Coordinated Exploration: Organizing Joint Search by Multiple Specialists to Overcome Mutual Confusion and Joint Myopia

Thorbjørn Knudsen and Kannan Srikanth

Administrative Science Quarterly 2014 59: 409 originally published online 29 May 2014
DOI: 10.1177/0001839214538021

JOHNSON
Cornell University

Administrative Science Quarterly
59 (3)409–441

DOI: 10.1177/0001839214538021
asq.sagepub.com
SAGE

Coordinated Exploration: Organizing Joint Search by Multiple Specialists to Overcome Mutual Confusion and Joint Myopia

Thorbjørn Knudsen[1] and Kannan Srikanth[2]

Abstract

In this paper, we use an agent-based simulation model to investigate how coordinated exploration by multiple specialists, as in new product development, is different from individual search. We find that coordinated exploration is subject to two pathologies not present in unitary search: mutual confusion and joint myopia. In joint search, feedback to one agent's actions is confounded by the actions of the other agent. Search therefore leads to increasing mutual confusion because agents are unable to learn from feedback to correct their faulty mental models of the search space. Incorrect beliefs held by one agent lead to mistakes, and because it is unclear which agent was wrong, this confuses the other agent, either into revising (correct) beliefs or holding on to (incorrect) beliefs. Sharing knowledge aligns specialists' mental models and counters mutual confusion by inducing coordination around particular search regions. Yet that very effort increases joint myopia, as agents prematurely reinforce each other into choosing from an increasingly narrow portion of the search space. In the extreme, high levels of shared knowledge induce agents to abandon their distinct search approach in favor of a lower common denominator. In coordinated exploration, increasing coordination efforts (such as by increasing communication) reduces mutual confusion but simultaneously increases joint myopia. Efforts to reduce joint myopia, such as by slow learning or lower levels of knowledge transfer, however, automatically increase mutual confusion. As modeled in our simulation, successful joint search needs to balance these two effects. Our results suggest that because unitary-searcher models abstract from epistemic interdependence, their predictions are potentially misleading for coordinated exploration.

[1] Strategic Organization Design Unit, University of Southern Denmark
[2] Strategy Area, Indian School of Business

Keywords: coordinated exploration, organization design, coordination, joint search, new product development, mutual adaptation, co-evolution of mental models

Models of search and learning are foundational to the behavioral theory of organizations (March and Simon, 1958). Typical models of organizational search consider firms to be unitary actors whose behavior is constrained by cognitive limitations (Cyert and March, 1963; Nelson and Winter, 1982). Their objective is to identify tall peaks on a search landscape—optimal choices—by balancing exploration of new domains and exploitation of known domains (Levinthal, 1997; Siggelkow and Levinthal, 2003; Ethiraj and Levinthal, 2004). These models' assumption that organizational search involves one actor is a useful abstraction and baseline that has facilitated a thorough examination of problems and solutions related to balancing exploration and exploitation. But such models fail to take into account that organizational search often involves a joint effort by specialists from different domains who need to coordinate their search efforts. These are problems of coordinated exploration.

In coordinated exploration, whereas the specialists are responsible for search in their own domains, their payoffs depend on the choices of other specialists. In traditional game theory, interdependent actors know all their choices, the choices available to others, and all joint payoffs. Boundedly rational actors, however, know neither all the available choices in their own or in others' domains nor their payoffs. For this reason, the choices of each can influence joint outcomes in ways that the specialists can neither anticipate nor understand fully, giving rise to the problem of coordinated exploration. Models of search performed by single actors ignore the influence of interdependence between agents that is fundamental to a theory of search involving multiple actors, and therefore they offer few insights into problems of coordinated exploration.

Problems of coordinated exploration are ubiquitous in organizations. New product development involving multiple specialists working together is an important example. Consider the problem of designing windmills. Rotor blade design is critical for the efficiency of a windmill and involves finding the right combination of structural and aerodynamic characteristics to achieve desired performance. Structural properties enable blades to withstand adverse weather conditions. Good aerodynamics is critical for efficiently converting wind energy to power. Light, curved blades generally have good aerodynamics but poor strength. Two specialists jointly solve this innovative search problem: a structural mechanics engineer and an aerodynamics engineer who have little knowledge of each other's domains.

Product positioning is another example of coordinated exploration. Marketing experts typically explore the domain of identifying the customer, dividing a market into segments that respond in homogenous ways to changes in the marketing mix, such as product, price, promotion, and place. At the same time, product experts explore the domain of what the firm can offer its customers, devising product or service offerings that deliver different combinations of costs and features. A value proposition lies at the intersection of these choices and needs to be discovered jointly by the two departments. Other organizational problems that display the properties of coordinated exploration include balancing synergies vs. responsiveness between business units that operate in different markets, and managing interagency service delivery during disaster situations.

The essential feature of these problems is that their solutions involve more than one actor, each searching in a different aspect of the problem space, and the actors' separate decisions are integrated to generate a joint payoff matrix that the actors do not know and cannot conceive in advance. Coordinated exploration therefore involves solving two interdependent problems: (1) search—the specialist problem of searching for new valuable alternatives in a particular domain, and (2) coordination—the problem of choosing an alternative from a particular domain that is jointly attractive to all agents, though it may not be optimal for any one of them.

Organizations frequently engage in problems of coordinated exploration because attempting to benefit from the division of labor is one of the fundamental reasons to organize. As soon as the environment becomes too complex for a single individual to comprehend, organizations take advantage of the economies of specialization offered by the division of labor. But the division of labor simultaneously generates the need for integration, or knowledge sharing. Typically, organizations rely on differentiation and then integration to manage this tradeoff (Lawrence and Lorsch, 1967). In these cases, experts or specialized organizational units search in individual domains and then integrate their results into a joint solution and observe its payoff.

Prior work in organization theory has implicitly assumed that results from individual search models can be generalized to problems of coordinated exploration. When multiple agents are engaged in coordinated exploration, however, it is a case of epistemic interdependence, a situation in which one agent's optimal choices depend on accurately predicting another agent's actions (Puranam, Raveendran, and Knudsen, 2012). It is well known that communication, or more generally shared knowledge, is necessary to coordinate epistemic interdependence, which has led prior theorists to suggest that coordinated exploration is nothing but individual search coupled with high levels of communication between the agents (cf. Tushman and Nadler, 1978). But empirical work suggests that a high level of communication between specialists is not reliably associated with good outcomes (Montoya-Weiss and Calantone, 1994; Brown and Eisenhardt, 1995; Sine, Mitsuhashi, and Kirsch, 2006). Therefore it is unclear how epistemic interdependence influences boundedly rational organizational search. In this sense, prior work on joint search is fundamentally incomplete because it does not account for epistemic interdependence or how coordination is achieved.

In contrast, models in game theory do take into account epistemic interdependence, but they neglect search. Game-theoretic models typically assume that agents make a choice over a known state-space and do not consider conditions in which the agents' knowledge of the nature of the state-space is evolving. Typical game-theoretic work does not model evolving game structures, which is an important property in many real-world problems. Models by Lounamaa and March (1987) and Puranam and Swamy (2011), though boundedly rational because actors do not know the actions/payoffs available to the interdependent agent, also assume that actors know all actions available to them. We thus have little knowledge about coordinated exploration, even though it is an important problem for organizations. While prior theory has concentrated on the impact of integration mechanisms—ways of knowledge sharing—on coordination among multiple specialists, here we use our model to explore the impact of integration on search itself to better understand coordinated exploration.

ORGANIZING SEARCH FOR COORDINATED EXPLORATION

Organization scholars have studied how to manage the specialization–coordination tradeoff for over 50 years. March and Simon (1958) suggested that organizations achieve coordination by two generic means: plan and feedback. When interdependence is stable and predictable, plan-based coordination mechanisms such as standard operating procedures, rules, and routines are effective and efficient. When the nature of interdependence is unknown or unstable, coordination is achieved by feedback or mutual adjustment (Thompson, 1967). Coordinated exploration is important only when the nature of interdependence is unknown (or even unknowable) because of bounded rationality.

The information processing view of organizations builds on these fundamental insights to understand how organizations can be designed to operate effectively in situations with differing levels of interdependence and uncertainty. This view suggests that the coordination capacity of the organization must match its coordination needs (Galbraith, 1977), and therefore highly interdependent work must be structured to maximize opportunities for information transfer (Lawrence and Lorsch, 1967; Tushman and Nadler, 1978).

Though it is well accepted that shared knowledge is necessary to coordinate under epistemic interdependence (March and Simon, 1958; Puranam, Raveendran, and Knudsen, 2012), we also know that, in practice, it is very difficult to develop such shared knowledge among specialists (Cronin and Weingart, 2007). This is because the boundaries of specialization are also natural interpretive barriers that make it difficult to generate shared knowledge (Lawrence and Lorsch, 1967; Dougherty, 1992; Heath and Staudenmayer, 2000). Organizations use different kinds of "integration mechanisms" to develop shared knowledge (Lawrence and Lorsch, 1967; Clark and Fujimoto, 1991; Iansiti, 1995; Hoopes and Postrel, 1999). But different integration devices likely generate shared knowledge in different ways. For example, frequent communication vs. infrequent communication may lead to different patterns of shared knowledge over time. This means that the type of integration mechanism used is likely to affect the outcomes from coordinated exploration in two ways: (1) emergent shared knowledge will direct search, and (2) the resulting sampling of the search space will influence what new knowledge is acquired and shared (Denrell and March, 2001). The aggregation of these effects across interdependent actors can fundamentally change joint search behavior from individual search in ways that are currently undertheorized.

Prior work on search provides little guidance on these issues. As Knudsen and Levinthal (2007) observed, most models of organizational search are nonorganizational—they assume a single actor. Prior models of multi-agent search ignore epistemic interdependence and have sidestepped issues of coordination (see Rivkin and Siggelkow, 2003; Siggelkow and Levinthal, 2003; Fang, Lee, and Schilling, 2010). For example, these studies underscore the importance of "slow learning" in balancing exploration and exploitation, but the concept of shared knowledge that dominates the empirical literature is entirely absent from extant models of joint search. Thus, in the case of mutual adjustment between two agents, it is unclear if an increase in knowledge transfer between these agents will increase or reduce their joint exploration and if it will improve or reduce their chances of identifying the global peak. It is also unclear whether agents' efforts to explore via slow learning perturb their efforts to maintain shared knowledge and achieve coordination. To answer these questions, we need a model that takes into account how increasing shared knowledge affects agents' search behavior. Because prior joint search models do not formally model agents' knowledge or spell out the procedures by which agents influence each other's knowledge, they do not offer any predictions about organizing coordinated exploration. We construct such a model and consider three important elements that influence the level of shared knowledge: the level of communication, the extent to which agents are specialists vs. generalists, and the amount of exploration that agents engage in.

