete—that is, intractable to algorithmic solution—as the degree of interaction among component choices passes a critical threshold.

Copyright © 2005 John Wiley & Sons, Ltd.

rule-based knowledge is not available.' Analogies allow actors to take the insight developed in one context and apply it to a new setting.

In this paper, we examine how managers can use analogy to tackle the twin challenges of complexity and novelty. The managers we consider discover new positions neither by reasoning from first principles of economics nor by undertaking unguided local search. Rather, when faced with a new and complex setting, managers identify the features of the setting that seem most pertinent, think back through their experiences in other settings with similar features, and recall the broad policies that worked well in those settings. These broad policies then form the starting point for a local search process. Analogies to other settings, drawn from direct or vicarious experience, guide the strategy-making process. The image of the strategist here is that of a wizened trail guide who might not know the details of the local landscape, but can recognize types of terrain well enough to give some general guidance. Analogical reasoning gives managerial cognition a significant hand in strategy making, and it emphasizes aspects of strategy making like pattern recognition, judgment, and even wisdom—aspects that, in our view, are prominent among practicing strategists but are understated in the academic literature on strategy.

To begin to address this gap, we develop an analytical structure with which we can explore both the power and limits of analogical reasoning in discovering strategic positions. We do so by first describing qualitatively what we view as the basic features of analogical reasoning in strategy making. We next formalize this argument in an agent-based simulation of firms that struggle to find superior competitive positions in unfamiliar industries. The simulation clarifies our depiction of analogical reasoning. Results of the simulation identify the managerial and structural characteristics that make analogical reasoning more or less powerful. Our analysis highlights the critical role played by a manager's mental representations, the conditions that make breadth and depth of managerial experience valuable, the dangers of holding too closely to the guidance offered by an analogy, and the power of analogy in the face of poorly decomposed decision problems.

In our concluding discussion, we suggest that the approach taken here helps to bridge two perspectives on strategy making: behavioral approaches that emphasize limits to cognition and the importance of search processes, and the positioning approach, with its emphasis on conscious, *a priori* strategic choice. Each approach encompasses important elements of reality. The strategy field, at an aggregate level, has recognized the importance of pluralism, but we often lack individual analytical frameworks that embrace and meld multiple perspectives. We suggest that our analytical structure has this property and therefore, beyond the value of the particular results we offer, we view the conceptual framework as a useful contribution that may facilitate a more integrated analysis of strategic decision making.

MAKING STRATEGY BY ANALOGY

Reasoning by analogy is a common form of logic among business strategists. Facing a novel opportunity or predicament, strategists think back to some similar situation they have faced or heard about, and they apply the lessons from that previous experience. Analogies—to the past, to other firms or industries, and to other competitive settings like sports or war—come up frequently in strategy discussions. Popular management writers hail pattern recognition and associated analogical thinking as the best way for managers to cope with rapid change (e.g., Slywotzky and Morrison, 1999). The case method, perhaps the most popular form of management education, is designed in part to give students a rich base from which to draw analogies. Through the case method, students 'are led to active consideration of a tremendous number of diverse and related real situations, which it would take them at least a lifetime of experience to encounter, and they are thus given a basis for comparison and analysis when they enter upon their careers of business action' (Gragg, 1940).

In this section, we lay out the key elements of analogical reasoning. A few examples illustrate how analogies are commonly used and give us concrete instances in which we ground the following conceptual discussion:

- The supermarket has served repeatedly as the basis for analogical reasoning. Charlie Merrill fashioned Merrill Lynch's distinctive approach to retail brokerage after the practices he witnessed as a long-time manager in the supermarket industry. 'Although I am supposed to be an

Copyright © 2005 John Wiley & Sons, Ltd.

investment banker,' he confessed, 'I think I am really and truly a grocery man at heart' (Perkins, 1999). Likewise, when Charles Lazarus founded the toy superstore Toys 'R' Us in the 1950s, he relied explicitly on an analogy to supermarket retailing. Indeed, he called his retail outlet the Baby Furniture and Toy Supermarket until the signage at a new site demanded a shorter name (Hast, 1992). And when Thomas Stemberg, a former supermarket executive, established the office supply superstore Staples, he posed the initial strategic insight as an analogical question: 'Could we create a Toys 'R' Us for office supplies?' (Stemberg, 1996). The basic supermarket formula—exhaustive selection, low prices and margins, and high volume—has been applied in a wide range of retail categories.

- Starting in the 1970s, Circuit City thrived by selling consumer electronics in large superstores. A wide selection of products, professional sales help, and a policy of not haggling with customers distinguished the company's stores. In 1993, Circuit City surprised investors by announcing that it would open CarMax, a chain of used car outlets. The company argued that the used car industry of the 1990s bore a close resemblance to the electronics retailing environment of the 1970s. The industry was dominated, for instance, by small Mom & Pop dealers with questionable reputations. The company hoped that its success formula from electronics retailing would work well in an apparently analogous setting.
- In 1997, the Board of the Internet portal Lycos decided to grow rapidly and become a full-fledged new-media company, in part by means of acquisition. After making its first acquisition—of homepage builder Tripod—Lycos's managers faced a decision that they later identified as the pivotal choice in the company's history: should Lycos maintain Tripod as a separate operation with its own brand name or integrate Tripod more fully into Lycos-branded operations? The management team argued the issue for weeks, and ultimately an analogy proved decisive in the debate. Traditional media companies, the managers observed, tended to maintain multiple divisions with separate brands but to coordinate back-office operations. Since Lycos wanted to become a large media company, the managers reasoned, the company should do likewise. Accordingly, the company kept alive the brands of Tripod and subsequent acquisitions, but integrated the back offices (Gavetti and Rivkin, 2004).
- Many factors contributed to Enron's startling failure, but headlong diversification based on loose analogies played an important role. After apparent success in trading natural gas and electric power, Enron executives moved rapidly to create markets for other goods ranging from coal, steel, and pulp and paper to weather derivatives and broadband telecom capacity. Analogical reasoning seemed to drive the expansion. Executives looked for markets with certain characteristics: fragmented demand, rapid change due to deregulation or technological progress, complex and capital intensive distribution systems, lengthy sales cycles, opaque pricing, and mismatches between long-term supply contracts and short-term fluctuations in customer demand (Salter, Levesque, and Ciampa, 2002). On the broadband opportunity, for instance, Enron Chairman Kenneth Lay said: '[Broadband]'s going to start off as a very inefficient market. It's going to settle down to a business model that looks very much like our business model on [gas and electricity] wholesale, which obviously has been very profitable with rapid growth' (*Gas Daily*, 2000). However, the analogical reasoning failed to appreciate important, deeper differences between the market for natural gas and the market for bandwidth. The broadband market was based on unproven technology and was dominated by telecom companies that largely resented Enron's encroachment. The underlying good, bandwidth, did not lend itself to the kinds of standard contracts that made efficient trading possible in gas and electricity. Perhaps worst, in broadband trading Enron had to deliver capacity the 'last mile' to a customer's site, an expensive chllenge that gas wholesalers did not face. Even according to Enron's misleadingly rosy financial reports, the broadband venture was disastrous (Salter *et al*., 2002).

In each of these cases, a management team borrowed a broad set of related choices from one industry and applied the system to a new industry that it believed to be similar on some crucial dimensions. In some instances (e.g., Merrill Lynch, Lycos), the team was motivated by a problem that required a solution. In other cases (e.g., Enron), the analogy took the form of a solution seeking a problem. In

Copyright © 2005 John Wiley & Sons, Ltd.

the following discussion, we focus on the problem-requiring-a-solution type of analogy, though we suggest that the solution-seeking-a-problem form of analogical reasoning is important as well, particularly as a source on entrepreneurial activity.

Well beyond the context of business, the use of analogical reasoning has been of long-standing interest among cognitive scientists (Gick and Holyoak, 1980; Holyoak and Thagard, 1995; Thagard, 1996). When reasoning by analogy, an individual starts with a situation to be handled—the target problem. The actor develops a mental representation of the target problem, a lower-dimensional sketch that, in the actor's view, captures the salient characteristics of the situation. She then uses some computational procedure to scour other settings with which she is familiar, due to either direct or vicarious experience, and identifies a setting that displays similar salient characteristics. This setting serves as a source of a candidate solution. The individual then transfers the candidate solution and applies it to the target problem. Clearly the power of analogy depends on the validity of the similarity mapping between source and target contexts as well as the quality of the solution suggested by the source context.

The danger of 'superficial mappings' between the source and the target is a widely studied phenomenon in cognitive psychology. Experimental work and field evidence strongly suggest that poor analogies are typically based on representations that capture only superficial features of the problems (Holyoak and Thagard, 1995). Indeed, research in this field shows that superficial similarity can readily induce experimental subjects, even well-educated individuals, to adopt poor analogies.[2] We suspect the same is true of practicing managers.

In the context of business strategy, the observable characteristics of an industry may constitute the dimensions of a representation. Three features of any industry, for instance, are the size of economies of scale, the size of customer switching costs, and the heterogeneity of customer tastes. Suppose a manager in a novel setting opts to represent her target problem along these dimensions. On the basis of an initial assessment of the target, the manager judges that the target industry is characterized by modest economies of scale, large switching costs, and diverse customer tastes. The manager then engages in a simple computational procedure: where has she seen modest economies of scale, large switching costs, and diverse tastes before, in other industries? The manager reviews her experience and realizes that, in a specific industry that was similar along these three dimensions, a particular niche provider of high-end, premium products was highly successful. She then transfers this solution to the target industry, adopting a small-scale manufacturing policy, a cream-skimming pricing policy, a targeted sales policy, and so forth. If the dimensions she chose to focus on are the ones that best summarize the true drivers of performance, the firm improves its odds of success.

