9-2014

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

Thorbjørn KNUDSEN
University of Southern Denmark

Kannan SRIKANTH
Singapore Management University, KSRIKANTH@smu.edu.sg

Coordinated Exploration:

Organizing Search by Multiple Specialists to Overcome Mutual Confusion and Joint Myopia[1]

Thorbjørn Knudsen*
Strategic Organization Design Unit
University of Southern Denmark
DK 5230 Odense M
DENMARK

+45-6550-3148
tok@sam.sdu.dk

Kannan Srikanth
Strategy Area
Indian School of Business
Gachibowli, Hyderabad 500032
INDIA

+91-40-2318-7165
Kannan_srikanth@isb.edu

Revision Date: 25 Nov, 2013

Published in Administrative Science Quarterly, 2014 Sep, vol. 59, issue 3, pp. 409-441 .
http://doi.org/10.1177/0001839214538021

[1] We gratefully acknowledge many valuable comments from the associate editor Philip Anderson, three anonymous referees, and on previous versions from Oliver Baumann, Giovanni Gavetti, Reddi Kotha, Tobias Kretschmer, Dan Levinthal, Jim March, Phanish Puranam, Bart Vanneste, Sid Winter, Jason Woodard and participants in the 2nd TOM society meetings.
*Authors contributed equally and are listed alphabetically.

Coordinated Exploration:

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

Abstract

The coordination of specialists' search efforts is one of the principal purposes of organization. Integration mechanisms enable joint search by allowing interdependent others to form shared mental models of the joint task. Whereas prior theory has concentrated on how integration mechanisms impact coordination among multiple specialists, they have not explored how integration impacts search. We call problems that emphasize both search and coordination as "problems of coordinated exploration". Using a computational model, we find that coordinated exploration is not simple scaling up of individual search. Coordinated exploration is subject to two problems – *mutual-confusion* and *joint-myopia* – that arise only when epistemic interdependence is coupled with uncertainty. Agents' attempts to reduce mutual-confusion automatically increases joint-myopia and vice-versa. Organizing coordinated exploration requires that agents' mental model alignment balances the need for both coordination and search in order to avoid these two pathologies in joint search.

Key Words: Coordinated Exploration, Organization Design, Coordination, Joint Search, New Product Development.

INTRODUCTION

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 by balancing exploration and exploitation (Levinthal, 1997; Ethiraj and Levinthal, 2004; Siggelkow and Levinthal, 2003). The unitary actor assumption is a useful abstraction and baseline that has facilitated a thorough examination of problems related to balancing exploration and exploitation and identified organizational mechanisms that can facilitate such search.

However such models abstract from the fact that organizational search in many cases involves *coordinated exploration*, where search is undertaken by multiple actors who need to coordinate their efforts. Many problems solved by firms involve joint effort by multiple domain experts (specialists). Interdependence between specialists implies the need for coordinating their search efforts. Whereas each specialist is responsible for search in their own domain, their payoffs are dependent 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, neither know all the available choices in the relevant domains, nor their payoffs. For this reason, the choices of each can influence their joint outcomes in ways that neither anticipates nor fully understands. This gives rise to the problem of *coordinated exploration*. Models of search performed by unitary actors ignore coordination that is fundamental to a theory of search involving multiple actors, and therefore offers few insights into problems of coordinated exploration. In this paper we aim to understand how organizations solve such problems.

Problems of coordinated exploration are ubiquitous in organizations because 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. However, division of labor is accompanied by the need for integration. Typically, organizations rely on differentiation and then integration to straddle this trade-off (Lawrence and Lorsch, 1967). In these cases, experts or specialized organizational units search in individual domains, the results of which are then integrated into a joint solution with an observable payoff. 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 such that it is jointly attractive to all agents, though it may not be optimal for any one of them.

As our canonical example, we adopt New Product Development (NPD) within firms, which typically involve multiple specialists working together. To make the rather abstract idea of coordinated exploration more concrete, consider the problem of designing windmills. Rotor blade design is critical for the efficiency of a windmill. Blade design 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. In general, light curved blades have good aerodynamics, but poor strength. This innovative search problem is jointly solved by two specialists: 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 "who should be the customer," portioning a market into segments that respond in homogenous ways to changes in the marketing mix (product, price, promotion, and place). At the same time, product experts explore the domain "what can we offer customers," devising product or service offerings that deliver different cost/feature combinations. A value proposition lies at the intersection of these choices, and needs to be discovered jointly by the two departments.