Effect of Communication between Agents

Lawrence and Lorsch (1967) argued that as task environments become more complex, specialized ''differentiated'' units become necessary to attend to specific environmental attributes. Differentiation refers to the differences across organizational subunits that arise as a consequence of their local adaptation to unit-specific tasks and environments. Depending on the demands of the environment, the actions of the differentiated units need to be more or less integrated for the organization to achieve desirable outcomes. The most complex environments demand both high levels of differentiation across subunits and high levels of integration among these units, giving rise to the problem of coordinated exploration.

According to the information-processing theory of organizations, highly interdependent work must be organized such that there is a high level of communication between the agents (Galbraith, 1977; Tushman and Nadler, 1978), which increases the level of shared knowledge among the agents, thereby promoting coordination (Simon, 1947; Schelling, 1960; Srikanth and Puranam, 2011). Even though this theory is about coordination rather than search, it has been extensively used to make predictions about coordinated exploration. For example, it is almost axiomatic in the new product development literature that a higher level of information transfer between agents is associated with better performance (see reviews by Brown and Eisenhardt, 1995; Krishnan and Ulrich, 2001). It should be noted that the arguments made by the information processing theory concern efficiency; organizing work with greater amounts of information transfer than necessary would be effective in uncovering good solutions but more expensive (Thompson, 1967; Galbraith, 1977; Tushman and Nadler, 1978). The prediction from this stream of work can be summarized in the hypothesis below:

Hypothesis 1a (H1a): In problems of coordinated exploration, specialist agents with a higher level of communication will be more likely to identify high-value combinations than specialist agents with a lower level of communication.

The above prediction, however, is not uncontroversial. For example, empirical work in the new product development literature suggests that intense communication between different specialists may lead to poor innovative outcomes (Tyre and Hauptman, 1992; Hauptman and Hirji, 1996; Song and Montoya-Weiss, 1998; Song, Thieme, and Xie, 1998; Song and Xie, 2000), though they do not always clarify why this may the case.[1] The literature on boundary spanners reaches similar conclusions, finding that projects with boundary spanners tend to perform better than projects without them (Tushman and Katz, 1980; Carlile, 2004), but even in large projects with high levels of interdependence, very few boundary spanners are required to achieve good outcomes (Tushman and Scanlan, 1981).

In contrast to H1a, these empirical findings suggest that a high level of information transfer is unnecessary and perhaps even harmful in coordinated exploration, but the mechanisms that underlie these findings are unclear. Many large-sample studies have hypothesized that a high level of communication should be associated with better performance but did not find that relationship. A plausible mechanism is that the difficulty in aligning mental models of different specialists leads to the pursuit of shortcuts and therefore lower performance (Tyre and Hauptman, 1992). Because significant effort is needed to transfer knowledge across specialists with incompatible mental models or ''thought worlds'' (Dougherty, 1992, 2001; Heath and Staudenmayer, 2000), such efforts are likely to be prone to conflict and delays. As a consequence, teams are more likely to pursue objectives that are minimally acceptable for all team members rather than to explore broadly to achieve more rewarding outcomes. For example, Davis and Eisenhardt (2011) explored innovations from high-technology alliances and found that a consensual leadership style promotes costly attempts at sharing information, which quickly leads these firms to adopt a lowest-common-denominator approach. These observations suggest the following competing hypothesis:

Hypothesis 1b (H1b): In problems of coordinated exploration, specialist agents with a higher level of communication will be less likely to identify high-value combinations than specialist agents with a lower level of communication.

Effect of Agents' Skills as Specialists vs. Generalists

One approach to the differentiation–integration tradeoff and the difficulty of communicating across specialists' boundaries is to employ agents with T-shaped skills—i.e., deep domain expertise in one domain, represented by the vertical bar of the T, and adequate knowledge in other domains, represented by the horizontal bar of the T (Iansiti, 1993; Leonard, 1995). Individuals with T-shaped skills have the ability to search for solutions to problems not only from their deep expertise but also taking into account how their choice is likely to interact with other constraints that a joint solution needs to satisfy. In the

[1] Montoya-Weiss and Calantone (1994) observed that few empirical studies use technological innovativeness as a measure of new product development success; instead, market share, financial success, and, most frequently, speed of development are used as proxies indicative of success. When the success of new product development projects is measured in terms of innovativeness, these studies suggest that facilitating very high levels of communication among the specialist agents may be associated with poor performance.

context of the windmill example, if the structural mechanics engineers have T-shaped skills, they are less likely to limit search for solutions to the strongest materials, such as steel, because they recognize that these also tend to be heavy and unlikely to generate much power. Employing individuals with T-shaped skills is therefore likely to be associated with successful problem solving across multiple domains (Madhavan and Grover, 1998). This suggests that agents with T-shaped skills are more likely to be successful than individual specialist searchers. Therefore we hypothesize:

Hypothesis 2 (H2): In problems of coordinated exploration, agents with T-shaped skills will be more likely to identify high-value combinations than will specialists.

Recent work, however, suggests that H2 may not always be true. Specifically, though the single searcher with T-shaped skills is likely to be more successful than the single specialist searcher, it is unclear whether a team of agents with T-shaped skills is more likely to be successful than a team of specialists. Recent empirical work finds that employing personnel with T-shaped skills (as opposed to specialists) is not necessarily associated with new knowledge creation or effective exploration in a new product development context (Lee and Choi, 2003; Tsai and Huang, 2008). In fact, though these studies hypothesized that employing personnel with T-shaped skills should have an impact on performance, they failed to find such an impact in their data. This suggests that H2 may not hold under some circumstances in problems of coordinated exploration.

Effect of Agent Exploration

The contrasting effects argued above arise from two problems inherent in joint search: the inability to coordinate and the lack of adequate exploration. Formal work that models firms' adaptation as search over a rugged landscape suggests that local adaptation traps firms in local peaks, and exploration is crucial for superior performance in such adaptation problems (Levinthal, 1997). Exploration aims to provide a basis for better choices in the future, as opposed to maximizing the immediate returns (Gittins, 1989). Unconstrained exploration, however, also leads to poor outcomes because the agents never exploit the promising alternatives that their exploration highlighted (Sutton and Barto, 1998). Studies of organizational search and learning have convincingly demonstrated the need to balance exploration with exploitation for superior performance.

Of course, agents need not explicitly engage in exploration activities. Contexts that undermine efficient adaptation by disrupting action-outcome-feedback linkages allow agents to "wander" in the search space, a process that automatically promotes exploration (March, 1991; Denrell and March, 2001). These "slow-learning" effects are considered to be particularly beneficial in complex environments that require broad exploration (March, 2006; Knudsen and Levinthal, 2007). These studies suggest the following hypothesis:

Hypothesis 3a (H3a): In problems of coordinated exploration, the extent of exploration has a curvilinear relationship to performance, such that agents who explore too little and agents who explore too much are less likely to find high-value solutions.

It is unclear, however, whether the above prediction derived from models of unitary search is accurate in problems of coordinated exploration, because the epistemic interdependence that characterizes coordinated exploration leads to two pathologies in search: mutual confusion and joint myopia. First, the feedback received by the agents is the joint payoff associated with both their and the other agent's actions. Therefore agents are unable to distinguish whether the positive or negative feedback outcomes are a consequence of their own action or the action of the interdependent other (Lounamaa and March, 1987; Puranam and Swamy, 2011). This impedes adaptation because agents are unable to learn from feedback to improve their mental models of action-outcome linkages. In other words, ambiguity in feedback promotes mutual confusion that in effect misleads agents into maintaining a flawed mental model of the task environment. Such flawed mental models may never be corrected in coordinated exploration because of enduring epistemic interdependence. This is in sharp contrast to models of unitary search in which epistemic interdependence is absent, and correction of mental models is possible.

Second, reducing mutual confusion will increase joint myopia. Mutual confusion can be countered if agents maintain mental models of the search space that are fully aligned with each other at every point in time. Alignment of mental models allows the interdependent actors to anticipate the others' expected actions, as in game-theoretic models. This is why it is commonly thought that high levels of communication among interdependent agents can facilitate coordination. But this comes at the cost of joint myopia—reducing exploration of the search space so the agents focus on a narrow portion of the landscape that both see as beneficial. As the agents receive more information about each other, they are more likely to choose actions that reliably take into account the others' preferences, thereby limiting search to areas known to be mutually beneficial. This narrowing of search, important for coordinating, necessarily comes at a cost: a more superficial understanding of other regions in the landscape that perhaps are more valuable. For example, behavioral economists have demonstrated that in coordination games, it is very difficult for the group to shift from a low-performing equilibrium to a high-performing equilibrium because it requires a coordinated shift among all participants (Van Huyck, Cook, and Battalio, 1997; Camerer, 2003). In the face of bounded rationality, agents do not know if any other better equilibrium exists and therefore cannot achieve a coordinated shift. According to theories of unitary search, such myopia can always be overcome with deliberate exploration strategies, as suggested by H3a. In problems of coordinated exploration, however, as in all coupled learning problems, exploration has the consequence of increasing mutual confusion by unintended interference in agents' learning and therefore is unlikely to lead to superior search outcomes. These arguments suggest the following competing hypothesis:

Hypothesis 3b (H3b): In problems of coordinated exploration, agents who explore individually are less likely to find high-value combinations than agents who do not explore individually.

The general mechanism that underpins our theory is the tradeoff between the need for agents to align their mental models in the face of epistemic interdependence and the need for adequate exploration. This tradeoff is challenging because of bounded rationality. Lack of aligned mental models results in poor performance because of mutual confusion, but alignment at the cost of joint myopia also leads to suboptimal outcomes. The greater the alignment in mental models, the more the team chooses options that are likely to result in a positive payoff given team members' (accurate) understanding of what the other actors are likely to choose. But this sensitivity to epistemic interdependence stifles exploration of the search space (whose payoff potential is unknown) by concentrating search efforts in the subspace that is of immediate mutual interest to all agents.

Coordinated exploration therefore needs to balance mutual confusion against joint myopia, which is likely influenced by the relationships among the following elements: (1) the agents' initial knowledge; (2) agents' learning based on each agent's own search efforts and the extent of integration with the other agent; and (3) the nature of the landscape. Understanding coordinated exploration therefore requires a careful trace of the evolving relationships among integration mechanisms, individual mental models, and the level of shared mental representations. As in much analysis of dynamic systems, a computational model is a suitable method for tracing these feedback-driven interacting relationships.