The difficulty that faces the analogizing manager is that there are innumerable dimensions along which one can form a representation and some dimensions may be misleading. Suppose, for instance, that a different manager facing the same target industry ignores economies of scale, switching costs, and customer heterogeneity. Instead he pays attention only to the prevalence of Internet technology in the industry. He notes that the target industry relies heavily on the Internet for sales, marketing, and distribution.

[2] In one study (Gilovich, 1981), for example, students of international conflict at Stanford were told of a hypothetical foreign policy crisis: a small democratic nation was being threatened by an aggressive, totalitarian neighbor. Each student was asked to play the role of a State Department official and to recommend a course of action. The descriptions of the situation were manipulated slightly. Some of the students heard versions with cues that were intended to make them think of the crises that preceded World War II. The President at the time, they were told for example, was 'from New York, like Franklin Roosevelt,' refugees were fleeing in boxcars, and the briefing was held in 'Winston Churchill Hall.' Other students heard versions that might have reminded them of Vietnam. The President was 'from Texas, the same state as Lyndon Johnson,' refugees were escaping in small boats, and the briefing took place in 'Dean Rusk Hall.' Clearly, there is little reason that the home state of the President, the vehicles used by refugees, or a briefing room name should influence one's recommendation. Yet the surface features caused the two groups to reach very different conclusions. Students in the first group were significantly more likely to apply the lessons of World War II—that aggression must be met with force—than were students in the second group, which veered toward a hands-off policy inspired by Vietnam.

Not only were the subjects of the experiment lured by superficial likeness, but—perhaps more disturbing—they were *not even aware* that they had been lured. After making a recommendation, each student was asked, 'How similar is the situation to World War II?' and 'How similar is the situation to Vietnam?' The two groups gave identical answers. Small, superficial, irrelevant features led well-educated students to draw their analogies from different sources. This caused them to recommend very different candidate solutions. Yet the students seemed unaware that they were drawing from different sources.

Copyright © 2005 John Wiley & Sons, Ltd.

In his experience with Internet-based industries, successful companies in such settings spend aggressively in order to get big fast. He deploys this candidate solution in the target industry. If the prevalence of Internet technology does not truly shape the target landscape, he may find that his mass-market product is far less appealing to diverse customers than the offerings of niche competitors and that his large-scale manufacturing operations do not lower his costs. By focusing on an irrelevant dimension and ignoring three other, more pertinent characteristics, he has been led to a poor analogy.

Analogical reasoning and strategic positioning

The target problem is the business situation the strategist hopes to resolve. In our particular setting, the challenge of interest is how to position a firm in an industry that is novel—novel either for the managers of the firm itself (e.g., the used car business for Circuit City's managers) or for all managers (e.g., the Internet portal industry). To position itself, the company must make a vast array of detailed choices about how to develop, design, produce, sell, deliver, and service products (Porter, 1985). These choices incur costs and generate buyer value, and therefore they shape the economic success of the firm.

We find it useful to conceive of the target problem in terms of a high-dimensional performance landscape (Kauffman, 1995; Levinthal, 1997; McKelvey, 1999, and references therein). Each detailed choice constitutes a horizontal dimension on this landscape, and the vertical dimension records the economic success associated with each combination of choices, thereby creating a surface. The task of management is to discover a combination of choices that, together, produce high performance; in graphical terms, the challenge is to find a high peak on the performance landscape. Because the decisions may interact with one another (that is, the choice made on each may alter the marginal costs and benefits associated with the others), the performance landscape may contain numerous local peaks. A local peak is a configuration of choices from which one cannot improve performance by altering any single choice, even though simultaneous change in many choices may reposition a firm to higher ground.

In many business settings, detailed choices cluster together to form policy domains. Choices about automation, lot sizes, work flow, factory scale, and so forth may aggregate into a manufacturing policy, for example, with combinations of these detailed choices taking on labels that summarize the policies (e.g., mass production, job shop operations, cell manufacturing, or batch production). Obviously, the choice of a particular approach for higher-order policies typically has an influence on detailed choices. Three aspects of these policy domains shape the modeling choices we make below. First, the interactions within the decisions that make up a policy domain may be more intense than the interactions across domains. This creates a type of interdependence that Simon (1962) has labeled near-decomposability. In our subsequent analysis, we consider the effect of decomposability on the efficacy of analogy.

Second, the effect of one policy domain on the efficacy of another may depend not on the detailed choices made in the first domain, but on the net effect of those choices. Take, for instance, the example of Circuit City's original success in consumer electronics. Circuit City's success in marketing and selling consumer electronics clearly had an effect on its policy for procuring electronics. When the marketing and sales policy produced strong demand among consumers in the 1970s, the marginal value of large-scale procurement increased dramatically. The effect of marketing and sales on procurement depended not on detailed marketing and sales choices such as pricing, sales tactics, and sales compensation, but only on the total demand that those choices generated in aggregate.

Third, the high-level policies adopted by a company are often more visible to outsiders and perhaps more memorable to insiders than the detailed decisions that underpin them. In Circuit City's case, the consumer-friendly marketing and sales policy that it created in the electronics retailing market was highly visible, but the compensation schemes and information technology that enabled that policy remained more obscure. We assume below that analogies provide guidance about high-level policies, not detailed decisions per se.

Quality of analogically based solutions

On the basis of representational similarity, the analogizing manager chooses a source industry from her library of experience. Beyond the choice of representation, we see three factors that

Copyright © 2005 John Wiley & Sons, Ltd.

influence the quality of the guidance provided by an analogy. First, there is *the breadth of the manager's experience*. The larger the number of other industries the manager has experienced, personally or vicariously, the greater is the chance that the manager will be familiar with an industry that matches the target well along the representational dimensions. It is this kind of breadth that strategy consultants often purport to bring to a client.

A second factor that affects the quality of an analogy's guidance is *the depth of the manager's experience* in the source industry. If the manager has spent a great deal of time in the source industry and understands deeply what distinguishes a good position from a poor one in that setting, it is more likely that the analogy will accurately guide the manager's firm to a favorable set of policies. On the other hand, deep experience in a source setting might be a double-edged sword (Levinthal and March, 1993). If a manager couples deep experience with a poorly chosen source, the result can be disastrous: profoundly convinced that she knows exactly what to do, the manager may persist in applying the policies that succeeded in a prior, but very different, setting and may ignore negative feedback from the environment. In this sense, Enron's extensive experience with gas and electricity trading may have impeded its ability to adapt in its broadband venture. More generally, it is useful to consider how seriously managers take their analogies and how closely they abide by the candidate solutions. In the formal model we develop below, we consider both orthodox analogizers, who consistently hold to the set of policies recommended by their analogy, and heterodox analogizers, who use the recommended set only as a starting point for further exploration.

Analogies typically give high-level guidance whose details must be worked out at the ground level of the target industry. This points to a third factor that influences the quality of the guidance provided by an analogy: the degree to which a high-level policy choice constrains the detailed choices that comprise it. In some settings, a policy choice may tightly constrain detailed choices, while in others a policy choice may leave wide latitude for what goes on below. Since candidate solutions are transferred at the level of policies, not detailed choices, an analogy provides sharper guidance in the former setting than in the latter. When an analogy is chosen well, we expect clearer guidance to produce higher performance.

The subtlety of analogizing and its implications for the establishment of effective competitive positions suggest that a formal model of the process may generate valuable insights. In the following section, we describe such a model, built as an agent-based simulation. The model incorporates the features we have described here: representations of variable quality; the possibility of analogies based on superficial similarity; target and source industries of variable complexity; a distinction between broad policies and detailed choices; interactions across choices within policies and interactions across policies; and managers who differ in the breadth and depth of their experience and in the orthodoxy with which they follow their candidate solutions. The model allows us to confirm some of the relatively straightforward arguments we have made here (e.g., that better representations make analogical reasoning more effective). It also reveals some less obvious factors, such as the role of the underlying interaction structure among policy choices, that make analogical reasoning a more or less powerful tool for discovering effective competitive positions in novel and complex settings.

A MODEL OF SEARCH IN NOVEL AND COMPLEX WORLDS

The simulation model undertakes two basic sets of operations. First, it generates a family of performance landscapes—a target landscape and a set of potential sources—in which the modeler can tune the relationship between the sources and the target. The performance landscape itself can be tuned with respect to the degree to which firm choices on these landscapes are decomposable. Second, the model permits firms to search for effective positions on the target landscape in a variety of ways. Firms may differ, for example, in terms of the mapping between targets and sources that each firm uses, the way in which attractive candidate solutions are identified on each source landscape, and the manner in which each firm couples analogical reasoning with incremental search. We discuss the generation of families of landscapes and the search for effective positions in turn. (Table 1 summarizes the model's parameters and symbols.)

Copyright © 2005 John Wiley & Sons, Ltd.