Many organizational problems display the properties of coordinated exploration, including these: coordinating input from legal experts and the HR-department in location decisions, coordinating response by business units that operate in different markets, or managing inter-agency communication during disaster situations. In these problems, the essential feature is that their solutions involve more than one actor, each searching in a different aspect of the problem space, and their separate decisions are integrated to generate a joint payoff matrix that the actors do not know and cannot conceive in advance.

Prior work in organization theory has implicitly assumed that results from individual search models can be generalized to problems of coordinated exploration. However, when multiple agents are engaged in coordinated exploration, it is a case of *epistemic interdependence*, a situation in which one agent's optimal choices depend upon 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 theory to suggest that coordinated exploration is nothing but individual search coupled with high levels of communication between the agents (Tushman and Nadler, 1978). However, empirical work suggests that a high level of communication between specialists is not reliably associated with good outcomes (Brown and Eisenhardt, 1995; Sine et al, 2007; Montoya-Weiss and Calantone, 1994). Therefore, it is currently unclear how epistemic interdependence influences boundedly rational organizational search.

In this paper we investigate how to organize coordinated exploration using an agent based model. We find that conclusions from unitary searcher models do not fit coordinated exploration; coordinated exploration, depending on how it is organized, can outperform or underperform search by a single agent. Specifically, coordinated exploration is subject to two pathologies that are not present in unitary search.

Joint search involves a situation where feedback to one agent's actions is confounded by the actions of the other agent. Search therefore leads to increasing *mutual-confusion* where agents are unable to learn from feedback to correct their faulty mental models of the search space. Incorrect mental models held by one agent leads to mistakes, which in turn confuse the other agent (perturb their mental model) and so on. Sharing knowledge aligns mental models and counters mutual-confusion by inducing coordination around particular search regions. However, that very effort increases *joint-myopia* in search, a situation where agents mutually reinforce each other into an increasingly narrow portion of the search space prematurely. In the extreme, high levels of shared knowledge induces agents to abandon their distinct search approach in favor of a lower common denominator.

In other words, in coordinated exploration, *increasing coordination effort reduces mutual confusion, but simultaneously increases joint myopia; efforts to reduce joint myopia automatically increase mutual confusion.* 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 in the context of coordinated exploration.

PRIOR WORK ON ORGANIZING FOR COORDINATED EXPLORATION

How to effectively manage the specialization-coordination trade-off has engaged organization scholars for over fifty years, if not longer. 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. However, when the nature of interdependence is unknown or unstable, coordination is achieved by feedback. Thompson (1967) extended these insights and argued that under conditions of unknown interdependence, coordination is achieved by mutual adjustment. Coordinated exploration is only important 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 and seeks to understand how organizations can be designed to effectively operate 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; Tushman and Nadler, 1978). Therefore, highly interdependent work must be structured to maximize opportunities for information transfer (Lawrence and Lorsch, 1967; Tushman and Nadler, 1978).

In other words, it is well accepted that shared knowledge is necessary to coordinate epistemic interdependence (March and Simon, 1958; Puranam, et al, 2012). However, we also know that, in practice, it is very difficult to develop such shared knowledge among specialists (Lawrence and Lorsch, 1967; Cronin and Weingart, 2007). This is because the boundaries of specialization are also natural interpretive barriers that make the formation of shared understandings difficult (Lawrence and Lorsch, 1967; Dougherty, 1992; Heath and Staudenmayer, 2000). Organizations employ different kinds of “integration mechanisms” for the purpose of developing shared knowledge among specialists (Lawrence and Lorsch, 1967; Clark and Fujimoto, 1991; Iansiti, 1995; Hoopes and Postrel, 1999).

However, different integration devices likely generate shared knowledge in different ways. For example, frequent communication vs. infrequent communication leads to different patterns of (shared) knowledge over time. This means that the type of integration mechanism employed is likely to impact the outcomes from coordinated exploration in two ways: (1) emergent shared knowledge directs search. (2) The resulting sampling of the search space influences what 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 under-theorized.