A MODEL OF COORDINATED EXPLORATION

To understand both the coordination and search aspects of coordinated exploration, we need to model an agent's knowledge or cognition as an information structure that bears some resemblance to the potentially unknowable real world. The agent's choices are informed by this information structure, or mental model, and it evolves over time with feedback. The heart of coordinated exploration is that agents are constrained by epistemic interdependence. In a dynamic perspective, this means that the evolution of one agent's information structure is significantly influenced by the actions of the other agent, which may be unobservable. In order to take into account such epistemic interdependence in a model of coordinated exploration, we need an approach to represent the current state of the agent's information structure that can direct search conditional on the actions of the other agent(s) and its evolution with feedback.

Partition Models: An Approach to Modeling Knowledge

The conception of knowledge as partitions in a state-space developed in Samuelson (2004) provides us with a handy tool to model the agent's evolving information structure. As explained in Samuelson (2004), economists model knowledge as a state-space, so what someone ''knows'' is represented as a set of partitions of that space. The more the partitions in an agent's information structure, the greater his or her knowledge about the space. Knowledge represents category learning: in the space where an ignorant agent sees only one category, a more knowledgeable agent can identify several nuanced categories, for example, distinguishing wood as pine, cherry, or oak. In search models, the task of the agent is to partition the information structure so that it is possible to identify elements in the knowledge space that likely correspond to objects that are actually of high value. Search proceeds by going through the current information partitions or if necessary by further partitioning the information structure.

For example, in the windmill design problem, bounded rationality implies that the two engineers do not know beforehand the full set of materials and shapes they could recombine. For instance, the initial knowledge of the structural materials engineer could be limited to three categories of candidate materials, wood, metal, and other—that is, only three partitions. Out of these, the engineer may select a promising candidate, say wood, and investigate it further—for example, discovering that there are two different types of wood, hard and soft. This represents an additional partitioning of his or her mental model of the search space as it pertains to wood. Among hardwoods, the engineer may discover that the oak behaves differently from elm, which in turn is different from pine. This increase in knowledge is represented as more partitions in the information structure. Note that in this example, the engineer's knowledge partitions are becoming more fine-grained in the subspace pertaining to wood, whereas his or her knowledge about other regions of the search space is unchanged.

The other interdependent agent in this task, the aerodynamics engineer, likely has partitioned the joint search space differently, depending on shapes, such as straight or curved. If the new partition available to the structural mechanics engineer, hard wood vs. soft wood, is not available to the aerodynamics engineer and vice versa, their knowledge partitions diverge. Because of this incongruence in mental models, the two agents may come to different conclusions about which region in the search space is attractive and therefore can make mutually inconsistent choices as each selects a solution that appears to be useful from one's own point of view but may in fact be jointly useless. This is the challenge of epistemic interdependence that agents involved in coordinated exploration need to solve.

As agents increasingly partition their knowledge structure, their mental models become increasingly incongruent, and they need to expend more effort in aligning their mental models. Partition models elegantly capture this tradeoff that increasing differentiation, modeled as more fine-grained partitions of the search space, requires increasing effort in integration, modeled as increasing alignment of the agents' partition structures.

Model Mechanics

To understand coordinated exploration, we model search in a two-dimensional landscape, such as structural mechanics and aerodynamics in the windmill example. The search landscape is a matrix in which each combination of the two technologies defines a coordinate with an associated payoff, as shown in figure 1. Two agents search in this landscape; the row agent chooses the row, and the column agent chooses the column. In our example, to create the next prototype of the new windmill, the structural mechanics engineer (row agent) chooses the material, such as wood vs. metal, and the aerodynamics engineer (column agent) chooses the shape, such as curved vs. straight. Once these agents have chosen in their own dimensions, a prototype is created with these joint properties (e.g., a windmill made of wood with straight blades), which is

Figure 1. Task environment (one-quarter the size of the actual task environment).

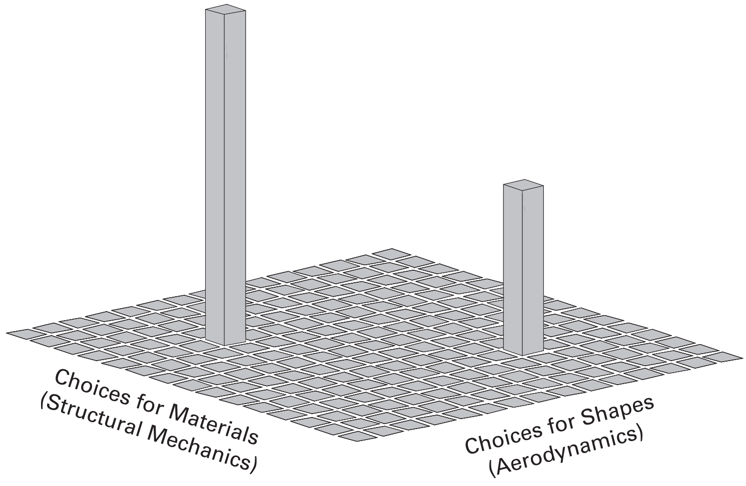

associated with a payoff. Figure 2 provides an overview of the baseline model and table 1 an overview of the parameters of the simulation. These are explained in greater detail below, and the parameters were chosen after numerous robustness checks to fine-tune the model.

Initial conditions: The search space. As shown in figure 1, in the baseline model, the search space is a matrix defined by 64 possible choices in two complimentary dimensions (row, column), and each of these 64 × 64 combinations is associated with a payoff.[2] We initially exercise our model with a landscape that contains two peaks of varying heights as shown in figure 1. This landscape emphasizes both search, because there are only two valuable peaks among the possible 4,096 combinations, and coordination, because each peak acts as a Nash equilibrium in this game. The shape of the landscape can materially affect successful strategies for coordinated exploration, so we ensure robustness using different landscapes that lay higher or lower emphasis on search vs. coordination.

Initial conditions: The agent's mental model of the search space. At t = 0, agents are endowed with a mental model of the search space. This mental model consists of two elements: (1) available decision choices and (2) payoffs associated with these choices. Our agents are boundedly rational and do not see all the 64 × 64 choices available to them beforehand, and as a consequence they do not accurately know the associated payoffs.

At the beginning of the simulation, agents do not have fine-grained partitions of the search space; they see only a very limited number of choices for each dimension. Figure 3 provides an example in the context of the windmill

[2] It may be helpful to draw a brief analogy of our model with the NK modeling structure. In our model, N = 2 because agents are searching only in two decision parameters, and K = 1. In the NK model, agents have a dichotomous choice, 0 or 1, for each decision variable. In our model, agents have 64 choices for each decision variable.

Table 1. Description of Simulation Parameters

Parameter	Range	Purpose
Search landscape	64 × 64	The agents' task environment. This is the total number of possible combinations in which the agents search for innovations. In each dimension (row, column), there are 64 possible alternatives the agent can choose.
Initial granularity (own dimension, other dimension)	1 × 1, 8 × 8, 2 × 1, 4 × 1, 8 × 1, 16 × 1, 24 × 1	The agents' information partitions at the start of search. The number of initial choices the agent sees in each domain can vary between 1 and 64 in each domain.
Exploration parameter in switching	τ: 0 – 0.1	The agents' propensity to engage in explorative activity. This governs the agent's move between choices he or she is aware of. The movement is governed by a Softmax algorithm. The higher the parameter (τ), the more uncorrelated the actual movement of the agent with payoff differences.
Communication frequency	0 – 0.5	Communication regulates the extent to which the agents' knowledge partitions are aligned. For each *dig* attempt, this is the probability with which the row (or column) agent receives new knowledge provided by the column (or row) agent. When the probability is zero, agents do not communicate. When the probability is 0.5, agents communicate approximately every other round.
Propensity to partition in other dimension	0 – 0.5	The agent's ability to bring forth new information in the complementary dimension. This is the probability with which a row (or column) agent's *dig* attempt results in a partition in the search space in the column (or row) dimension.

example, in which (both) the agents see a 3 × 2 choice set instead of reality (the full 64 × 64 matrix). The sharpness of the agents' initial vision (knowledge)—the number of choices in each dimension that an agent can see at the beginning of the game—is specified as a parameter in the model and may vary from 1 (most blurred) to 64 (sharpest) along each dimension. The agents' limited vision of the choice set also limits their understanding of the performance consequences of the choice set. The payoff the agents associate with each perceived cell in the matrix is the average of payoffs for the "real" combinations that are latent in that cell. For example, as shown in figure 4, for the other-curved combination, they see the average for all other materials and all curved shapes.[3]

Agents' initial mental models get more refined with time as their knowledge partitions become more fine-grained. As the choice set becomes more refined, the payoff associated with each element in the choice set also becomes more accurate.[4] Refining mental models involves two actions: first choosing the portion of the landscape to further explore (step 1) and then actually exploring (gaining a sharper vision of) that region (steps 2 and 3). The *switch* operation

[3] By assumption, agents have correct expectations about the attractiveness of each choice they see. This treatment is similar to the payoff matrix seen by agents in Gavetti and Levinthal (2000). In addition, agents have commensurate mental models. They agree that there are two dimensions, and they both see the same payoffs for identical subspaces, with no idiosyncratic distortions or filtering errors.

[4] We do not model noise in payoffs. When our agent achieves perfect vision of the choice set, the agent simultaneously achieves perfect information about the payoffs of each choice set.

Figure 2. Baseline search algorithm.

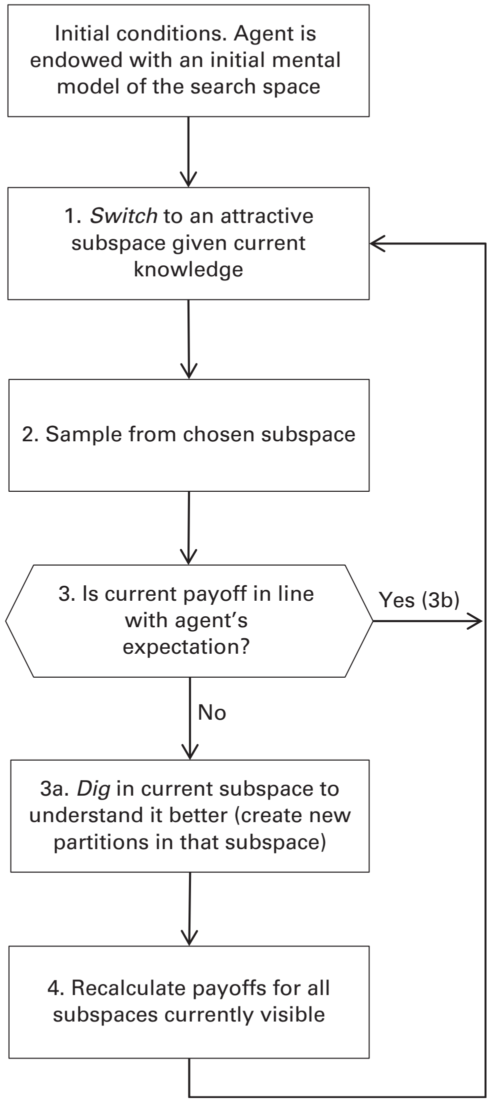

chooses the region for further exploration, and the *dig* operation refines the agents' current mental model in the specific location determined by *switch*.