Table 1. Parameters and symbols

Parameters related to generation of families of landscapes	P	The number of high-level policy decisions that each firm faces on a target or source landscape.
	D	The number of detailed decisions each firm must make within each high-level policy. $P \times D$ is the total number of decisions each firm makes. Each decision is binary so there are $2^{P \times D}$ possible configurations of detailed decisions.
	X	The total number of observable industry characteristics. Each characteristic is binary, so there are 2^X industries or landscapes in a family: one target and $2^X - 1$ potential sources.
	K_w	The probability that the performance contribution of a focal detailed decision is affected by the resolution of each of the other $D - 1$ detailed decisions within the focal decision's policy. The 'w' in K_w signifies 'within'.
	K_b	The probability that the performance contribution of a focal detailed decision is affected by the resolution of each of the other $P - 1$ high-level policies. 'b' signifies 'between'.
	$\boldsymbol{X}_{\mathbf{PROB}}$	An X-digit vector of probabilities. The first element is the probability that the first observable characteristic affects the performance contribution of each detailed decision. Likewise for elements 2, 3, ..., X.
Parameters related to firm search	DEPTH	The portion of each source landscape that is evaluated in order to generate the library of most promising policy configurations. The higher this value is, the more likely it is that the library truly identifies the best policies for each source.
	BREADTH	The portion of the library of most promising policy configurations that is available to a particular analogizing firm.
	⟨# # #⟩	Numbers within angle brackets refer to the representation of each analogizing firm. The first number is the observable characteristic the firm's managers deem the most important; the second number is the second most important; and so forth.
	Orthodoxy	Each analogizing firm is set to be orthodox or heterodox. During incremental search, orthodox firms never violate the policy guidance provided by the analogical source, while heterodox firms may ignore that guidance.
Symbols used to describe the model	**d**	A particular configuration of detailed decisions; a $P \times D$ string of zeros and ones.
	d_i	The resolution, zero or one, of a specific detailed decision. i is between 1 and $P \times D$.
	{# # #}	Numbers within braces consistently refer to lists of $P \times D$ detailed decisions. Within a list, vertical lines separate the P policy domains.
	[# # #]	Numbers within square brackets refer to lists of P high-level policy choices.
	(# # #)	Numbers within parentheses refer to lists of X observable industry characteristics.

Families of performance landscapes

We generate families of landscapes using an adaptation of Kauffman's (1993) NK model.[3] The NK structure highlights the interdependency among choices. Settings in which choices are more interdependent (a higher K value in the standard structure) result in more rugged, multi-peak performance surfaces because a change in one choice will have repercussions for many of the N choices.

[3] Developed in the biological sciences, Kauffman's model has now been used to explore a variety of issues related to organizational and technological search. See, for instance, Kauffman (1995), Westhoff, Yarbrough, and Yarbrough (1996), Levinthal (1997), McKelvey (1999), Rivkin (2000), Gavetti and Levinthal (2000), Rivkin (2001), and Rivkin and Siggelkow (2002, 2003). Fleming and Sorenson (2001, 2004) have found that subtle predictions of the NK model concerning patterns in technological search are borne out in patent data.

Copyright © 2005 John Wiley & Sons, Ltd.

In our elaboration of Kauffman's basic framework, we introduce a hierarchy of choices. Each firm is assumed to face P high-level policy decisions, each of which includes D detailed choices. A competitive position, then, is defined by $P \times D$ detailed choices. For simplicity, we assume that each choice offers two options. A strategic position can then be represented as a $P \times D$-digit string of zeros and ones: $\mathbf{d} = \{d_1\ d_2 \ldots d_{P\times D}\}$ with $d_i = 0$ or 1 for all i. Therefore, a firm has $2^{P\times D}$ positions available to it in the target industry. Likewise, in each source industry, there are $2^{P\times D}$ possible configurations of choices.

Policy choices are related to detailed choices by a simple majority rule. Suppose, for instance, that $D = 3$. We say that a particular policy choice is 1 if the detailed decisions within the policy are configured such that the majority of them are 1: {111}, {110}, {101}, or {011}. It is 0 if most of the detailed decisions are 0: {000}, {001}, {010}, or {100}. Each configuration of detailed choices gives a unique configuration of policies, but each configuration of policies is consistent with many configurations of detailed choices.[4]

In order to explore analogical reasoning and possible similarity between source and target contexts, we also need to characterize and be able to tune the degree to which one business context is like another. We postulate that associated with all industries, including the target industry and each potential source industry, is a set of observable characteristics. These are the possible dimensions along which representations can be drawn. In particular, we assume that there are X observable characteristics, each of which offers two options. As a result, there are 2^X combinations of observable characteristics. Given a particular specification of the X characteristics for the target industry, there are $2^X - 1$ possible combinations of characteristics for potential source industries.[5]

Thus each industry has an X-digit 'tag.' Recall, for instance, the hypothetical situation we described earlier in which industries had four observable characteristics: the size of economies of scale (large or small), the size of customer switching costs (large or small), the heterogeneity of customer tastes (high or low), and the prevalence of Internet technology (high or low). The particular industry our two managers faced, with small economies of scale, high switching costs, diverse customers, and prevalent Internet technology, might be represented as (0 1 1 1). In choosing an analogical source, each manager paid attention to only part of this vector; we return to this consideration below.

Our goal is to create families of landscapes in which (a) we control the pattern of interaction among detailed choices within and across policy domains, (b) pairs of landscapes with similar observable characteristics—similar tags—are more alike than are pairs with very different characteristics, and (c) some observable characteristics have more influence than others on landscape topography.

To meet this goal, we assume that the contribution to performance of each detailed decision on each landscape depends not only on the resolution of that decision (0 or 1), but also, possibly, on the resolutions of other decisions within the focal decision's policy, the resolutions of other policies, and the states of observable characteristics. The pattern of dependence is influenced by three parameters. K_w, a number between 0 and 1, is the probability that each detailed choice's contribution depends on the resolution of each other choice *within* the focal choice's policy domain. K_b, a number between 0 and 1, controls interdependence *between* policy domains. Specifically, it is the probability that the focal choice's contribution depends on the resolution of each other policy. $\boldsymbol{X}_{\mathbf{PROB}}$, a vector of X numbers each between 0 and 1, controls the probability that each observable characteristic influences the contribution of the focal choice. An observable characteristic with a large component in $\boldsymbol{X}_{\mathbf{PROB}}$ has a large impact on landscape topography.

For a given set of values for P, D, X, K_w, K_b, and $\boldsymbol{X}_{\mathbf{PROB}}$, the simulating computer generates an influence matrix—a matrix that records which choices, policies, and observable characteristics influence the contribution of each detailed choice. (See Figure 1 for an example.) This influence

[4] Suppose, for instance, that $P = 3$ and $D = 3$. Then the detailed configuration {011|111|001} implies the policy choice [1 1 0]. But [1 1 0] is also consistent with {101|011|000}, as well as many other configurations of detailed choices. For clarity, we place policy configurations in square brackets and detailed configurations in braces. Within detailed configurations, we will separate policy domains by vertical bars. We require D to be an odd number so that the majority rule provides a unique policy choice for each configuration of detailed decisions.

[5] There are 2^X possible combinations of characteristics; however, one of them signifies the target landscape itself.

Copyright © 2005 John Wiley & Sons, Ltd.

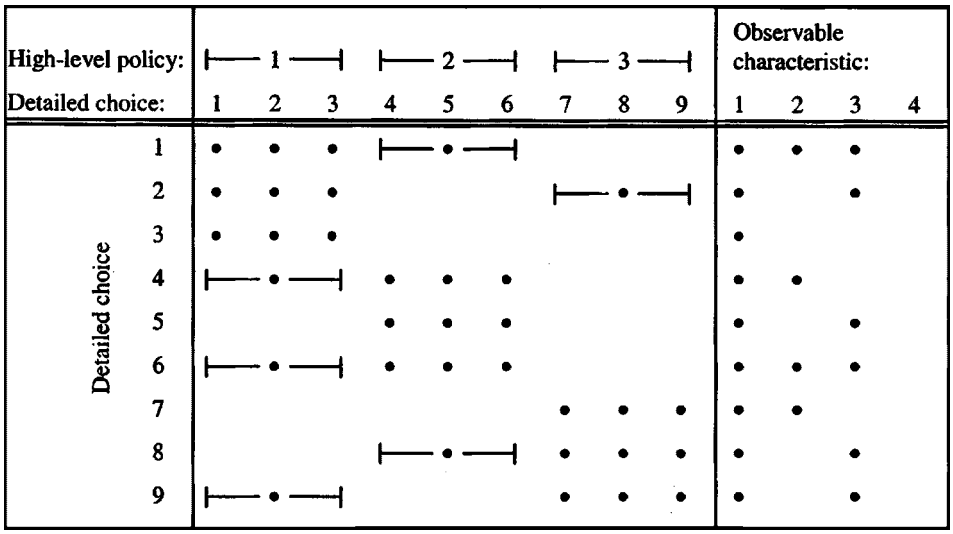


Figure 1. Example of an influence matrix

matrix is then used to generate target and source landscapes. We lay out the details of this procedure in an Appendix, and we describe the upshot of the procedure in intuitive terms in the following two paragraphs.