Prior formal work on search has not addressed these issues, and therefore provides little guidance. As Knudsen and Levinthal (2007) observe, most models of organizational search are non-organizational – they assume a unitary actor. Prior models of multi-agent search ignore epistemic interdependence. Specifically, extant models of joint search attempt to understand how to balance exploration vs. exploitation and provide some subtle insights regarding how this tradeoff is achieved in organizations (see Rivkin and Siggelkow, 2003; Siggelkow and Levinthal, 2003; Siggelkow and Rivkin, 2006; Fang, Lee and Schilling, 2010), but have side-stepped issues of coordination. For example, these studies underscore the importance of “slow learning” in balancing exploration concerns with demands for exploitation. However, the concept of ‘shared knowledge’ that dominates the empirical literature is entirely absent from extant models of joint search.

Consider the case of mutual adjustment between two agents: will increase in knowledge transfer and shared knowledge between these agents increase or reduce their joint exploration? Will it improve or reduce their chances of identifying the global peak, i.e. the maximal point in the task environment? In contrast, will agents' efforts to explore via slow learning perturb their efforts to maintain shared knowledge? To answer these questions, we need a model that takes into account how increasing "*shared*" knowledge impacts agents' search behavior. Since the above-cited 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 the way two or more specialists best organize a process of joint search.

In this sense, prior work on joint search is fundamentally incomplete since 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 neglect search. Game-theoretic models typically assume that agents make a choice over a known state-space, and do not consider conditions where the agents' knowledge regarding the nature of the state-space is evolving. Typical game theoretic work does not model evolving game structures (Brandenburger, 2008; Heifetz, 2008), which is an important property in many real world problems. Models by Lounamaa and March (1987) and Puranam and Swamy (2012), though boundedly rational, since actors do not know the actions/payoffs available to the interdependent agent, also assume that actors know all actions available to them.

In other words we have very little knowledge about coordinated exploration, even though it appears to be a very important problem for organizations.

THEORY AND HYPOTHESES

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 sub-units 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 sub-units and high levels of integration among these units, giving rise to the problem of coordinated exploration.

The information processing theory of organizations suggests that highly interdependent work must be organized such that there is a high level of communication between the agents (Galbraith, 1977; Tushman and Nadler, 1978; Nadler and Tushman, 1998). High levels of communication increases the level of common ground – knowledge that is shared and known to be shared – among the agents, thereby promoting coordination (Simon, 1947; Schelling, 1960; Srikanth and Puranam, 2011). Even though this theory is about coordination and not search, it has been extensively used to make predictions in situations of coordinated exploration. For example, it is almost axiomatic in the new product development (NPD) literature that a higher level of information transfer between agents is associated with better performance (see reviews by Krishnan and Ulrich, 2001; Brown and Eisenhardt, 1995). 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*:** Specialist agents with higher levels of communication will be *more* likely to identify high value combinations than agents with lower levels of communication in problems of coordinated exploration.

The above prediction, however, is not uncontroversial. For example, empirical work in the NPD literature suggests that more 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 and Xie, 2000; Song, Thieme and Xie, 1998; Bettenhausen, 1991; Shenhar et al, 2002; Gomes et al, 2003). Montoya-Weiss and Calantone, (1994) observe that few empirical studies use technological innovativeness as a measure of NPD success; instead market share, financial success, or most frequently, speed of development are used as proxies indicative of NPD success. When *innovativeness* is the key criterion for defining success of NPD projects these studies suggest that facilitating very high levels of communication among the specialist agents may be associated with poor performance. The literature on boundary spanners reaches similar conclusions. They find that projects with boundary spanners tend to perform better than projects without these (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).

These empirical findings suggest that a high level of information transfer is unnecessary and perhaps even harmful in situations of coordinated exploration. However, the mechanisms that underlie these findings are unclear. Many large sample studies typically hypothesize that high levels of communication is associated with better performance, but do 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). For example, Davis and Eisenhardt (2011) explore innovations from high-technology alliances, and find that a consensual leadership style promotes significant and costly attempts at sharing information, which leads these firms to quickly adopting a 'lowest common denominator' approach. In other words, since significant effort is needed to transfer knowledge across specialists with incompatible mental models or 'thought worlds' (Heath and Staudenmayer, 2000; Dougherty, 1992; 2001), such efforts are likely prone to conflict and delays (Ancona and Caldwell, 1989). As a consequence, teams are more likely to pursue objectives that are minimally acceptable for all team members rather than explore broadly to achieve more rewarding outcomes. These observations lead to a prediction that is in contrast to the previous hypothesis:

> ***Hypothesis 1b***: Specialist agents with higher levels of communication will be *less* likely to identify high value combinations than agents with lower levels of communication in problems of coordinated exploration.