Step 1: *Switch* to attractive subspace given current knowledge. Switching captures the logic of how agents change the focus of their attention from one region in the landscape to another. Initially, our agents are positioned at random in the landscape. They observe the payoff to their current choice and the payoff to all the other choices available to them based on their current mental model. Our agents are profit seeking and therefore switch to the most promising alternative they currently perceive. This is similar to Simon's (1962) conception of choosing between branches of the search tree for further exploration depending on the agents' expectations about which

Figure 3. Perfect vs. imperfect initial vision of task environment (one-quarter of actual size).

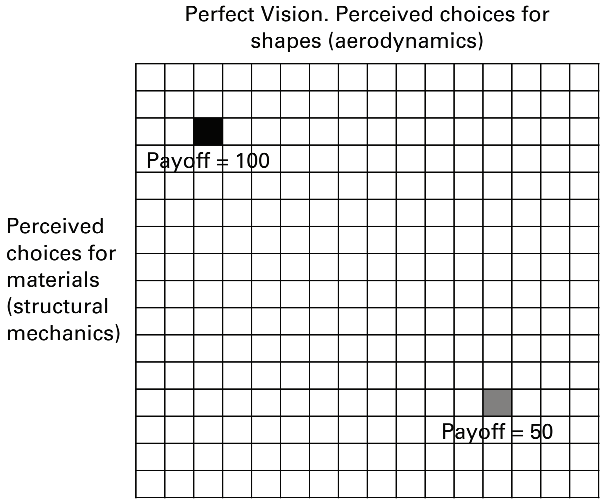

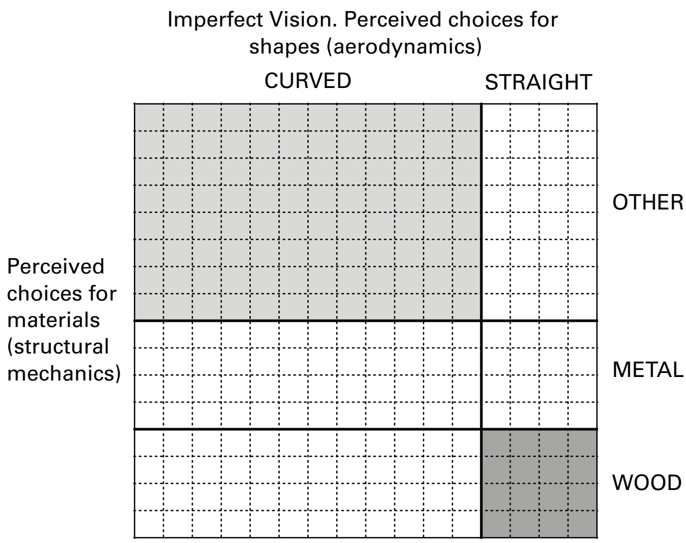

In perfect vision, all the 16 choices in the row and column dimension and the 256 (16x16) associated payoffs are visible to the agents. In imperfect vision, the agents only see a 3x2 matrix instead of the 16x16 matrix.

branch appears most attractive. For instance, in figure 4, the agent currently positioned in other–curved may instead choose to focus on the option wood–straight for further investigation based on the perceived payoffs. Note that the agents' perceptions of attractiveness depend on their current (imperfect) mental models and, as shown in figure 4, they may switch away from a region that contains the global peak.

Switch is accomplished as follows. Assuming that both the row and the column agent start with identical mental models, as shown in figure 4, the row agent chooses wood for material and the column agent chooses straight for shape

Figure 4. The agents' vision and switching operation to a more promising alternative (one-quarter of actual size).

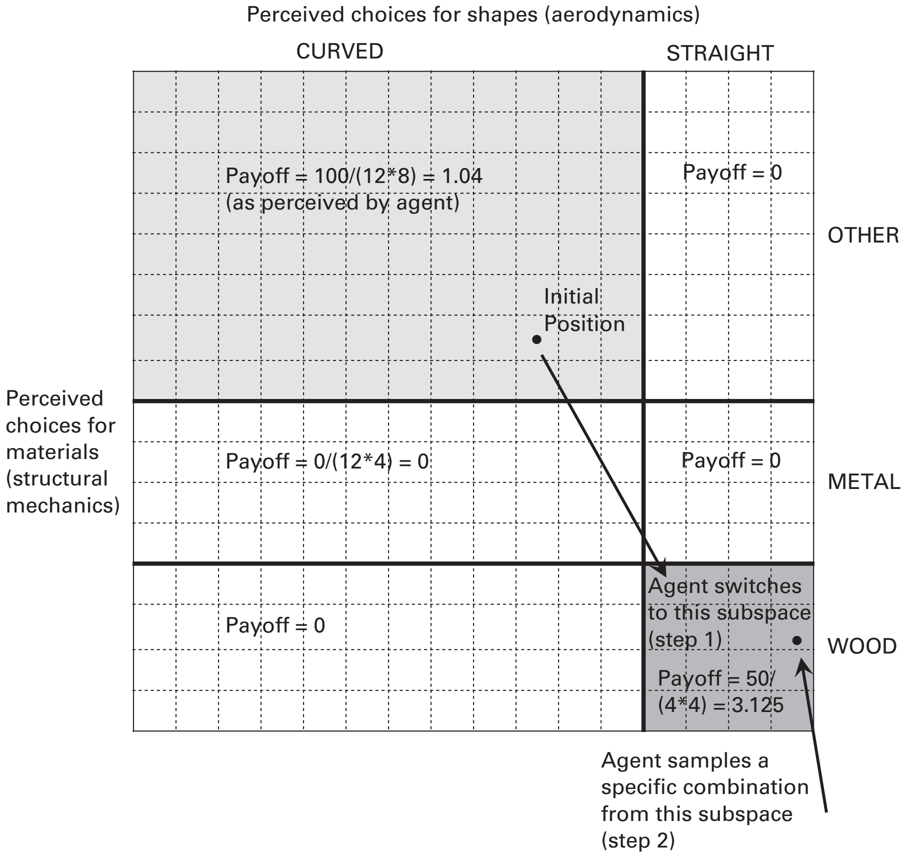

because each independently believes that this is the best subspace. When each of the agents makes their choice, they jointly switch to the wood–straight subspace.

In the baseline model, as illustrated in figure 4, the agents switch to the subspace with the highest perceived payoff, conditional on their mental model. In other specifications, we relax this assumption by allowing the agents to explore—sometimes they investigate subspaces that are not the most attractive as they currently see them. The higher the exploration parameter, the more the agents choose to investigate spaces at random without regard to their immediate attractiveness.[5] This is implemented using a Softmax algorithm.[6] The temperature τ in the Softmax algorithm is the exploration parameter—it determines the probability with which agents choose an alternative that does not have the maximum payoff as they currently perceive.

[5] This implements the typical strategy for modeling exploration in individual search models. We systematically vary this exploration parameter to understand the effect of individual exploration on joint search outcomes.

[6] The Online Appendix (http://asq.sagepub.com/supplemental) provides further mathematical details.

Step 2: Sample from chosen subspace. When both agents switch to their chosen subspace, they sample a combination from within that subspace. Because each alternative the agent is aware of (e.g., wood) contains multiple latent coordinates (e.g., oak, pine, cherry, etc.), sampling is achieved by placing the agent at random in one of these latent coordinates (e.g., oak for the row agent and a specific shape for the column agent; see figure 4). Because agents do not have any knowledge of the specific coordinates that make up a subspace, they have no control of their actual location within the chosen (coarse-grained) search space.

Step 3: Is current payoff in line with agents' expectations? With sampling, the agents become aware of the payoff to their joint solution. Each combination within a given coarse partition maps onto a particular payoff. Unlike game-theoretic models, our agents have less-than-perfect knowledge of any given subspace, and the payoffs they expect may be different from the payoff they receive from the particular combination they sample. This is because the expectation is the average of the payoffs of all the latent choices within that subspace. For example, in figure 4, the agent expects a payoff of 3.125, but the specific payoff actually received is zero.

Step 3a: If payoff is not in line with expectation, *dig* in current subspace to understand it better. The agents realize that any significant mismatch between expected and received payoff is a consequence of their imperfect knowledge of the subspace, which they then try to improve. We refer to an agent's propensity to gain further fine-grained partitions (sharper vision) in the chosen subspace as *dig*.[7] *Dig* implies that the agent expends effort in uncovering new knowledge such as by thinking about the problem, reading about it, or talking to others. Increased partitions allow the agent to distinguish between more nuanced categories. The idea is similar to Simon's (1962) conception of choice set expansion or refinement of the search tree.

In our model, when an agent decides to *dig*, a new knowledge partition occurs in the agent's mental model. In the windmill example, if the row agent decides to invest in understanding the subspace ''Wood'' more minutely, the subspace splits into hardwood and softwood (see figure 5), i.e., a new knowledge partition occurs. The increased partitions imply that the agent's mental model of the subspace is now more fine-grained in the row dimension. In figure 5, the agent now perceives four subspaces (hardwood–straight, hardwood–curved, softwood–straight, softwood–curved and the four associated payoffs) when earlier only two were perceived (wood–straight, wood–curved).

Similarly, when the column agent *digs*, the column dimension splits into two. Note that the new partition uncovered by the row agent is not visible to the column agent and vice versa. In figure 4, both the row and column agent perceive a 3×2 matrix. After *dig*, the row agent perceives a 4×2 matrix (as in figure 5), and the column agent would perceive a 3×3 matrix. Over time,

[7] In the baseline model, we implemented a surprise-driven search function (Cohen and Axelrod, 1984). In robustness checks, we implemented a version of the model in which the agent digs only when the payoff is less than the aspiration level (March, 1988). Our results are qualitatively unchanged.