- First consider observable characteristics. Each of the 2^X combinations of characteristics defines a particular business context. In the example we used above, the target context was characterized by the vector (0 1 1 1). There are 15 possible source settings with profiles of characteristics such as (0 0 0 1), (0 0 1 0), (0 0 1 1), and so forth. Suppose the influence matrix looks like Figure 1. The fourth observable characteristic does not influence the contribution of any choice. Therefore, on source landscape (0 1 1 0), the contribution of every decision is the same as it is on the target landscape (0 1 1 1); the difference in the fourth characteristic alters nothing. Source landscape (0 1 1 0) is *identical* to the target landscape in every way and is therefore an excellent source for analogy. In contrast, the first observable characteristic influences the contribution of every choice. On source landscape (1 1 1 1), which differs from the target (0 1 1 1) only along the first characteristic, the contribution of every decision is different from what it is on the target. The difference in the first characteristic shifts each contribution altogether so that (1 1 1 1) bears no resemblance whatsoever to the target (0 1 1 1). More generally, the larger is an element of $\boldsymbol{X}_{\mathbf{PROB}}$, the more profoundly does a difference in that characteristic alter the landscape. In choosing a source for an analogy, it is crucial to pay attention to observable dimensions with high values of $\boldsymbol{X}_{\mathbf{PROB}}$. Our procedure for constructing families of landscapes ensures that landscapes with similar observable characteristics have similar shapes. Moreover, it allows us to control which characteristics strongly delineate groups of landscapes with similar topographies and which characteristics are weak, or even irrelevant, delineators.
- Next consider the effects of K_w and K_b. When $K_w = K_b = 0$, the contribution of each choice depends only on the resolution of that choice and the state of observable characteristics. On any given landscape, i.e., for any particular set of observable characteristics, the decision problem facing a firm is easy to address; a change in any choice alters the contribution of no other choice so a firm can find the optimal configuration of choices simply by adjusting each choice to the resolution, 0 or 1, that makes the greater contribution in isolation. The landscape is smooth, single peaked, and easy to scale. When K_w is high and K_b remains low, the system is nearly decomposable (Simon, 1962). A change in any choice alters the contributions of other choices within the same policy, but it is possible for the firm to tackle its decision

Copyright © 2005 John Wiley & Sons, Ltd.

problem policy by policy. As K_b rises, however, decomposition becomes more difficult because a change that alters a policy now has ramifications for choices in other domains. Graphically, we find it useful to think of K_w and K_b in terms of plateaus. When K_w and K_b are low, landscapes are smooth. As K_b rises, they develop distinct plateaus, with the edges of plateaus defined by changes in policies. As K_w rises, the surfaces of individual plateaus become internally rugged.

Having generated a family of landscapes—target and potential sources—whose relations to one another are well controlled, the computer pinpoints an attractive candidate solution on each of the $(2^X - 1)$ potential sources. Specifically, for a given source landscape, the computer considers each of the 2^P combinations of policies, surveys a fraction DEPTH of the detailed configurations consistent with each policy configuration, assesses the performance of each detailed configuration it surveys, and remembers the policy that produces the best results on average for the detailed configurations considered. The result is an exhaustive library of the most promising policy configurations on each and every potential source landscape. Table 2 shows what the library might look like for the example used above. The library includes the information, for instance, that the policy [1 0 0] produced the best results on the source landscape with observable characteristics (0 0 0 1). The promising policy configurations will later serve as candidate solutions. DEPTH parameterizes the depth of experience on which candidate solutions are based.

Table 2. Example of a library of source landscapes and promising policies

Source landscape	Most promising policy configuration	Average performance within most promising configuration
(0 0 0 0)	[1 0 0]	0.65
(0 0 0 1)	[1 0 0]	0.65
(0 0 1 0)	[0 1 1]	0.62
(0 0 1 1)	[0 1 1]	0.62
(0 1 0 0)	[0 0 0]	0.55
(0 1 0 1)	[0 0 0]	0.55
(0 1 1 0)	[0 1 0]	0.74
(1 0 0 0)	[1 0 0]	0.61
(1 0 0 1)	[1 0 0]	0.61
(1 0 1 0)	[1 1 0]	0.65
(1 0 1 1)	[1 1 0]	0.65
(1 1 0 0)	[0 1 0]	0.73
(1 1 0 1)	[0 1 0]	0.73
(1 1 1 0)	[1 1 0]	0.70
(1 1 1 1)	[1 1 0]	0.70

Note that in this example, pairs of source landscapes that differ only in the last characteristic, which has no influence on performance levels, always have the same most promising policy configuration and the same average performance within the most promising policy configuration (assuming DEPTH = 1).

Search for effective positions

The stage is now set for a discussion of how firms search the target landscape. We consider two basic classes of search strategies: local search and analogical reasoning. *Local searchers* rely on simple hill-climbing and make no use of analogy. Each is given an initial configuration of detailed choices and a corresponding set of policy choices by chance. That is, each is released at a random location on the target landscape. In each subsequent period, each local searcher considers the $P \times D$ alternatives to its current configuration that involve a change in a single detailed choice. If it spots opportunities for improvement, the local searcher pursues one of the opportunities. Otherwise, it has arrived atop a local peak and remains there for the duration of the simulation.

Analogizers, in contrast, base their initial configurations on the library that was generated earlier. Each analogizer has access to a fraction BREADTH of the full library of potential source landscapes. The particular source landscapes to which an analogizer has access are chosen at random. In the example above with 15 possible sources, an analogizer with BREADTH = 0.33 might be familiar with landscapes whose observable characteristics are (1 1 1 0), (1 1 0 1), (1 0 1 0), (0 1 0 0), and (1 1 0 0). A less experienced firm with BREADTH = 0.2 might know only the landscapes with the observable characteristics (0 1 1 0), (1 0 0 1), and (1 0 1 0).

Each analogizing firm also has an ordered list of the representational dimensions, or observable characteristics, that it considers to be important. We consider this list to be the firm's representation of its problem context. The representation ⟨3 2 1⟩, for instance, implies that the firm considers the third observable characteristic to be most important, the second characteristic to be second-most important, the first characteristic to be third-most important, and the fourth characteristic altogether irrelevant. The representation ⟨4⟩ reflects

Copyright © 2005 John Wiley & Sons, Ltd.

a belief that only the fourth characteristic matters. The first manager we described before—who considered scale economies, switching costs, and customer heterogeneity—might have representation ⟨3 2 1⟩ while the second manager—who paid attention only to the prevalence of Internet technology—might have representation ⟨4⟩.

The analogizer uses its representation to choose among source landscapes within its breadth of experience. It does so in a lexicographic manner. Take, for instance, a firm familiar with source landscapes (1 1 1 0), (1 1 0 1), (1 0 1 0), (0 1 0 0), and (1 1 0 0). Assume that the firm has the representation ⟨3 2 1⟩, and the target landscape is (0 1 1 1), as in our earlier example. The firm first looks for a source landscape that matches the target perfectly on characteristics 3, 2, and 1. It finds none. It then disregards the characteristic it considers least important, characteristic 1, and looks for a source landscape that matches the target on characteristics 3 and 2. It finds one such landscape, (1 1 1 0), and chooses to focus on that source.

Having chosen a source, the analogizer transfers the candidate solution—in this example and in line with Table 2, [1 1 0]—back to the target. Consistent with our discussion in the previous section, candidate solutions carry only high-level policy guidance. The firm chooses, at random, a set of detailed choices that abide by the policy guidance. In the example, the firm may choose {110|111|010}. This configuration serves as the firm's starting point for subsequent search. In following periods, the analogizing firm engages in incremental improvement just as a local searcher would. If we force the firm to abide by its analogy in an orthodox manner, it never considers incremental improvements that violate the initial policy guidance. From {110|111|010}, for instance, an orthodox analogizer would never consider {010|111|010} because to do so would switch the first policy from 1 to 0. On the other hand, if we allow the firm to be heterodox in its use of analogy, its incremental improvement efforts are not constrained in this manner.

Note the crucial role that the firm's representation plays in this process. If the representation matches $\boldsymbol{X}_{\mathbf{PROB}}$ closely, in the sense that observable characteristics with high influence are early in the firm's list, the firm is likely to focus on a source landscape that closely resembles the target. The candidate solution will then guide the firm to a promising starting point. On the other hand, if the representation causes the firm to pay attention to industry characteristics that are secondary or irrelevant, the source will typically bear little resemblance to the target, and the candidate solution will likely provide poor guidance.

RESULTS

Our simulation results identify factors that make analogical reasoning especially powerful as well as relationships among those factors. We start by looking at factors related to the analogizing management teams: the quality of their representations, the breadth and depth of their experience, and the orthodoxy with which they use the candidate solution. We then turn to structural characteristics of the industries they face, especially the decomposability of the problem space and the breadth of each policy domain. Throughout this section, we examine families of landscapes in which there are four observable characteristics ($X = 4$), three of which truly affect landscape structure and one of which is irrelevant (i.e., $\boldsymbol{X}_{\mathbf{PROB}} = (0.5, 0.5, 0.5, 0.0)$).

We compare the position-seeking success of five types of firms: a local searcher and four analogizers that differ in their representations. The first analogizer, with representation ⟨1 2 3⟩, correctly surmises that the first three observable characteristics affect the choice/payoff mapping. The second, with ⟨1⟩, heeds only one of the three relevant characteristics. The third, with ⟨4 3 2⟩, gives primacy to the irrelevant characteristic, but also heeds some relevant factors. The final analogizer, with ⟨4⟩, heeds only the characteristic that has no bearing on the landscapes. Among representations that include a subset of observable characteristics, these four cover the spectrum from the longest (⟨1 2 3⟩ and ⟨4 3 2⟩) to the shortest (⟨1⟩ and ⟨4⟩) as well as the range from the most accurate given length (⟨1 2 3⟩ and ⟨1⟩) to the least accurate (⟨4 3 2⟩ and ⟨4⟩).[6] Results for firms with other representations fall between results for these four firms in an intuitive way; the performance of a firm with representation ⟨1 2⟩, for instance, lies between that of the firm with ⟨1⟩ and that of

[6] Our premise is that boundedly rational managers cannot attend to all aspects of their environment. For this reason, we do not examine firms that consider all four observable characteristics.

Copyright © 2005 John Wiley & Sons, Ltd.