One approach to the differentiation-integration trade-off and the difficulty of communicating across specialist boundaries is the employment of 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 Barton, 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 context of our windmill example, if the structural mechanics engineers possess T-shaped skills, they are less likely to limit search for solutions to the strongest materials such as steel, since they recognize that these also tend to be heavy and therefore are unlikely to generate much power. The employment of 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 that:

> ***Hypothesis 2a***: Agents with T-shaped skills will be *more* likely to identify high value combinations in problems of coordinated exploration when compared to specialists.

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 are more likely to be successful than a team of specialists. First, successful coordination is not determined by the amount of total knowledge available to the agents, but by the level of common ground (Clark, 1996; Puranam et al, 2012). Though a group of agents with T-shaped skills may possess higher levels of knowledge than a group of specialists, such agents acting autonomously, without taking active steps to coordinate and share their knowledge, are still likely to fare poorly in problems of coordinated exploration.

One could argue that agents with T-shaped skills, because they share some understanding of the complementary domains are likely to interact with each other more fruitfully (Madhavan and Grover, 1998; Iansiti, 1993). However, according to hypothesis 1b, perhaps interaction between such agents may also lead to poor innovative outcomes. Recent empirical work bolsters this premise by suggesting that employment of personnel with T-shaped skills (as opposed to specialists) is not necessarily associated with new knowledge creation or effective exploration in a NPD context (Lee and Choi, 2003; Tsai and Huang, 2008). It is interesting to note that these studies hypothesize that employing personnel with T-shaped skills should have an impact on performance, but fail to find one empirically. Therefore, we hypothesize:

> ***Hypothesis 2b***: (teams of) Agents with T-shaped skills will *not* be more likely to identify high value combinations in problems of coordinated exploration when compared to (teams of) specialists.

The contrasting hypotheses argued above arise from two pathologies: the inability to coordinate on the one hand and the lack of adequate exploration on the other. Formal work that models firm adaptation as search over a rugged landscape suggests that local adaptation traps firms in local optima, and exploration is crucial for superior performance in such 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, on the other hand, also leads to poor outcomes, since 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 (Denrell and March, 2001; March, 1991). These "slow-learning" effects are considered to be particularly beneficial in complex environments that require broad exploration of the state-space (Knudsen and Levinthal, 2007; March, 2006). Based on these studies, one would predict that:

Hypothesis 3a: Agents who explore moderately are *more* likely to find high value combinations in problems of coordinated exploration than agents who do not explore.

It is however unclear whether the above prediction derived from models of unitary search is accurate in problems of coordinated exploration. This is because the epistemic interdependence condition that characterizes coordinated exploration leads to two pathologies in learning. 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 others (Lounamaa and March, 1987; Puranam and Swamy, 2011). This impedes adaptation by allowing agents to continue searching in an unprofitable region of the landscape, because they may attribute poor payoffs to actions of the other agent rather than to their own misperception of action-outcome linkages. In other words, feedback ambiguity promotes *mutual-confusion* and in effect misleads agents into maintaining a flawed mental model of the landscape.

These effects of epistemic interdependence can be countered if agents maintain fully aligned mental models of the search space 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 the reason why it is commonly thought that high levels of communication among interdependent agents can facilitate coordination. However, 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. In other words, as the agents receive more information regarding each other, they are more likely to choose actions that reliably take into account others' preferences, thereby limiting search to areas known to be mutually beneficial. This narrowing of search, important for coordinating, necessarily comes at the cost of 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, since it requires a coordinated shift among all participants (Camerer, 2003; Van Huyck, Battalio and Cook, 1997). In the face of bounded rationality, agents do not know if any other better equilibrium exists, and therefore cannot achieve a coordinated shift. Theories of unitary search argue that such myopia can always be overcome with deliberate exploration strategies, as suggested by H3a. However, in problems of coordinated exploration, as in all coupled learning problems, exploration has the consequence of increasing unintended interference in agent learning and therefore is unlikely to lead to superior search outcomes. Therefore we argue that:

> ***Hypothesis 3b***: Agents who explore moderately are *less* likely to find high value combinations in problems of coordinated exploration than agents who do not explore.