Figure 5. ***Dig*** **operation: New information revealed to row agent.**

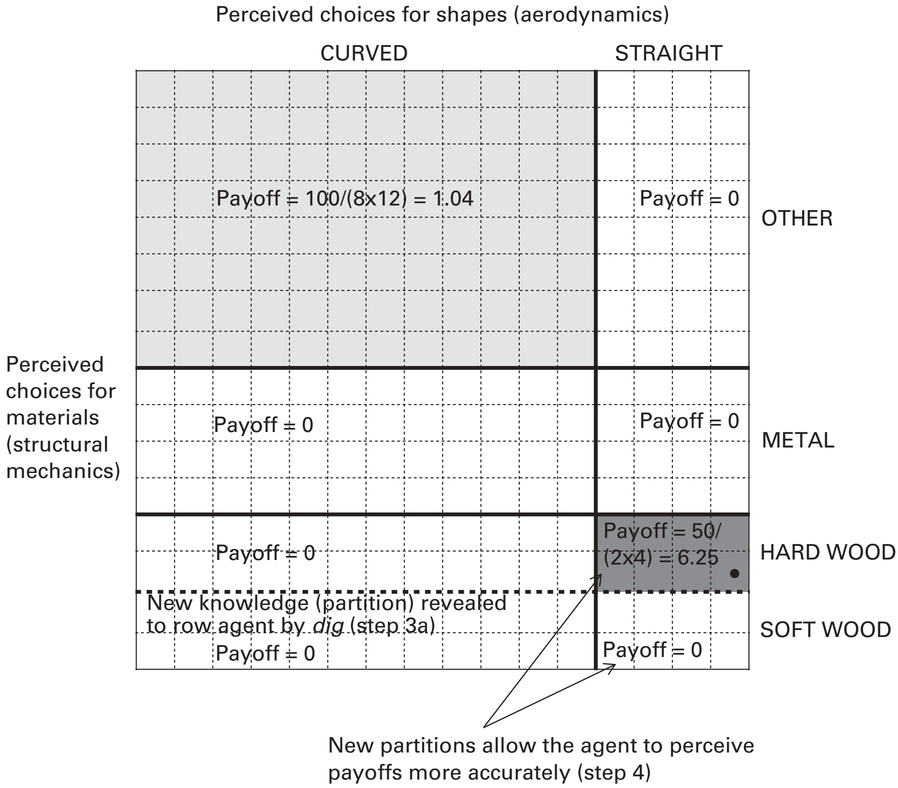

as the agents become aware of more partitions in the search space, their mental models increasingly diverge unless the agents take specific steps to align them.

To preserve bounded rationality and the logic of discovery in our model, we have imposed the following restrictions on the way digging leads to the refinement of mental maps. In the baseline model, each *dig* operation splits a subspace in two along the agent's specialist dimension (row or column). The exact point at which the split occurs in the subspace is chosen at random because the agent has no prior access to the latent choices within that subspace. Also, agents do not know when they have reached maximum granularity of vision in that subspace. When this is achieved, the agent may *dig* but does not become aware of any new partitions.

Step 3b: If payoff is in line with expectation, move back to step 1. If the actual payoff meets expectations, the agent does not expend effort in further partitioning the subspace but simply searches again, moving back to step 1. In the baseline model, under this condition, the agent samples again in the current subspace. Resampling is accomplished by placing the agent in a random combination within the current subspace (as explained in step 2 above).[8] In the baseline model, the agents effectively stop digging when their expected payoff is equal to what they actually receive, which happens only if they have identified the precise subspace that contains the peak. This implies that the agents have achieved perfect granularity in that subspace, but it is unlikely that they have maximal granularity in any other region of the landscape. An alternative assumption to the baseline is that if the received payoff exceeds the expected payoff, the agent decides not to search any further. We examined this alternative and found that our results are robust to this assumption.

Step 4: Recalculate payoffs to all subspaces currently visible. The agent recalculates expected payoffs for all known choices. If the *dig* operation in step 3a results in new partitions, the agent now has more choices available, and the payoffs the agent imputes to these choices have also grown more accurate, as shown in figure 5. At this point, the agent retraces the sequence of steps from step 1. The simulation ends after 500 discrete time steps. The value of 500 time steps was chosen because by then all simulations had approached a steady state.[9]

Organizing Joint Search

In unitary search, one agent searches in both dimensions and therefore automatically has fully aligned mental maps, as well as aligned actions. In joint search, however, there is a division of labor, and the *dig* operation leads to asymmetric mental models between these interdependent agents. To coordinate, the agents need to align their mental models. Different organization designs engender different patterns of interaction among the specialists, which determines the rate and the level of mental model alignment over time. This, in turn, affects agents' subsequent search locations (Denrell and March, 2001). We adopted three of the integration conditions proposed by Gavetti (2005), autonomous, top-down, and coordination, to understand their effects on coordinated exploration.[10] In all these conditions, the row agent determines the row position and the column agent the column position, and the agents switch to the portion of the matrix identified by their joint choice.

[8] Further sampling randomly repositions the agents within a subspace because agents who locate in a subspace have no knowledge about the underlying latent combinations. The agent is limited by its current granularity and consequently has no control over positioning within the chosen subspace. We can think of this as if the agents are performing experiments, but because experimental noise is not entirely eliminated (imperfect granularity), they get different results every time. Random repositioning allows the agent to continue search—if it happens to locate in a combination whose payoff is different enough from what is expected, the agent digs and is rewarded with more knowledge (more partitions). This procedure implements behavior that is consistent with the knowledge conditions we impute to agents. In principle, an agent cannot choose to stick to a particular point within a subspace unless the agent can actually see that point.

[9] We applied difference tests to the values of behavioral variables, as well as obtained payoffs. When these tests approach constant values for differences between successive time steps, the dynamics approach steady state.

[10] The fourth type of integration mechanism proposed by Gavetti (2005), "circulation of cognition," is not very meaningful in our setting because it involves complete transfer of knowledge from one agent to another.

Autonomous. In this case, the agents search in parallel but do not make any attempt to align their mental models.

Top-down. In this case, we consider the situation in which two agents search in parallel and senior management attempts to achieve coordination by imposing the same mental model on both the agents. This is accomplished such that both agents have identical (fully aligned) partitions of the search space initially (at t = 0). This initial alignment ensures that both agents begin by identifying the same region in the landscape as attractive and concentrate their search efforts in that region. After this initial alignment, search is identical to the autonomous regime. The initial partitioning is fairly limited and made at random, which allows us to understand the impact of initial shared knowledge on joint search.

Coordination. In this case, we consider the situation in which two agents search in parallel but make some attempt to align their mental maps of the search space. Coordination is an attempt by one agent to understand the world precisely as viewed by the other agent. In this condition, agents attempt to partition the search space in identical divisions by communicating their knowledge partitions to each other.

Coordination is modeled as follows: the row agent requests new knowledge about the column dimension from the column agent. With each request, the column agent provides the row agent with one new column partition that the row agent does not already know.[11] The more frequent these requests, the more aligned the knowledge partitions become. But there is an opportunity cost to communication. Because gaining new knowledge, by digging or by communicating, takes effort, each time an agent requests information from the other agent, the requesting agent forgoes the opportunity to further improve granularity in his or her own dimension in that time period. That is, an agent can improve the granularity in only one dimension at a time.[12] The frequency of communication is a parameter in the model and does not change with time. In this setup, we assume that communication effectively increases the alignment of mental maps. By assumption, there is no fundamental incongruence between agents' mental maps; differences between the two actors' knowledge partitions are the only source of misalignment. With infrequent communication, coordination approaches the autonomous case, whereas it approaches the opposite—fully shared knowledge structures—with increasing communication.[13]

[11] It works identically if the column agent initiates the request for a new partition from the row agent. Empirically communication between specialists is very difficult as knowledge does not easily transcend the different thought worlds that these specialists occupy (Dougherty, 1992; Iansiti, 1995). Therefore we have restricted communication to provide the agent with only one new knowledge partition. If agents communicated every partition they know about every time they communicated, they would no longer be specialists. If this were possible, mental models would be fully aligned in every time period, and the problem of mutual confusion would vanish, though the problem of joint myopia might still exist.

[12] In robustness checks, we relaxed this assumption and found that it does not qualitatively change our results.

[13] We model only perfect communication. The degree to which mental maps overlap is strictly governed by the frequency of communication. The autonomous case logically approximates imperfect communication.

FINDINGS

Tables 2–4 summarize the final payoffs received by the agents in the different treatment conditions in this simulation model while searching the two-peaks landscape as shown in figure 1. In all of the results, the exploration parameter $\tau = 0$. Initial granularity is set at 1×1 in all results except for the top-down search regime, where it is set at 8×8. Results are obtained at T = 500. Reported results are averages obtained from 300 runs for each condition, and reported differences are statistically significant at conventional levels. In general, the findings are inconsistent with the traditional hypotheses (H1a, H2a, H3a) that were derived from theory developed for search by a single agent.

Specialists with low (non-zero) levels of communication perform better than agents with high levels of communication. Table 2 shows that frequency of communication between specialists has a non-monotonic relationship with search outcomes. No communication between the agents (autonomous search) results in very poor outcomes (payoff of 0.70), while communicating only once in 20 rounds results in quite favorable outcomes (payoff of 0.97). Communicating very frequently (once in two rounds) results in comparatively poorer outcomes (payoff of 0.91). These results are consistent with H1b.

We find this pattern because no communication results in mutual confusion and lack of coordination whereas too much communication causes joint

Table 2. Impact of Level of Communication on Joint Search Performance by Specialists

Search regime	Propensity to partition in other dimension	Communication frequency	Final payoff (at t)
Coordination	0	0	0.70
	0	0.05	0.97
	0	0.10	0.93
	0	0.50	0.91

Table 3. Unitary vs. Joint Search Performance by Specialists vs. Agents with T-shaped Skills

Search regime	Propensity to partition in other dimension	Communication frequency	Final payoff (at t)
Unitary search	0	N/A	0.01
	0.50	N/A	0.90
Autonomous	0	0	0.70
	0.50	0	0.60

Table 4. Impact of Coordination by Hierarchy on Joint Search Performance by Specialists vs. Agents with T-shaped Skills

Search regime	Propensity to partition in other dimension	Communication frequency	Final payoff (at t)
Top-down	0	0	0.85
	0.50	0	0.85

myopia—a premature focus on local peaks. As argued before, we find that there is a tradeoff between search and coordination. Because specialists do not partition in the other dimension, if they do not communicate, then their choices are not coordinated. Communicating enables some alignment of mental maps and allows the agents to coordinate their choices. The more the mental maps are aligned, however, the more the agents influence each other in concentrating on a narrow portion of the landscape that both see as beneficial, but at the necessary cost of blurred vision of other regions in the landscape. Communication influences the region in which sharper vision is achieved because it regulates where knowledge partitions are increased. Therefore once a promising region is jointly identified, the agents concentrate on increasing their knowledge of that specific subregion and neglect exploration.