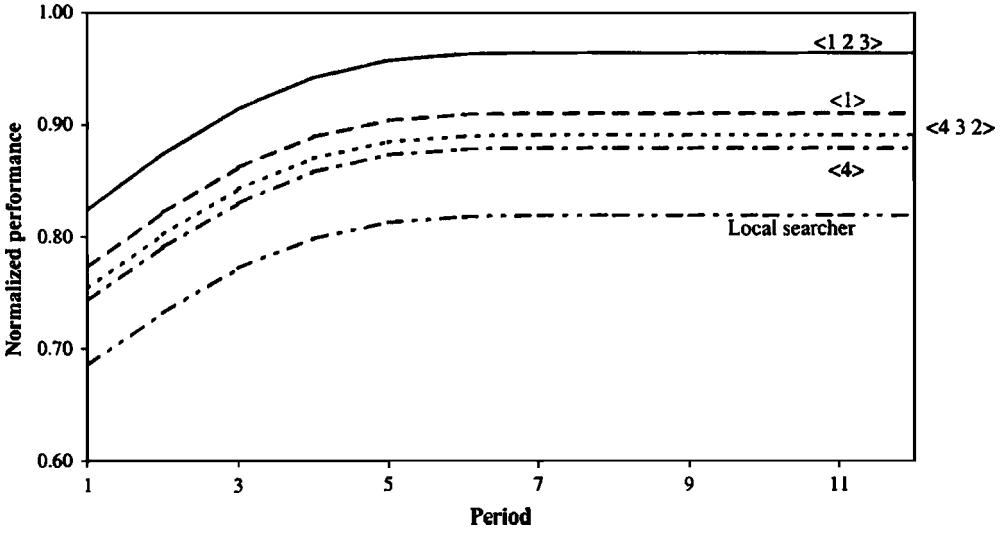


Figure 2. Firm performance over time

the firm with ⟨1 2 3⟩. Results for firms with other representations are available from the authors on request.

We report firm performance as a portion of the highest performance attainable on the target landscapes that were explored. By this measure, a type of firm that always attains the global maximum will achieve a performance of 1.00. Unless otherwise noted, each result is an average over 1,000 families of landscapes, each explored by 10 firms of each type. Each of the 1,000 families of landscapes shares the same structural parameters but constitutes an independent draw from the same underlying distribution.

Management characteristics

We first examine a situation in which $P = 3$ and $D = 3$. To highlight the effect of choosing the correct observable characteristics on which to base one's representation, we choose a set of baseline parameter settings in which cognition is most powerful. DEPTH = BREADTH = 1 for all firms, and all analogizing firms use candidate solutions in an orthodox manner. In addition, we set the probability of interaction across policy choices, K_b, equal to one and the probability of interaction among decisions within a policy, K_w, to zero.[7] For the sake of comparability, we model a local searcher that abides by its initial, randomly assigned policy choices in an orthodox manner.

Figure 2 shows the average performance of each type of firm over the course of the simulation. On average, firms with better representations achieve higher performance in the initial period than do firms with worse representations because their initial policy choices are guided by a source that more closely resembles the target. Within the constraints of these initial policy choices, the sets of detailed decision choices are randomly specified in the initial period. As a result, there is considerable room for each firm to adapt and improve its performance over subsequent time periods, even while it maintains its set of overarching policies. Higher-quality representations place firms in more opportune locales, as measured by their initial superior fitness as well as by their higher asymptotic value. While all firms search the landscape to identify more favorable sets of decisions, firms with better representations carry out this search in more favorable regions of the landscape.

An intriguing aspect of Figure 2 is that the firm with representation ⟨4⟩, the firm that heeds only an irrelevant dimension, fares substantially better than the local searcher. Even a 'false'

[7] The results we obtain from this first set of analyses are robust to different specifications for the interaction structure. We explore the impact of the interaction structure in the next set of analyses. We choose this particular combination of structural parameters because it yields the greatest power of cognition. Below we explore how the power of cognition depends on the structure of the decision problem.

Copyright © 2005 John Wiley & Sons, Ltd.

analogy has power. This result arises because, in choosing a source landscape on the basis of an irrelevant characteristic, the firm may, by chance, focus on a source landscape that happens to share *relevant* characteristics with the target as well. Suppose, for instance, that the target landscape has characteristics (0 1 1 1), as in the example in the previous section. The firm with representation ⟨4⟩ might choose as its source a landscape such as (1 1 1 1) or (0 0 1 1) or (0 1 0 1)—a landscape that *does* overlap with the target on some relevant dimensions. Based on this 'false' analogy, the firm identifies a set of choices coherent with this representation which, purely by chance, may prove useful in the target problem. In the discussion section below, we return to the virtues of false analogies.

We turn next to the effects of experience on the power of analogy. Table 3 shows the long-run differential between the performance of each analogizer and the performance of the local searcher for various levels of BREADTH, the portion of the source library to which the firm has access. Table 4 presents similar results for DEPTH, the portion of each policy domain that is explored to generate the library of candidate solutions. There are at least two striking features of these results. First, while greater BREADTH of experience steadily improves performance, the marginal returns to DEPTH of experience diminish rapidly; a management team with very limited experience on each source landscape fares nearly as well on the target as does a team with extensive experience.[8] This result suggests that it is far more important to identify an analogical source that shares the structure of the target problem than to pinpoint the absolute best solution within the source problem context. Second, greater breadth and depth of experience boost performance more when a firm has a high-quality representation. Put differently, broader and deeper experience does not improve a firm's lot if it draws from that experience on the basis of less relevant characteristics.

Tables 3 and 4 obscure a hazard of experience we discussed above: broad and especially deep experience may motivate a management team to abide by its candidate solution even though the subsequent local search process identifies superior

[8] In each row, the results for firms with DEPTH = 0.25, DEPTH = 0.50, and DEPTH = 1.00 are statistically indistinguishable.

Table 3. Effect of breadth of experience on analogizer's long-run performance advantage

Long-run performance advantage over a local searcher for an analogizer with representation ...	BREADTH =				
	0.10	0.25	0.50	0.75	1.00
⟨1 2 3⟩	0.069 (0.002)	0.086 (0.002)	0.111 (0.001)	0.132 (0.001)	0.146 (0.001)
⟨1⟩	0.065 (0.002)	0.074 (0.002)	0.081 (0.002)	0.086 (0.002)	0.090 (0.002)
⟨4 3 2⟩	0.055 (0.002)	0.055 (0.002)	0.066 (0.002)	0.069 (0.002)	0.073 (0.002)
⟨4⟩	0.054 (0.002)	0.053 (0.002)	0.056 (0.002)	0.057 (0.002)	0.059 (0.002)

$K_b = 1$, $K_w = 0$, DEPTH = 1, all firms are orthodox. Each figure is an average over 10,000 firms on 1,000 families of landscapes. Each number in parentheses is the standard deviation of the average figure to the left of it.

Table 4. Effect of depth of experience on analogizer's long-run performance advantage

Long-run performance advantage over a local searcher for an analogizer with representation ...	DEPTH =				
	0.01	0.10	0.25	0.50	1.00
⟨1 2 3⟩	0.119 (0.001)	0.141 (0.001)	0.146 (0.001)	0.146 (0.001)	0.146 (0.001)
⟨1⟩	0.083 (0.001)	0.090 (0.001)	0.093 (0.001)	0.093 (0.001)	0.092 (0.001)
⟨4 3 2⟩	0.060 (0.001)	0.072 (0.001)	0.072 (0.001)	0.071 (0.001)	0.072 (0.001)
⟨4⟩	0.054 (0.001)	0.059 (0.001)	0.062 (0.001)	0.061 (0.001)	0.060 (0.001)

$K_b = 1$, $K_w = 0$, BREADTH = 1, all firms are orthodox. Each figure is an average over 20,000 firms on 2,000 families of landscapes. Each number in parentheses is the standard deviation of the average figure to the left of it.

Copyright © 2005 John Wiley & Sons, Ltd.

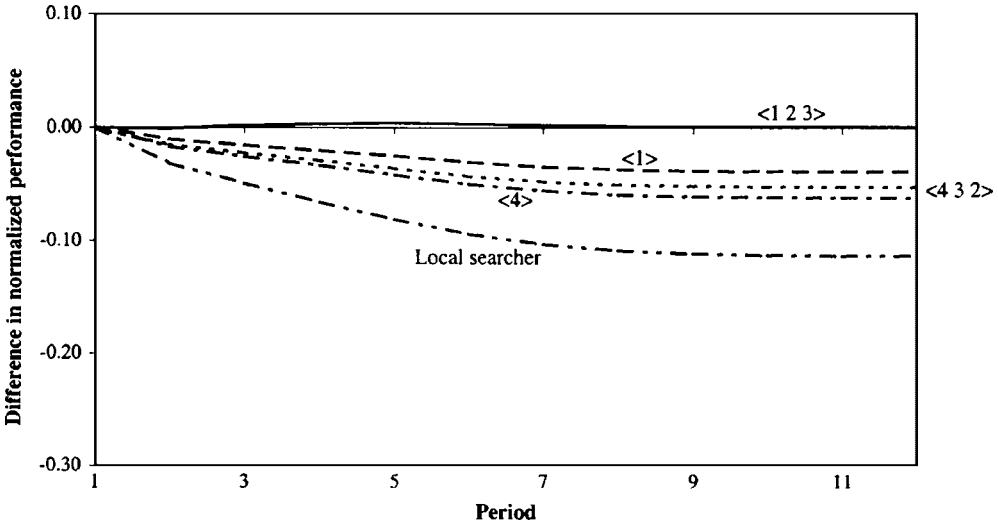


Figure 3. Performance differential due to orthodoxy

solutions. To size up this danger, we repeat the experiment shown in Figure 2, but now allow firms to be heterodox; that is, when improving incrementally, firms are permitted to make changes in detailed decisions that alter high-level policies. We then plot in Figure 3 the difference between the performance of the orthodox firm and the performance of the corresponding heterodox firm. One can think of this difference as the hidden cost a firm bears if it takes its candidate solution too seriously. The figure shows that this cost depends sensitively on the quality of a firm's representation. Firms with very good representations incur no costs of orthodoxy; they are in such promising parts of the target landscapes that they can perform well without crossing policy borders. Heterodox firms have little incentive to cross policy borders, so the search behaviors of orthodox and heterodox firms are effectively the same. Firms with poorer representations, in contrast, pay a heavy price for holding on too firmly to their analogies. Indeed, in results not reported here, we find that firms that take poor analogies seriously typically fare worse than heterodox local searchers; the penalty of orthodoxy outweighs the benefit of analogical guidance.