The general mechanism that underpins our theory is the tradeoff between the need to align 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 sub-optimal outcomes. The greater the alignment in mental models, the more the team chooses options that is likely to result in a positive payoff given their (accurate) understanding of what the other actors are likely to choose. However, this sensitivity to epistemic interdependence stifles exploration of the search space (whose payoff potential is unknown) by concentrating search effort in the sub-space that is of immediate mutual interest to all agents.

Therefore outcomes to coordinated exploration needs to balance mutual-confusion against joint-myopia, which is likely influenced by the relationships between the following elements: (1) the agent's initial knowledge; (2) agents' learning based on (2a) the agent's own search efforts and (2b) 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 between integration mechanisms, individual mental models, and the level of shared mental representations. As in much analysis of dynamical systems, a computational model is a suitable method to trace these feedback-driven interacting relationships.

A MODEL OF COORDINATED EXPLORATION

In order to understand both the coordination and search aspects of coordinated exploration, we need to model agent knowledge or cognition as an information structure that bears some (potentially crude) 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 (potentially unobservable) actions of the other agent. This property makes learning from experience much more difficult in joint search than for the single searcher. Therefore, to model coordinated exploration, we need an approach to represent the current state of the agent's information structure 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 in coordinated exploration. The study of knowledge has become quite central in economics because "questions of who knows what play a role in examining basic economic issues, such as when gains from trade exist and when these gains can be realized" (Samuelson, 2004, p. 367). Economists are typically interested in what people know, and what they know about what others know (mutual knowledge). We use this basic apparatus to model the 'uniqueness' and 'sharedness' in agents' knowledge structures to understand how search and coordination proceeds in parallel.

Until recently, the question of cognition was not considered in models of organizational search and learning. Instead of explicitly modeling cognition, the typical approach is to assume that searching agents learn fitness values in various regions of a (NK) landscape that represents the agents' task environment[2]. While this abstract characterization of cognition has proven useful for understanding fundamental properties of search in complex task environments, it is rather incomplete since it cannot readily account for epistemic interdependence.

To understand how two or more agents are actually able to coordinate exploration of some search space, it is critical to understand how their cognitive representations include information that can direct search conditional on the actions of the other agent. It is here that the knowledge as partition approach comes in handy. As explained in Samuelson (2004), economists model knowledge as a state-space, so what someone "knows" is some state of a partition of the space. The finer the partition of an

[2] Of course, one may argue that the agent's cognition, in this case, is represented by a simple mental map with a fixed structure comprising N policy attributes whose values are changed according to the information uncovered when the agent is searching the landscape.

agent's information structure, the greater her knowledge regarding 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, distinguish 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 then, proceeds by going through the current information partitions or, if necessary by further partitioning (or fine-graining) the information structure.

To make this approach concrete, consider again the windmill design problem from the introduction. Bounded rationality implies that *ex ante*, the two engineers do not know the full set of materials and shapes that they could recombine. For example, the initial knowledge of the structural materials engineer can be limited to three categories of candidate materials, wood, metal and other; i.e., her initial knowledge consists of only three partitions. Out of these, she may select a promising candidate, say wood, and investigate it further. For example, she may discover that there are two different types of wood, hard or soft. This represents an additional partition or fine-graining of her mental model of the search space in so far as it pertains to wood. Among hardwoods she may discover that the oak behaves differently from elm, which in turn is different from pine. This further fine-graining of her knowledge is represented as more partitions in her information structure. Note that in this example, the engineer's knowledge partitions are becoming more fine-grained in the sub-space pertaining to wood, whereas her knowledge regarding 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 in a different manner, depending on shapes, such as straight or curved. For example, 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, the two agents may come to different conclusions regarding which regions in the search space are attractive and therefore may make mutually inconsistent choices; i.e., each selects a solution that appears to be useful from their own point of view, but are in fact *jointly* useless. This is the challenge of epistemic interdependence that agents involved in coordinated exploration need to solve.

In other words, as agents increasingly fine-grain their partition structure, their mental models become increasingly incongruent, and therefore they need to expend more effort in aligning their mental models. Partition models elegantly capture this trade-off 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 our windmill example. The search landscape is a matrix where each combination of the two technologies defines a coordinate with an associated payoff (see 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 associated with a payoff. Figure 2 provides an overview of the baseline model and Table 1 provides an overview of all parameters. The parameters in Table 1 are explained in greater detail below. These were chosen after numerous robustness checks to fine-tune the model. Next, we briefly explain how the steps in Figure 2 are implemented.