Premature focus is a powerful detriment to search because it prevents agents from exploring the high-value region, which is a natural attractor. We find that in the low-communication case, the agents quickly identify an interesting region from their own viewpoint and then slowly try to understand the space from the other's viewpoint. This enables them to explore effectively; we find that their probability of digging is above zero even after 250 time steps. The high-communication condition, in contrast, stifles exploration. We find that agents' probability of increasing granularity is about the same in both dimensions. Because of this, agents converge very quickly on a mutually desirable solution and cease exploration; the probability of digging falls to almost zero within 50 time steps. This leads to suboptimal search outcomes in this case.

T-shaped skills are useful for unitary search but problematic when agents are interdependent. Table 3 shows the average payoff of a unitary searcher who is a specialist vs. an agent with T-shaped skills. For specialists, the propensity to partition in the other dimension is 0; for agents with T-shaped skills it is 0.50. As expected, the agent with the T-shaped skills performs much better. Whereas the specialist agent achieves a payoff of 0.01, an agent with T-shaped skills achieves a payoff of 0.90; on average, 80 percent of the agents identify the high peak (= 1.00) and 20 percent of agents identify the low peak (= 0.50). Figure 6 shows the location of agents over time. The top graph shows that no agent with T-shaped skills has a final payoff of zero; they reliably identify one of the two valuable peaks.[14]

Table 3 shows the results for coordinated exploration achieved by teams a team of specialist agents vs. a team of agents with T-shaped skills. First, for specialists, table 3 shows, as expected, that when agents are autonomous, their payoff (0.70) is higher than the unitary specialist searcher but lower than the unitary searcher with T-shaped skills. The bottom of figure 6 shows that several agents achieve a payoff of zero because one agent focuses on the tall peak whereas the other agent focuses on the short peak; they confuse each other and jointly land on a solution that has zero value.

[14] An individual searcher need not be equally good in both dimensions. In robustness checks, we found that a searcher who partitions in the other dimension approximately only once every twenty rounds achieves approximately the same final payoff as one who partitions in the other dimension once every two rounds.

Figure 6. Average position of agents in the two-peaks landscape.

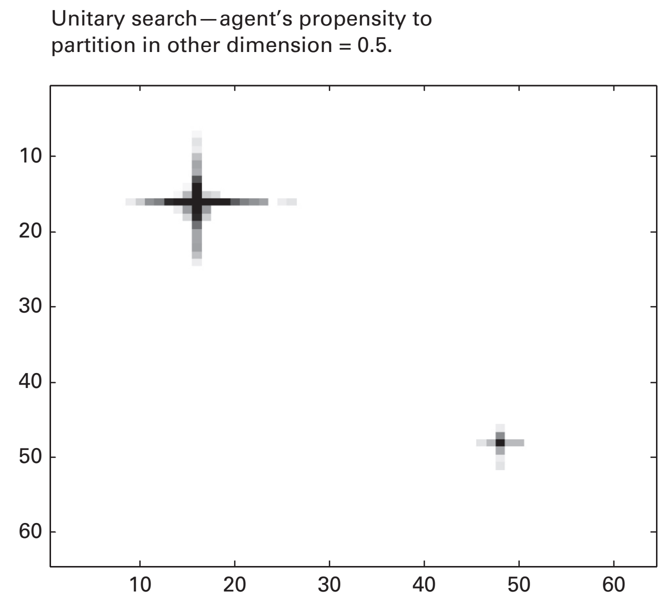

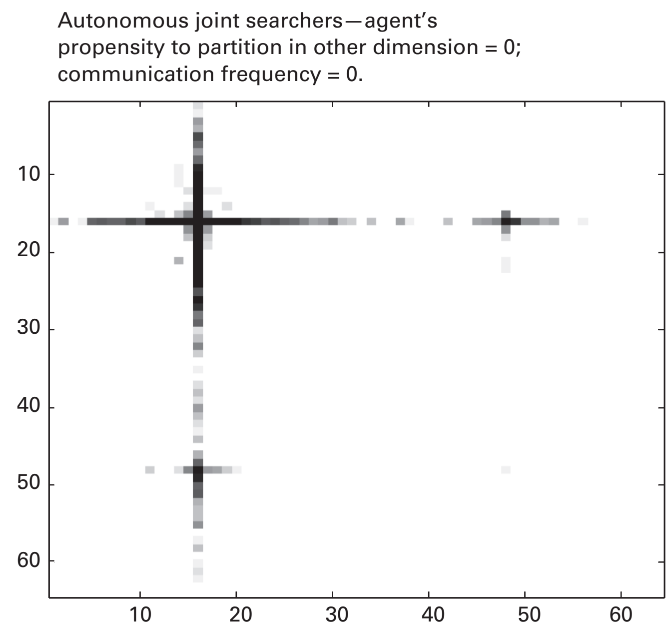

This evidence is consistent with our argument that agents in coupled learning problems are subject to mutual confusion, a condition that allows agents to persist with a flawed mental model of the search space. In this condition, at steady state, the *dig* probability of the average specialist agent is almost 0.9, that is, the agents are still attempting to find the valuable combination despite the negative feedback. Such confounding is impossible in search by a unitary agent. Because agents do not realize when maximal granularity is reached, they continue to experiment in the hope of finding better solutions.[15]

We argued that agents with T-shaped skills, though possessing superior knowledge of the landscape, are still subject to mutual confusion. Table 3 shows that their performance, with a payoff of 0.60, is *worse* than the autonomous specialist searchers. This finding demonstrates the challenge of epistemic interdependence in coordinated exploration. Because these two agents do not have aligned knowledge partitions, they are unable to make choices that are jointly valuable. This occurs even though both agents have almost as fine-grained partitions of the search space (average of 8.4×8.4) as the unitary searcher (8.6×8.4) has at steady state. In other words, the agents achieve very poor outcomes in joint search even though they are equipped with traits that give them the same potential as the individual searcher to acquire knowledge. This result is an interesting contrast to joint search models without epistemic interdependence. In prior work, joint searchers reached good solutions if they both had very good processing power (i.e., power to consider also the complementary dimension) as long as they did not prematurely weed out solutions (Rivkin and Siggelkow, 2003; Siggelkow and Rivkin, 2006).

Hierarchy as a coordination device. The problem of mutual confusion in coupled learning problems should lessen when agents are provided with a coordination device that gives them some ability to align their mental models of the search space. Communication is one means of aligning mental models. Another is hierarchy, which may impose identical mental models on the agents.

We implemented hierarchy following the top-down model proposed by Gavetti (2005) by providing two specialist agents with identical partitions in both dimensions at the beginning of search but having them search autonomously. Table 4 shows that when these agents are provided with an 8×8 partition space, they are able to achieve a payoff of 0.85, which is closer to the payoff achieved by the unitary searcher (0.90) and higher than the payoffs for two autonomous specialist searchers (0.70). The initial vision of the subspace in this case is sufficiently fine-grained for both the autonomous agents to target the same peak, and consequently the agents are less likely to mutually confuse each other. Similar results were achieved also for the agents with T-shaped skills.

This shows that integration mechanisms that rely on setting up an initial shared frame of reference, such as common culture or common processes, can be a powerful coordinating mechanism even under conditions of uncertainty. This finding is inconsistent with some of the earlier work that suggests that only feedback is useful for coordinating in situations of uncertainty (Galbraith, 1977; Tushman and Nadler, 1978) but consistent with more recent work that suggests that shared frames of reference help in coordinating by

[15] Situations in which agents continue to invest effort without realizing that a search region is useless are fairly common. Examples include the astronomers who spent their lives tweaking the geocentric model of the universe or the chemists who based their experiments on the phlogiston theory. Perhaps the most famous example is Pasteur, who tried to create a vaccine for rabies by using blood samples from infected animals, an ultimately fruitless quest because the rabies virus did not travel by blood like anthrax and the other pathogens he had worked with previously.

increasing the predictability of actions (Okhuysen and Bechky, 2009; Srikanth and Puranam, 2011).

Impact of exploration on unitary search vs. coordinated exploration. An agent explores when he or she chooses actions without regard to the immediate payoff. In our model, we investigated the impact of exploration by varying the exploration parameter, τ. As τ increased, the agent was more likely to switch between subspaces without regard to the immediate expected payoff from that subspace.

The first result we observed is that agents with $\tau = 0$ outperform agents with $\tau > 0$ across all conditions; the higher the exploration (τ), the worse the performance. Whereas the unitary searcher suffers relatively less than joint searchers for all values of τ tested, the performance of the autonomous specialists decreases the most. The results further suggest that the negative consequences of increasing mutual confusion dominate the slow-learning effect as agents' exploration increases (starting from low values of τ). Our results also support our intuition that increasing the alignment of mental models should help reduce mutual confusion. As agents increase exploration, the top-down condition performs better relative to the autonomous condition, and agents who communicate more frequently outperform those who communicate less often. These results are consistent with our argument in hypothesis 3b, at least for the two-spike landscape.

Robustness checks. As described in Online Appendix B, we performed a number of checks to assess the robustness of the findings to our assumptions and to clarify the underlying mechanisms. Specifically, we checked our results for (1) different types of landscapes, (2) different initial conditions, and (3) different search algorithms. These robustness checks in general strengthened our intuition about the above results.

DISCUSSION

Organizations exist to manage the tradeoff that arises with the division of labor: benefits from increasing specialization vs. losses arising from the need for coordination. Coordinated exploration—the condition in which specialist searchers need to coordinate their choices—is a significant problem for organizations but it is inadequately addressed by prior work. Prior theories on coordination ignore search, while prior work on organizational search has ignored the need for coordination; the bulk of this work characterizes the organization as a unitary actor (Cyert and March, 1963; Levinthal and March, 1981; Levinthal, 1997). Neither approach is helpful in understanding coordinated exploration.

Our contention is that predictions from prior theories are incorrect when applied to situations of coordinated exploration because the simplifying assumptions used abstract away from the fundamental problem posed by the coupling of uncertainty and epistemic interdependence. Theories of coordination assume that the specialist agents have complete knowledge in their own search domains and recommend strategies that swiftly increase common ground to achieve high performance (Galbraith, 1977; Tushman and Nadler, 1978). They ignore the effect of increasing common ground on subsequent search, i.e., increasing joint myopia that actually decreases performance in coordinated exploration. Similarly, theories of organizational search suggest that slow learning, which promotes moderate exploration, is important for achieving good search outcomes (Denrell and March, 2001; Siggelkow and Levinthal, 2003; Ethiraj and Levinthal, 2004; Knudsen and Levinthal, 2007; Fang, Lee, and Schilling, 2010). These theories ignore the effect of individual exploration on coordination, i.e., mutual confusion.