Structural characteristics

The power of analogy depends not only on the representations, experience, and orthodoxy of the management team, but also on the structure of the target landscape. We explore the effect of both 'vertical' and 'horizontal' structural changes on our results. By vertical structure, we mean the relationship between the number of detailed decisions per policy (D) relative to the number of higher-order policies (P). By horizontal structure, we mean the intensity of interaction within each policy domain (K_w) and the intensity of interaction across domains (K_b).

To explore vertical structure, we compare the performance of analogizing firms to that of local searchers in two settings: $P = 5$, $D = 3$ vs. $P = 3$, $D = 5$. The total number of detailed decisions is 15 in both cases, but the high-level guidance given by analogy is more thorough when policy domains are numerous and narrow ($P = 5$, $D = 3$). Analogical reasoning is more powerful in this situation, as the results in Table 5 show. Comparing across rows in Table 5, we see that the differential created by more thorough guidance diminishes as the quality of the representation deteriorates. Detailed guidance is less valuable if it is based on misleading keys.

To explore horizontal structure, we first examine the effects of K_b and K_w on landscape topography and then turn to their effects on the performance of analogizing firms. One can think of each landscape as consisting of 2^P policy domains, with 2^D configurations of detailed decisions within each domain. Tables 6a and 6b reveal typical topography within and across policy domains, as a function of K_b and K_w. To construct the tables,

Copyright © 2005 John Wiley & Sons, Ltd.

Table 5. Effect of policy breadth on analogizer's long-run performance advantage

Long-run performance advantage over a local searcher for an analogizer with representation ...	$P = 5$ $D = 3$	$P = 3$ $D = 5$	Differential
⟨1 2 3⟩	0.162 (0.002)	0.109 (0.002)	0.053*
⟨1⟩	0.102 (0.002)	0.073 (0.002)	0.029*
⟨4 3 2⟩	0.085 (0.002)	0.054 (0.002)	0.031*
⟨4⟩	0.066 (0.002)	0.048 (0.002)	0.018*

$K_b = 1$, $K_w = 0$, BREADTH = DEPTH = 1; all firms are orthodox. Each figure is an average over 3000 firms on 300 families of landscapes. Each number in parentheses is the standard deviation of the average figure above it. * Indicates that a differential is statistically significant at the 0.1% level.

Table 6a. Absolute change in performance associated with a within-policy change in a detailed decision

	$K_w = 0$	$K_w = 0.5$	$K_w = 1$
$K_b = 0$	0.056	0.074	0.087
$K_b = 0.5$	0.051	0.069	0.085
$K_b = 1$	0.050	0.068	0.083

$P = 3$, $D = 3$; each figure an average over 40 firms walking at random for 125 periods on 100 landscapes.

Table 6b. Standard deviation across average performance in each policy domain for a typical landscape

	$K_w = 0$	$K_w = 0.5$	$K_w = 1$
$K_b = 0$	0.069	0.060	0.055
$K_b = 0.5$	0.092	0.081	0.061
$K_b = 1$	0.110	0.088	0.071

$P = 3$, $D = 3$; each figure an average over 40 firms walking at random for 125 periods on 100 landscapes.

we release firms in each policy domain on a large number of landscapes, let each walk at random within a domain, and record performance during the random walk. Table 6a shows the average absolute value of the performance change caused by altering one detailed decision within a policy domain; this reflects within-policy ruggedness. Table 6b reports the standard deviation across the average level of performance within each domain for the typical landscape; this captures across-policy ruggedness. Within-policy ruggedness is driven upward by increases in K_w, while across-policy ruggedness is boosted by K_b and mitigated by K_w.

Together, Tables 6a and 6b paint a clear picture of landscape topography. When $K_b = K_w = 0$, there are no interactions among decisions and the landscape is relatively smooth, with changes in individual decisions having relatively little impact on overall performance. As K_b rises with K_w still 0, the surface becomes more rugged. In particular, plateaus defined by policy configurations arise. Within a policy configuration, the surface is relatively smooth; a change in a detailed choice that does not change a policy alters only the contribution of that decision in isolation. In contrast, a change in a detailed choice that does alter a policy—a step over the edge of the plateau—causes many contributions to change. The result is a surface with internally smooth plateaus of quite different elevations. In contrast, as K_w rises with K_b fixed, each plateau becomes internally rugged.

This sets the stage for understanding the effects of K_b and K_w on analogical reasoning. The surface plot in Figure 4 shows, for various combinations of K_b and K_w, the long-run performance differential between an analogizing firm with representation ⟨1 2 3⟩ and a local searcher. ($P = 3$, $D = 3$, and DEPTH = BREADTH = 1.) When $K_b = K_w = 0$, the landscape is smooth, both the analogizer and the local searcher fare well, and as a result the differential is modest. Indeed, a differential exists only because firms are constrained to abide by their original policies. If firms were heterodox, both would attain the global peak, and the differential would be precisely zero.[9] As K_b rises with K_w

[9] While it is true that analogy offers no long-run advantage when target landscapes are smooth ($K_b = K_w = 0$) and firms are heterodox, this outcome masks important dynamics. Though all firms eventually scale the global peak in such a situation, analogical reasoning helps firms with good representations to do so quickly. As a result, analogizing firms enjoy a transient advantage. The better is the firm's representation, the bigger is this advantage. (Simulation results that illustrate this effect are available from the authors on request.) There are numerous, unmodeled reasons to believe that, in reality, such an advantage might be of lasting importance. Discovering a strong competitive position first (Lieberman and Montgomery, 1988) may be especially important if the first firm to arrive at the position can deter others from copying its configuration of choices; if customers are loyal and reluctant to switch once others reproduce the position; if the selection environment is so vigorous that slow

Copyright © 2005 John Wiley & Sons, Ltd.

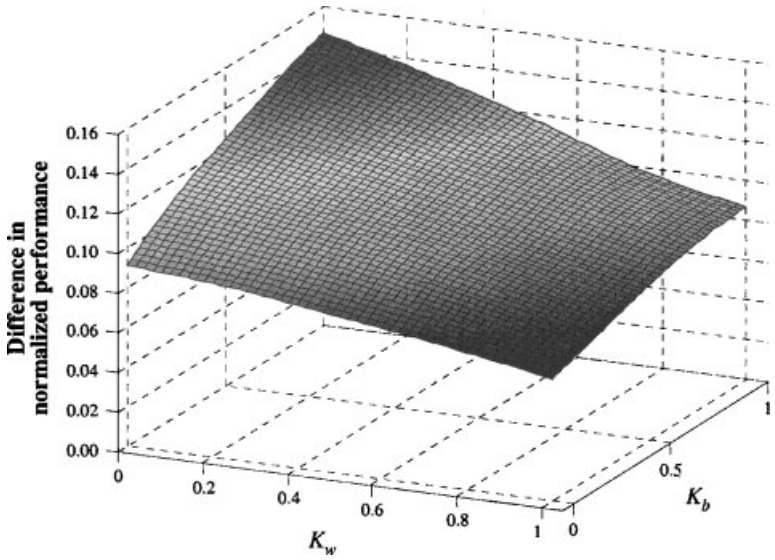


Figure 4. Long-run advantage of analogizer with representation ⟨1 2 3⟩ as a function of K_w and K_b

low, internally smooth plateaus of different heights emerge on the target landscape. In such a setting, it is quite valuable to start one's local search in a favorable policy domain—on a promising plateau. Accordingly, analogy is very powerful for high K_b and low K_w. In contrast, an increase in K_w undermines the power of analogy: interactions within policy domains make each plateau rugged, and even a firm with a good representation, which begins its local search in a favorable policy domain, is likely to get stuck on a low local peak. Good analogies in our model give policy-level guidance, but if guidance at that level is not sufficiently specific to lead a firm to success, then analogical reasoning loses much of its force.

Overall, we find that analogy is least powerful in fully decomposed settings ($K_b = 0$, $K_w = 1$) or nearly decomposed settings (K_b low, K_w high); in such settings, the primary challenge facing management is to achieve success within each of several, fairly independent policy domains, and analogy—at least as we have modeled it—offers little toward this end. Analogical reasoning is also of limited use in situations of full independence ($K_b = K_w = 0$), where incremental logic alone will deliver a good configuration of choices, and in situations of very high interdependence of detailed decisions ($K_w = 1$), where proliferating local peaks trap analogizers and local searchers alike. Analogy is most powerful when cross-policy interactions dominate (K_b high, K_w low)—that is, when the plateaus defined by policy configurations are internally smooth, but heights are sufficiently varied across plateaus that there is a danger of being stranded on a low plateau. A number of scholars have argued that piece-by-piece, subsystem-by-subsystem learning is most effective in nearly decomposable systems (Simon, 1962; Baldwin and Clark, 2000; Frenken, Marengo, and Valente, 1999). Our results suggest that analogical reasoning offers its greatest relative benefit precisely where these approaches fail, in systems that are not nearly decomposable.