INSERT FIGURE 1, FIGURE 2 AND TABLE 1 HERE

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 64x64 combinations is associated with a payoff.[3] 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 – since there are only two valuable peaks among the possible 4096 combinations, and coordination – since each peak acts as a Nash equilibrium in this game. Since the shape of the landscape can materially affect successful strategies for coordinated exploration, 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 64x64 choices available to them ex-ante, and as a consequence do not accurately know the associated payoffs.

[3] It may be helpful to draw a brief analogy of our model with the NK modeling structure. In our model, N=2, since agents are searching only in two decision parameters. 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.

At the beginning of the simulation, agents do not have fine-grained partitions of the search space, i.e. they only see a very limited number of choices for each dimension. Figure 3 provides an example in the context of the windmill example, where the agents see a 3x2 choice set instead of reality (the full 64x64 matrix). The sharpness of the agents' initial vision (knowledge) – i.e., 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 choices set. The payoff the agent associates with each perceived cell in the matrix is the average of payoffs for the "real" combinations that are latent in that cell. For example, in Figure 4, for the wood/straight combination – they see the average for all woods and all straight shapes.[4]

FIGURE 3 and FIGURE 4

Agents' initial mental models get more refined with time as they generate more fine-grained knowledge partitions. As the choice set expands, the payoff associated with each element in the choice set also becomes more accurate.[5] 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 chooses the region for further exploration and the DIG operation

[4] By assumption, agents have correct expectations regarding the attractiveness of each choice they see. This treatment is similar to the payoff matrix seen by agents in Gavetti and Levinthal (2000). This baseline assumption could be refined in future research by adding noise to expectations. In addition, agents have commensurate mental models. They agree that there are two dimensions, and they both see the same payoffs for identical sub-spaces, with no idiosyncratic distortions or filtering errors.

[5] We do not model noise in payoffs. When our agent achieves perfect vision of the choice set, the agent simultaneously achieves perfect information regarding the payoffs of each choice-set. Introducing noisy payoffs may be a very interesting extension to this model (see also previous footnote).

refines the agents' current mental model in the specific location determined by SWITCH. These are explained in detail below.

Step 1: SWITCH to attractive sub-space 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 as currently perceived. This is similar to Simon's (1962) conception of choosing between branches of the search tree for further exploration, depending on the agents' expectations regarding which 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 agent's perception of attractiveness is dependent on their current (imperfect) mental model, 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, since each independently believes that this is the best sub-space. When each agent makes its choice, they jointly switch to the Wood/Straight sub-space.

In the baseline model, as illustrated in Figure 4, the agents switch to the sub-space with the highest perceived payoff, conditional on their mental model. In other specifications, we relax this assumption by allowing the agents to explore, i.e. sometimes they investigate sub-spaces 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.[6] This is implemented using a Softmax algorithm.[7] The temperature in the Softmax algorithm is the exploration parameter – it determines the probability with which the agent chooses an alternative that is not immediately payoff-maximizing.

Step 2: Sample from chosen sub-space: When both agents switch to their chosen sub-space, they sample a combination from within that sub-space. Since 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). Since agents do not have any knowledge of the specific coordinates that make up a sub-space, they have no control of their actual location within the chosen (coarse-grained) search space.

Step 3: Is current payoff in-line with agent expectation: 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 sub-space, 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 sub-space. For example, in Figure 4, the agent expects a payoff of 3.125, but the specific payoff they actually receive 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

[6] 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.
[7] The Appendix provides further (mathematical) details on the mechanics of the model.

expected and received payoff is a consequence of their imperfect knowledge of the sub-space, which they then try to improve. We refer to the agent's propensity to gain further fine-grained partitions (sharper vision) in the chosen sub-space as DIG.[8] DIG implies the agent expends effort in uncovering new knowledge such as by thinking about the problem, reading about it, talking to others etc. Increased fine-graining allows 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 sub-space 'Wood' more minutely, the sub-space splits into hard-wood and soft-wood (see Figure 5). The increased partitions imply that the agent's mental model of the sub-space is now more fine-grained in the row dimension; in Figure 5, it now sees four average payoff values where before it only perceived two. 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. Over time, 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.