Managing the Scylla and Charybdis of mutual confusion vs. joint myopia distinguishes problems of coordinated exploration from problems of unitary search and from pure coordination problems. Mutual confusion arises because feedback to interdependent searchers confounds the consequences of their actions with the actions of others, preventing them from forming accurate mental models of the search space. Aligning mental models has the consequence of reducing mutual confusion but at the cost of increasing joint myopia, which is agents' tendency to concentrate on that portion of the landscape that is perceived as jointly attractive while ignoring the need to broadly explore the search space.

Limitations

Although we have developed an informative model of coordinated search, this work is subject to a number of limitations. First, the landscape of innovation is exogenous to the model. Because the optimal organization of joint search is contingent on the type of innovation landscape, how do managers know what type of landscape they are searching in? We do not address this, but prior work on belief formation may be helpful here, such as the work on analogical reasoning (Gavetti and Rivkin, 2007). Second, we have assumed that agents can switch seamlessly anywhere in the landscape, which may not be feasible because of limited rationality. A related problem is that knowledge of the agent sequentially increases—there is no forgetting in this model. Third, we have not systematically modeled the performance of agents with asymmetric abilities. We have also not explored sequencing, such as first searching with generalists and then with specialists, or coordinating by one type first and then by the other type. This is interesting future work. Fourth, we have not modeled hierarchy, which may be an important mechanism to align mental models or direct search without such alignment. Our top-down form of coordination is related but can be extended. As a final limitation, our analysis is focused on joint search involving two specialists. There is no reason to believe that including more dimensions would alter the results, though the analysis would be more complicated. Future research could examine whether the dynamics we identify remain unaltered for higher-dimensional problems.

Contributions and Future Research

Despite its limitations, our model makes some important contributions. Significant organizational phenomena call for the need to jointly consider the search for solutions by individuals with diverse knowledge (i.e., a specialist activity) and achieving coordination between these interdependent searchers. Our novel contribution is to extend prior theory on organizational search to include problems of coordinated exploration and understand the mechanisms by which problems of coordinated exploration differ from unitary search.

This is one of the first efforts to model search by considering both the cognition and organization of multiple agents and their joint impact on search outcomes. It is not possible to understand the role of epistemic interdependence on search without modeling agents' cognition. We show that joint search is not scaling up of individual search but is qualitatively different. By employing a richer modeling strategy, we were able to refine predictions from previous theory and illustrate a novel mechanism—the tradeoff between mutual confusion and joint myopia—that makes joint search problems very different from individual search. Moreover, our mechanism is robust to a number of checks, including the nature of the landscape (Rivkin and Siggelkow, 2007), differences in the initial cognition of the agents (Gavetti, 2005), and the specific assumptions of the search algorithm—surprise-driven vs. payoff-driven (Simon, 1962; Cohen and Axelrod, 1984; March, 1988). As a methodological contribution, we also provide an alternative modeling platform in which it is possible to understand the consequences of different kinds of assumptions about common knowledge. Finally, our work has some very interesting implications for game theory. Prior games have considered either perfect information or imperfect information to the extent that payoffs to action choices are noisy. Our model can be interpreted as a game whose structure unfolds with time. This feature of games is rather unexplored, though search is fundamental to organizations.

The novel mechanisms explicated in this study throw some light on resolving long-standing empirical contradictions and offer some novel predictions. Some empirical work in new product development suggests that high levels of communication improve innovativeness in performance, whereas other studies suggest the opposite (Tyre and Hauptman, 1992; Montoya-Weiss and Calantone, 1994). This contingency effect is not well understood theoretically. We offer novel predictions by suggesting a contingency when this relationship between communication volume and innovation performance is true. Our simulation results suggest that the effect of communication depends on the nature of the task environment. The more the landscape emphasizes search over coordination, the more detrimental are the effects of too much communication. In contrast, the more the landscape emphasizes coordination over search, the greater the need for communication. Future empirical research should take into account the nature of the problem space when determining the impact of organizational mechanisms that promote high levels of interaction, such as cross-functional teams, on new product development performance.

Our simulation model also suggests that the impact of communication volume on innovation performance depends on the initial knowledge held by the agents. The more knowledgeable the specialist agents are in their own domains, the less need for search and the more the joint search problem resembles a coordination problem. Under these conditions, more communication should have a beneficial effect. This suggests that the more deep specialists communicate, the greater the likelihood of achieving a highly innovative outcome, whereas the reverse should hold true for ''shallow'' specialists. From an analysis of patent data, Fleming (2007) suggested precisely such a relationship. Future empirical work could test this relationship in other contexts.

Our results may also have interesting implications for the organization of innovation. Srikanth and Puranam (2014) argued that higher levels of common ground are found within firm boundaries than across them. Coupled with the findings from our study, this suggests that alliances or other market-based organizations may be a more effective way to organize innovations that require significant levels of search, whereas organization under hierarchy and tight communication may be more effective for innovations that require high levels of coordination. For example, Kotha and Srikanth (2013) argued in the context of Boeing's 787 program that greater coordination effort between suppliers would help balance the need for both innovation and coordination. This argument is consistent with our model, though it is contrary to popular opinion that perhaps Boeing should not have outsourced the critical design tasks.[16] Clarifying these relationships could be an interesting avenue for future work.

Despite a family resemblance, the mechanism we identify is distinct from slow learning in typical exploration–exploitation models involving unitary agents (Denrell and March, 2001; Knudsen and Levinthal, 2007; Fang and Levinthal, 2009). The slow-learning result is typically achieved by decreasing the sensitivity of agents' actions to their performance consequences. In our model, however, this strategy leads to poor outcomes because of mutual confusion. We also find that reducing communication improves performance by reducing joint myopia, though only in landscapes that require exploration. Unlike prior work (Lounamaa and March, 1987; Puranam and Swamy, 2011), our result is not achieved because agents' sensitivity to payoffs decreases but because search precedes coordination. Lazer and Friedman (2007), in a model of network search based on the NK framework, reached a conclusion similar to ours with respect to frequency of communication. In their model, actors would mimic other successful actors, and when there was no one to mimic, they would attempt to adapt their status quo configuration. Their core result was that systems with higher levels of connectivity and communication frequency would perform better in the short run, at the expense of long-run performance. While we reach a similar conclusion about communication frequency, the underlying mechanism is very different. In Lazer and Friedman's (2007) model, actors could learn from each other about what does and does not work, but it was a pure search model without the need to coordinate actions in the face of epistemic interdependence. In contrast, our model captures not just the evolution of mental models but also their convergence in a coordination process and the impact on subsequent exploration. Typical slow-learning models do not capture these dynamics.

The contrast with slow-learning models does raise the question of why the agents are unable to acquire more information and then use this information in a sensible or optimal way to guide both exploration and coordination. The answer to this question lies in the assumptions we make: (1) agents initially have few partitions, i.e., they initially have little (or no) understanding of the task environment; (2) agents have limited overlap in their partitions, i.e., they see the search space from different positions; (3) agents have no common knowledge, i.e., even when they have overlapping partitions, they do not know that they do; (4) agents act in parallel, i.e., there is no principal–agent relationship so that one agent can explicitly guide the other; and (5) agents do not know what the optimal payoff is.

[16] For example, see ''Why Boeing's 787 was a nightmare waiting to happen'' (The Guardian, 18 Feb. 2013, http://www.theguardian.com/business/2013/jan/18/boeing-787-dreamliner-grounded) and ''Nightmareliner'' (The Economist, 3 Sept. 2011, http://www.economist.com/node/21528275).

To use information to optimize both exploration and coordination, one or more of these assumptions must be lifted. For example, if the agents faced a task environment with known maximum payoffs, they could simply use this information to define a sensible stopping point (not necessarily the global maximum). As our robustness results show, if the agents have very high levels of initial knowledge about the search space, coordination concerns outweigh exploration concerns, and this tradeoff can be managed. But this ''explore, then coordinate'' approach can be fairly time consuming, potentially expensive (some experiments are costly to conduct), and perhaps inconclusive (what is a very high level of knowledge in the real world, and where did it come from?).

If agents had common knowledge about their knowledge partitions, they perhaps could keep taking samples from the wider space to see if some distant points were superior. Our results suggest that exploration with overlapping knowledge partitions, but without common knowledge, leads only to mutual confusion. But establishing such common knowledge in effect implies the lack of specialization. One way to prevent this may be if the agents were to establish a principal–agent relationship or rules for sequential search (e.g., see Selten and Warglien, 2007). We have not further examined this option because it would dramatically increase the configuration space of the model. This is an excellent avenue for future research.

A related question is whether hierarchy can simply solve these coordination problems. We investigated the top-down condition, in which a superior directs search by aligning mental models up front, which has only a limited impact on improving performance because the hierarchy needs to be informed in advance either of the location of the global peak or of the emergent knowledge (new partitions) of both the specialists to coordinate effectively. The first condition implies that search is largely unnecessary. The second puts extreme demands on the coordination capacity of the organization; it is perhaps easier to inform the other specialist directly than to inform the supervisor who then in turn directs the search efforts. Hierarchy cannot simply solve problems of coordinated exploration, for the simple reason that the specialist agents have much more immediate knowledge than the supervisor, and this knowledge is difficult to transmit. Hierarchy could potentially have a role in solving these problems either by sequencing actions or by appointing agents with different skills at different stages of the problem. Examining these options is a good avenue for future research. We hope that our model and its conclusions will encourage others to examine the role of hierarchy and further broaden our understanding of the relationship between organization and joint search.

Acknowledgments

We gratefully acknowledge many valuable comments from Associate Editor Philip Anderson, three anonymous *ASQ* referees, and on previous versions from Oliver Baumann, Giovanni Gavetti, Reddi Kotha, Tobias Kretschmer, Daniel Levinthal, James March, Phanish Puranam, Bart Vanneste, Sidney Winter, Jason Woodard, participants in the Theoretical Organizational Models society meetings, and seminar participants at the Institute for Strategy, Technology and Organization, Ludwig-Maximilians University, Munich, National University of Singapore and Singapore Management University. Kannan Srikanth gratefully acknowledges support from the Strategic Organization Design Unit, University of Southern Denmark. Thorbjørn Knudsen benefitted from a

Sapere Aude grant from the Danish Research Council for Independent Research (FSE). The authors contributed equally and are listed alphabetically.

REFERENCES

Brown, S. L., and K. M. Eisenhardt
1995 ''Product development: Past research, present findings, and future directions.'' Academy of Management Review, 20: 343–378.