DISCUSSION AND CONCLUSION

Discovering an effective competitive position is a difficult endeavor. It is difficult because the mapping from strategic decisions to performance is typically complex and, especially in novel settings, unknown to the decision-makers. Though cognizant of such difficulties, positioning scholars emphasize the role of deductive reasoning and rational choice in the origin of positions (Porter, 1996; Ghemawat, 1999). In contrast, evolutionary theorists highlight the bounds of individual

movers don't survive their temporary disadvantage; or if the environment changes so rapidly that one must reap the benefits of an advantage quickly. We speculate that analogical reasoning might be especially valuable in such settings, even in the absence of interactions among decisions.

Copyright © 2005 John Wiley & Sons, Ltd.

rationality and posit that effective positions emerge through a mix of luck and experiential, local search, thus leaving little space for the cognition of managers (Nelson and Winter, 1982). Although recognizing the merits of both such perspectives, we hesitate to ascribe all effective positions to either the omniscience assumed in economic analysis or the myopia of experiential learning.

We put forth a model of the origin of strategies that, while recognizing the limits of managerial rationality and the intelligence of local search, also respects the power of managerial cognition. Indeed, the notion of bounded rationality, which we take as a cornerstone of our conceptual apparatus, does not rule out the possibility of intendedly rational choice (March and Simon, 1958). Bounded rationality suggests that thinking is typically premised on simplified cognitive representations of the world (Simon, 1991). As boundedly rational actors, managers create cognitive simplifications of their decision problems and come up with solutions on the basis of such simplifications. These solutions, in turn, may imprint subsequent efforts at local search, thus playing a central role in the discovery of strategic positions (Gavetti and Levinthal, 2000). This perspective, which represents a middle ground between positioning and evolutionary arguments, suggests that the roots of superior competitive positions may lie in the cognition of managers, particularly in the way they represent the world.

The effect of cognition on managerial action is an enormous area of inquiry. We cut into this topic by looking at a particular, important type of action: the creation of a competitive position in an unfamiliar industry. This type of action confronts managers with daunting complexity so it is natural to expect cognition—a fundamental instrument to handle complexity—to play a major role. It also forces decision-makers to cope with novelty. In novel situations, wisdom from prior experience in other contexts can be particularly powerful. For this reason, we focus on a particular vehicle for transporting experience across contexts: analogy.

Our conceptual model of analogical reasoning is straightforward. We surmise that, in any set of industries, a large number of underlying characteristics drive the relationship between firm action and performance. There are so many characteristics and their effects are so difficult to discern that boundedly rational actors focus their reasoning efforts on a subset of the characteristics. These subsets form representations. Some representations are more effective bases of reasoning than are others. A good representation distinguishes sets of similar payoff functions from one another, while a poor representation leaves very different payoff functions in the same category. In other words, a representation is a classification scheme. An effective scheme puts similar objects in the same class and different objects in distinct classes.

Armed with a good representation and adequate experience, a firm is well prepared to draw a candidate solution from a germane source and apply it to a target industry. Accordingly, our simulation model shows the best performance among firms with high-quality representations—good classification schemes. That said, we also find it better to take guidance from some source, even one based on a poor representation, than to start one's local search at a randomly assigned configuration. There is some chance that, purely by luck, the source will prove to be a good one even though it is based on a poor representation. This power of 'false' analogies is reminiscent of Weick's (1990) tale of a Hungarian military reconnaissance unit. Lost in the snowy Alps, the troops prepare for the worst. Then one soldier discovers a map in his pocket. Once equipped with the map, the unit outlasts the storm, finds its bearings, and returns to safety. Only later does a commanding officer realize that the map is of the Pyrenees, not the Alps. The tale is often interpreted in terms of presence of mind and motivation: the map calmed the troops and moved them to take coordinated action. We would suggest a second possibility: perhaps parts of the Pyrenees and parts of the Alps truly do resemble one another. The soldiers may have received a good candidate solution, though only by chance, just as do some (though not all) of our firms with poor representations. In business also, it is possible to have a poor representation yet obtain a good candidate solution. For instance, Lycos's partial integration of Tripod worked out well, but even some supporters of the decision felt the analogical reasoning itself was dubious (Gavetti and Rivkin, 2004). An intriguing avenue for future research is to examine what makes some analogies useful even if based on similarity along irrelevant dimensions. We suspect the key is to focus on dimensions that, even if irrelevant, are correlated with true drivers of success.

Copyright © 2005 John Wiley & Sons, Ltd.

Our results also shed light on the roles of broad and deep experience. We find breadth and depth of experience to be valuable only if a manager has a good representation—a valid system for categorizing environments and classifying lessons learned. In addition, beyond a modest level of depth, performance is not sensitive to depth. As a result, if top management of a diversified firm believes that business unit managers have a good grasp of what factors drive the choice–performance relationship in their particular business, they should be more willing to rotate managers across divisions, to invest in and exploit breadth of experience. Otherwise, it may be better to keep each business unit manager in a single business for a longer period, to develop depth.

A possibility we do not model, worthy of future research, is that experience in a wide range of industries may help a manager to build an understanding of what drives the relationship between choice and performance. Put simply, breadth may improve representations. As business school instructors who teach by the case method, we are sympathetic to this perspective. Much of what goes on in business strategy courses, we believe, is representation building. We aim to focus student attention on dimensions of environments that shape the relationship between action and outcome, and we do so in part by exposing students to a broad variety of cases. A framework such as the Five Forces (Porter, 1980) is helpful because it synthesizes across cases and identifies the dimensions that are most salient across the range of particulars. In doing so, the framework enhances pattern recognition. Frameworks that classify environments and firms along different dimensions—say, by the durability of underlying resources (Williams, 1992) or by the character of capabilities (Teece, Pisano, and Shuen 1997) rather than by the nature of competitive forces—might lead managers to see very different patterns.

More broadly, the pervasive impact of representations in our findings affirms the importance of research that examines the mental models of strategists (e.g., Porac, Thomas, and Baden-Fuller, 1989; Huff, 1990; Huff and Jenkins, 2002). Especially intriguing is the question of where representations come from (other than business school curricula). Future empirical efforts might examine the processes by which management teams develop, track, and alter their beliefs about what characteristics distinguish similar settings from different ones. Modeling efforts might expose simulated firms to a series of decision problems, not the single problem we model here, and allow management teams to alter their representations as they update their beliefs about the true $\boldsymbol{X}_{\mathbf{PROB}}$. Models might also incorporate shocks to $\boldsymbol{X}_{\mathbf{PROB}}$. Our speculation is that established teams with solid beliefs about $\boldsymbol{X}_{\mathbf{PROB}}$ will find such shocks especially hazardous. Such a team may very well persist in paying attention to observable characteristics that used to be relevant but are no longer. As a result, they may draw analogies based on similarities that have become superficial.

Assessing the role and performance implications of analogies also requires analyzing how analogies are used. In the context of our model, we distinguish between orthodox and heterodox uses of analogy. Not surprisingly, orthodoxy is most costly for firms with poor representations. More surprisingly, orthodoxy provides no advantage over heterodoxy even when the analogy is based on a good representation. Good representations seed subsequent search efforts on such promising areas of the landscape that heterodox analogizers are not motivated to violate the policies identified by the analogy. This result suggests that analogy is more effective as an instrument to seed search efforts than as a means to constrain them. The psychological literature on analogy focuses on representations and experience as central elements underlying the quality of analogies (Holyoak and Thagard, 1995). Particularly in the context of organizational search efforts, we believe, that in addition to these fundamental elements, managers should pay explicit attention to the orthodoxy with which analogies are used.

Our final set of results shows that analogical reasoning produces the greatest long-run advantage over incremental local search in settings that are poorly decomposed. In these contexts, the high-level policy guidance offered by analogy is necessary to identify consistent policy configurations—policy configurations that cannot be recreated via incremental search at the level of detailed choices. Others have touted the virtues of decomposable or nearly decomposable systems of choices (Simon, 1962; Baldwin and Clark, 2000). In nearly decomposable or modular systems, one can learn by means of parallel experiments at the subsystem level (Baldwin and Clark, 2000). Moreover, boundedly rational managers might be

Copyright © 2005 John Wiley & Sons, Ltd.

able to apply deductive reasoning at the subsystem level even if the system as a whole outstrips their processing abilities. However, not all systems of choices are neatly decomposable. We suggest that, for dealing with systems of choices that are inherently non-decomposable or have not yet been decomposed, reasoning by analogy is particularly powerful. Good representations underlie the transfer of powerful strategic solutions from managers' past experience. These solutions, in turn, allow a rich appreciation of the architecture of the strategic problem (Henderson and Clark, 1990). It is in poorly decomposed systems that this kind of architectural wisdom, and therefore the quality of the cognitive representations underlying the analogy, play a particularly important role. These results emphasize that analogy and, more generally, cognition are especially powerful in settings where parallel local search fails to guide organizational adaptation effectively.

More broadly, our analyses may help us better understand observed variation in the processes that underlie the origins of successful strategic positions. Our findings lead us to hypothesize, for instance, that cases in which local search is the main process guiding firms toward successful positions may well reflect some select structural characteristics of the industries where such companies succeed. In particular, such settings are likely to have intense interactions among decisions within policies. Rich sets of interactions among the detailed decisions lying below the surface of higher-level policies limit the power of *a priori* cognition in these settings. We believe that this linkage between the structural characteristics of the context and the mechanisms underlying the development of successful positions is a crucial one—one that future empirical studies should consider carefully.