Camerer, C.
2003 Behavioral Game Theory. Princeton, NJ: Princeton University Press.

Carlile, P. R.
2004 ''Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries.'' Organization Science, 15: 555–568.

Clark, K. B., and T. Fujimoto
1991 Product Development Performance: Strategy, Organization and Management in the World Auto Industry. Boston: Harvard Business School Press.

Cohen, M. D., and R. Axelrod
1984 ''Coping with complexity: The adaptive value of changing utility.'' American Economic Review, 74: 30–42.

Cronin, M. A., and L. R. Weingart
2007 ''Representational gaps, information processing, and conflict in functionally diverse teams.'' Academy of Management Review, 32: 761–773.

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

Davis, J. P., and K. M. Eisenhardt
2011 ''Rotating leadership and collaborative innovation: Recombination processes in symbiotic relationships.'' Administrative Science Quarterly, 56: 159–201.

Denrell, J., and J. G. March
2001 ''Adaptation as information restriction: The hot stove effect.'' Organization Science, 12: 523–538.

Dougherty, D.
1992 ''Interpretive barriers to successful product innovation in large firms.'' Organization Science, 3: 179–202.

Dougherty, D.
2001 ''Re-imagining the differentiation and integration of work for sustained product innovation.'' Organization Science, 12: 612–631.

Ethiraj, S. K., and D. Levinthal
2004 ''Bounded rationality and the search for organizational architecture: An evolutionary perspective on the design of organizations and their evolvability.'' Administrative Science Quarterly, 49: 404–437.

Fang, C., J. Lee, and M. A. Schilling
2010 ''Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning.'' Organization Science, 21: 625–642.

Fang, C., and D. Levinthal
2009 ''Near-term liability of exploitation: Exploration and exploitation in multistage problems.'' Organization Science, 20: 538–551.

Fleming, L.
2007 ''Breakthroughs and the 'long tail' of innovation.'' MIT Sloan Management Review, 49 (1): 69–74.

Galbraith, J. R.
1977 Organization Design. Reading, MA: Addison-Wesley.

Gavetti, G.
2005 ''Cognition and hierarchy: Rethinking the microfoundations of capabilities' development.'' Organization Science, 16: 599–617.

Gavetti, G., and D. A. Levinthal
2000 ''Looking forward and looking backward: Cognitive and experiential search.'' Administrative Science Quarterly, 45: 113–137.
Gavetti, G., and J. W. Rivkin
2007 ''On the origin of strategy: Action and cognition over time.'' Organization Science, 18: 420–439.
Gittins, J.
1989 Multi-armed Bandit Allocation Indices. New York: Wiley.
Hauptman, O., and K. K. Hirji
1996 ''The influence of process concurrency on project outcomes in product development: An empirical study of cross-functional teams.'' IEEE Transactions on Engineering Management, 43: 153–164.
Heath, C., and N. Staudenmayer
2000 ''Coordination neglect: How lay theories of organizing complicate coordination in organizations.'' In R. I. Sutton and B. M. Staw (eds.), Research in Organizational Behavior, 22: 153–193. Greenwich, CT: JAI Press.
Hoopes, D. G., and S. Postrel
1999 ''Shared knowledge, 'glitches,' and product development performance.'' Strategic Management Journal, 20: 837–865.
Iansiti, M.
1993 ''Real world R&D: Jumping the product generation gap.'' Harvard Business Review, 71 (3): 138–147.
Iansiti, M.
1995 ''Technology integration: Managing technological evolution in a complex environment.'' Research Policy, 24: 521–542.
Knudsen, T., and D. A. Levinthal
2007 ''Two faces of search: Alternative generation and alternative evaluation.'' Organization Science, 18: 39–54.
Kotha, S., and K. Srikanth
2013 ''Managing a global partnership model: Lessons from the Boeing 787 'Dreamliner' Program.'' Global Strategy Journal, 3 (1): 41–66.
Krishnan, V., and K. T. Ulrich
2001 ''Product development decisions: A review of the literature.'' Management Science, 47: 1–21.
Lawrence, P. R., and J. W. Lorsch
1967 Organization and Environment: Managing Differentiation and Integration. Boston: Harvard University Press.
Lazer D., and A. Friedman
2007 ''The network structure of exploration and exploitation.'' Administrative Science Quarterly, 52: 667–694.
Lee H., and B. Choi
2003 ''Knowledge management enablers, processes, and organizational performance: An integrative view and empirical examination.'' Journal of Management Information Systems, 20: 179–228.
Leonard, D.
1995 Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation. Boston: Harvard Business School Press.
Levinthal, D. A.
1997 ''Adaptation on rugged landscapes.'' Management Science, 43: 934–950.
Levinthal, D. A., and J. G. March
1981 ''A model of adaptive organizational search.'' Journal of Economic Behavior and Organization, 2: 307–333.

Lounamaa, P. H., and J. G. March
1987 ''Adaptive coordination of a learning team.'' Management Science, 33: 107–123.
Madhavan, R., and R. Grover
1998 ''From embedded knowledge to embodied knowledge: New product development as knowledge management.'' Journal of Marketing, 62: 1–12.
March, J. G.
1988 ''Variable risk preferences and adaptive aspirations.'' Journal of Economic Behavior and Organization, 9: 5–24.
March, J. G.
1991 ''Exploration and exploitation in organizational learning.'' Organization Science, 2: 71–87.
March, J. G.
2006 ''Rationality, foolishness and adaptive intelligence.'' Strategic Management Journal, 27: 201–214.
March, J. G., and H. A. Simon
1958 Organizations. New York: Wiley.
Montoya-Weiss, M. M., and R. Calantone
1994 ''Determinants of new product performance: A review and meta-analysis.'' Journal of Product Innovation Management, 11: 397–417.
Nelson, R. R., and S. G. Winter
1982 An Evolutionary Theory of Economic Change. Cambridge, MA: Harvard University Press.
Okhuysen, G. A., and B. A. Bechky
2009 ''Coordination in organizations: An integrative perspective.'' Academy of Management Annals, 3: 463–502.
Puranam, P., M. Raveendran, and T. Knudsen
2012 ''Organization design: The epistemic interdependence perspective.'' Academy of Management Review, 37: 419–440.
Puranam, P., and M. Swamy
2011 ''Expeditions without maps: Why faulty initial representations may be useful in joint discovery problems.'' Working paper, London Business School; http://papers.ssrn.com/abstract=1153142.
Rivkin, J. W., and N. Siggelkow
2003 ''Balancing search and stability: Interdependencies among elements of organizational design.'' Management Science, 49: 290–311.
Rivkin, J. W., and N. Siggelkow
2007 ''Patterned interactions in complex systems: Implications for exploration.'' Management Science, 53: 1068–1085.
Samuelson, L.
2004 ''Modeling knowledge in economic analysis.'' Journal of Economic Literature, 42: 367–403.
Schelling, T. C.
1960 The Strategy of Conflict. Cambridge, MA: Harvard University Press.
Selten, R., and M. Warglien
2007 ''The emergence of simple languages in an experimental coordination game.'' Proceedings of the National Academy of Sciences, 104 (18): 7361–7366.
Siggelkow, N., and D. A. Levinthal
2003 ''Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation.'' Organization Science, 14: 650–669.
Siggelkow, N., and J. W. Rivkin
2006 ''When exploration backfires: Unintended consequences of multilevel organizational search.'' Academy of Management Journal, 49: 779–795.

Simon, H. A.
1947 Administrative Behavior. New York: MacMillan.
Simon, H. A.
1962 The Sciences of the Artificial. Cambridge, MA: MIT Press.
Sine, W. D., H. Mitsuhashi, and D. A. Kirsch
2006 ''Revisiting Burns and Stalker: Formal structure and new venture performance in emerging economic sectors.'' Academy of Management Journal, 49: 121–132.
Song, M., and J. Xie
2000 ''Does innovativeness moderate the relationship between cross-functional integration and product performance?'' Journal of International Marketing, 8: 61–89.
Song, X. M., and M. M. Montoya-Weiss
1998 ''Critical development activities for really new versus incremental products.'' Journal of Product Innovation Management, 15: 124–135.
Song, X. M., R. J. Thieme, and J. Xie
1998 ''The impact of cross-functional joint involvement across product development stages: An exploratory study.'' Journal of Product Innovation Management, 15: 289–303.
Srikanth, K., and P. Puranam
2011 ''Integrating distributed work: Comparing task design, communication, and tacit coordination mechanisms.'' Strategic Management Journal, 32: 849–875.
Srikanth, K., and P. Puranam
2014 ''The firm as a coordination system: Evidence from software services offshoring.'' Organization Science, http://dx.doi.org/10.1287/orsc.2013.0886.
Sutton, R. S., and A. G. Barto
1998 Reinforcement Learning. Cambridge, MA: MIT Press.
Thompson, J. D.
1967 Organizations in Action. New York: McGraw-Hill.
Tsai, M. T., and Y. C. Huang
2008 ''Exploratory learning and new product performance: The moderating role of cognitive skills and environmental uncertainty.'' Journal of High Technology Management, 19 (2): 83–93.
Tushman, M. L., and R. Katz
1980 ''External communication and project performance: An investigation into the role of gatekeepers.'' Management Science, 26: 1071–1085.
Tushman, M. L., and D. A. Nadler
1978 ''Information processing as an integrating concept in organizational design.'' Academy of Management Review, 3: 613–624.
Tushman, M. L., and T. J. Scanlan
1981 ''Boundary spanning individuals: Their role in information transfer and their antecedents.'' Academy of Management Journal, 24: 289–305.
Tyre, M. J., and O. Hauptman
1992 ''Effectiveness of organizational responses to technological change in the production process.'' Organization Science, 3: 301–320.
Van Huyck, J. B., J. P. Cook, and R. C. Battalio
1997 ''Adaptive behavior and coordination failure.'' Journal of Economic Behavior and Organization, 32: 483–503.

Authors' Biographies

Thorbjørn Knudsen is a professor of organization design at the University of Southern Denmark (SDU) where he is heading the Strategic Organization Design Unit and is co-director of the Danish Institute for Advanced Study (DIAS), Campusvej 55, DK 5230 Odense Denmark (e-mail: tok@sam.sdu.dk). His research interests center around evolutionary and adaptive processes in organizations and the way organization design can

shape these processes. He received his Ph.D. in business economics from the University of Southern Denmark.

Kannan Srikanth is an assistant professor of strategy in the Indian School of Business, Gachibowli, Hyderabad 500032, India (e-mail: kannan_srikanth@isb.edu). His research focuses on learning and coordination aspects of organizational adaptation and innovation. He received his Ph.D. in strategic and international management from the London Business School.