Cognition in complex worlds inevitably involves simplification. The precise basis of simplification, however, is not inevitable. As academics and practitioners, we are often apologetic about operating in the space of such simplifications—the commonly ridiculed 'two-by-twos' of the strategy field. But as boundedly rational individuals, we cannot think in high-dimensional spaces. The relevant question is not whether we conceive of complex strategic problems in terms of a few overarching variables, but rather what those variables will be. The choice of variables can have a major impact on performance, both directly and by altering the influence of other factors such as managerial experience and orthodoxy. Our findings support this perspective in the case of analogical reasoning, and we speculate that it will also hold true for other forms of managerial cognition. Our hope is that rigorous analysis of cognition will help bridge the chasm between rational, positional perspectives on strategy and behavioral, evolutionary approaches. Understanding how firms identify effective competitive positions requires both perspectives. With the current work, we try to provide some substantiation of that link and a platform on which others can build.

ACKNOWLEDGEMENTS

For helpful comments, we are grateful to Nicolaj Siggelkow, anonymous referees, and seminar audiences at Babson, Harvard Business School, MIT's Sloan School, the University of Calgary, UCLA, the University of Toronto, and the Wharton School. We also thank Howard Brenner for computer programming efforts, Simona Giorgi and Elizabeth Johnson for research assistance, and the Division of Research of Harvard Business School for generous funding. Errors remain our own.

REFERENCES

Baldwin CY, Clark KB. 2000. *Design Rules: The Power of Modularity*. MIT Press: Cambridge, MA.

Cyert R, March J. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall: Englewood Cliffs, NJ.

Fleming L, Sorenson O. 2001. Technology as a complex adaptive system: evidence from patent data. *Research Policy* **30**: 1019–1039.

Fleming L, Sorenson O. 2004. Science as a map in technological search. *Strategic Management Journal*, Special Issue **25**(8–9): 909–928.

Frenken K, Marengo L, Valente M. 1999. Interdependencies, near-decomposability and adaptation. In *Computational Techniques for Modelling Learning in Economics*, Brenner T (ed). Kluwer: Dordrecht; 145–165.

Gas Daily. 2000. Broadband to be as big as Enron in five years. 5 April.

Gavetti G, Levinthal D. 2000. Looking forward and looking backward: cognitive and experiential search. *Administrative Science Quarterly* **45**: 113–137.

Gavetti G, Rivkin J. 2004. Rationality and plasticity over time: toward a grounded theory of the origin of strategies. Working paper, Harvard Business School.

Ghemawat P. 1999. *Strategy and the Business Landscape*. Addison-Wesley: Reading, MA.

Copyright © 2005 John Wiley & Sons, Ltd.

Gick ML, Holyoak KJ. 1980. Analogical problem solving. *Cognitive Psychology* **12**: 306–355.
Gilovich T. 1981. Seeing the past in the present: the effect of associations to familiar events on judgments and decisions. *Journal of Personality and Social Psychology* **40**: 797–808.
Gragg CI. 1940. Because wisdom can't be told. *Harvard Alumni Bulletin* **43**(2): 78–84.
Hast A (ed). 1992. *International Directory of Company Histories.* St James Press: London.
Henderson RM, Clark KB. 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly* **35**: 9–30.
Holyoak K, Thagard P. 1995. *Mental Leaps: Analogy in Creative Thought*. MIT Press: Cambridge, MA.
Huff AS. 1990. *Mapping Strategic Thought.* Wiley: New York.
Huff AS, Jenkins M. 2002. *Mapping Strategic Knowledge.* Sage: Thousand Oaks, CA.
Kauffman SA. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press: Oxford.
Kauffman SA. 1995. *At Home in the Universe.* Oxford University Press: Oxford.
Levinthal D, March J. 1993. The myopia of learning. *Strategic Management Journal*, Winter Special Issue **14**: 95–112.
Levinthal D. 1997. Adaptation on rugged landscapes. *Management Science* **43**: 934–950.
Lieberman MB, Montgomery DB. 1988. First-mover advantages. *Strategic Management Journal*, Summer Special Issue **9**: 41–58.
March J, Simon HA. 1958. *Organizations*. Wiley: New York.
McKelvey B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: strategies for rugged landscapes. *Organization Science* **10**: 294–321.
Mintzberg H. 1978. Patterns in strategy formation. *Management Science* **24**: 934–948.
Nelson R, Winter S. 1982. *An Evolutionary Theory of Economic Change*. Belknap: Cambridge, MA.
Neustadt RE, May ER. 1986. *Thinking in Time: The Uses of History for Decision-Makers*. Free Press: New York.
Pascale RT. 1984. Perspectives on strategy: the real story behind Honda's success. *California Management Review* **26**(3): 47–72.
Perkins EJ. 1999. *Wall Street to Main Street*. Cambridge University Press: New York.
Porac JF, Thomas H, Baden-Fuller C. 1989. Competitive groups as cognitive communities: the case of Scottish knitwear manufacturers. *Journal of Management Studies* **26**: 397–416.
Porter ME. 1980. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. Free Press: New York.
Porter ME. 1985. *Competitive Advantage: Creating and Sustaining Superior Performance*. Free Press: New York.
Porter ME. 1996. What is strategy? *Harvard Business Review* **74**(6): 61–78.
Rivkin JW. 2000. Imitation of complex strategies. *Management Science* **46**: 824–844.
Rivkin JW. 2001. Reproducing knowledge: replication without imitation at moderate complexity. *Organization Science* **12**: 274–293.
Rivkin JW, Siggelkow N. 2002. Organizational sticking points on NK landscapes. *Complexity* **7**: 31–43.
Rivkin JW, Siggelkow N. 2003. Balancing search and stability: interdependencies among elements of organizational design. *Management Science* **49**: 290–311.
Salter MS, Levesque LC, Ciampa M. 2002. Innovation corrupted: the rise and fall of Enron. Harvard Business School Working Paper 02–102.
Siggelkow N. 2002. Evolution towards fit. *Administrative Science Quarterly* **47**: 125–159.
Simon HA. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**: 99–118.
Simon HA. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* **106**: 467–482.
Simon HA. 1991. Bounded rationality and organizational learning. *Organizational Science* **2**: 125–134.
Slywotzky AJ, Morrison DJ. 1999. *Profit Patterns: 30 Ways to Anticipate and Profit from Strategic Forces Reshaping Your Business*. Times Business: New York.
Stemberg TG. 1996. *Staples for Success: From Business Plan to Billion-dollar Business in Just a Decade*. Knowledge Exchange: Santa Monica, CA.
Teece DJ, Pisano G, Shuen A. 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* **18**(7): 509–533.
Thagard P. 1996. *Mind: Introduction to Cognitive Science*. MIT Press: Cambridge, MA.
Weick K. 1990. Cartographic myths in organizations. In *Mapping Strategic Thought*, Huff AS (ed). Wiley: New York; 1–10.
Westhoff FH, Yarbrough BV, Yarbrough RM. 1996. Complexity, organization, and Stuart Kauffman's *The Origins of Order*. *Journal of Economic Behavior and Organization* **29**: 1–25.
Williams JR. 1992. How sustainable is your competitive advantage? *California Management Review* **34**(3): 29–51.

APPENDIX: DETAILED DESCRIPTION OF HOW FAMILIES OF LANDSCAPES ARE GENERATED

For a given set of values for P, D, X, K_w, K_b, and $\boldsymbol{X}_{\mathbf{PROB}}$, the computer generates an influence matrix—a matrix that records which choices, policies, and observable characteristics influence the contribution of each detailed choice. Figure 1 shows what the matrix might look like for a case in which $P = 3$, $D = 3$, $X = 4$, $K_w = 1$, $K_b = 0.33$, and $\boldsymbol{X}_{\mathbf{PROB}} = (1, 0.5, 0.5, 0)$. A bullet (•) in the matrix indicates that the column element of

Copyright © 2005 John Wiley & Sons, Ltd.

the matrix influences the contribution of the row element. The second row, for instance, indicates that the contribution of the second detailed choice is affected by the resolutions of choices 1, 2, and 3, the resolution of policy 3, and the states of the first and third observable characteristics. Because $K_w = 1$, there is thorough interaction within each policy domain. Because $K_b = 0.33$, the typical decision is affected by one-third of the other policies. The first observable characteristic is highly influential, affecting the contributions of all detailed decisions because the first element of $\boldsymbol{X}_{\mathbf{PROB}}$ is 1. The fourth observable characteristic is altogether irrelevant since the fourth element of $\boldsymbol{X}_{\mathbf{PROB}}$ is 0.

Once the influence matrix is set, the computer generates target and source landscapes. That is, it assigns a pay-off to each of the $2^{P \times D}$ possible configurations of choices on each of the 2^X landscapes. The contribution of each detailed choice depends on the resolution of that choice, the resolution of other choices and policies, and the state of observable characteristics. For each possible combination of the influential factors, the computer draws a contribution at random from a uniform U[0, 1] distribution. In Figure 1, for instance, the second choice is affected by six factors (as indicated by the six bullets in the second row). Each of these factors can be resolved in two ways, so the computer draws 2^6 possible contributions for the second choice: one corresponding to each of the possible configurations of the six factors. The overall performance associated with a particular configuration of detailed choices and observable characteristics is then the average over the $P \times D$ contributions:

$$\text{Performance}(\mathbf{d} \text{ given a set of observable characteristics}) = \left[\sum_{i=1}^{P \times D} \text{contribution of decision } i \text{ given } \mathbf{d} \text{ and observable characteristics}\right] \Big/ (P \times D).$$

Copyright © 2005 John Wiley & Sons, Ltd.