FIGURE 5

To preserve bounded rationality and the logic of discovery in our model, we have imposed the following restrictions on the way digging leads to refinement of mental maps. In the base-line model each dig operation splits a sub-space in two along the

[8] In the baseline model, we implement a surprise driven search function (Cohen and Axelrod, 1982). In robustness checks, we implement a version of the model where the agent digs only when their payoff is less than their aspiration level, based on the aspiration level model (March, 1988). Our results are qualitatively unchanged.

agent's specialist dimension (row agent fine-grains in row, column agent fine-grains in column). The exact point where the split occurs in the sub-space is chosen at random, since the agent has no prior access to the latent choices within that subspace. Also, agents do not know when they have reached maximum fine-graining. 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 actual payoff meets expectations, the agent does not expend effort in further fine-graining the sub-space, but simply searches again; i.e., they move back to step 1. In the baseline model, under this condition, the agent samples again in the current sub-space. Resampling is accomplished by placing the agent in a random combination within the current subspace (as explained in step 2 above).[9] 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 sub-space that contains the peak. This implies the agent has achieved perfect granularity in that sub-space; however, it is unlikely that they have maximal granularity in any other region of the landscape. An alternative assumption to the baseline is that if received payoff exceeds expected payoff, the agent decides not to search any further. We examine this alternative and find 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

[9] 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 since 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 precision). 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. But "seeing that point" would mean that the agent had much finer granularity (a subspace is just another term for a knowledge partition).

results in new partitions, the agent now has more choices available, and consequently, the payoffs imputed to these choices by the agent 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 it ensures that all simulations had approached a steady state.[10]

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 division of labor, and the DIG operation leads to asymmetric mental models between these interdependent agents. In order 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 have adopted three of the integration conditions – autonomous, top-down and coordination – proposed by Gavetti (2005) to understand their effects on coordinated exploration.[11] In all these conditions, the row agent determines the row position and the column agent the column position in each time step. The agents switch to the portion of the matrix identified by their joint choice.

Autonomous: The key to this design is that the agents search in parallel, but do not make any attempt to align their mental models.

[10] 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 approaches steady state.
[11] The fourth type of integration mechanism proposed by Gavetti (2005), "circulation of cognition" is not very meaningful in our setting since it involves complete transfer of knowledge from one agent to another at some point in time.

Top-Down: In this case, we consider the situation where two agents search in parallel, and where senior management attempts to achieve coordination by imposing the same mental model on both the agents; i.e., at t=0, both agents have identical (fully aligned) partitions of the search space. This ensures that both agents identify 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. Note that the initial partitioning is fairly limited and made at random. This regime allows us to understand how initial shared knowledge impacts joint search.

Coordination: In this case, we consider the situation where 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 (column) agent requests the column (row) agent for new knowledge regarding the column (row) dimension. With each request, the column (row) agent provides the row (column) agent with one new column (row) partition that the row (column) agent does not already know.[12] The more frequent these requests, the more aligned the knowledge partitions become. However, there is an (opportunity) cost to communication. Since gaining new knowledge, by digging or by communicating, is an effortful accomplishment, each time period an agent requests the other agent for information, the requesting agent forgoes the opportunity to further improve the granularity in its own dimension in that time period. That is, an

[12] Empirically communication between specialists is very difficult, and 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 communicate every partition they know about every time they communicate, they are no longer specialists! If this were possible, mental models will be fully aligned, and the problem of coordinated exploration will not exist.

agent can only improve the granularity in one dimension at a time.[13] The frequency of communication is a parameter in the model and does not change with time.

It is important to note that in this set-up, we assume that communication effectively increases alignment of mental maps. By assumption, there is no fundamental incongruence between agents' mental maps -- *differences* between the granularities of 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.[14]

FINDINGS

Tables 2-4 summarize the final payoffs received by the agents in the different treatment conditions in this simulation model while searching the 2-peaks landscape (as shown in Figure 1). In general, the findings are consistent with the alternative hypotheses (H1b, H2b, H3b) argued for in the theory section.

<u>*1. Agents with low (non-zero) level of communication perform better than agents with high levels of communication*</u>: From table 2, we see that frequency of communication has a non-monotonic relationship with search outcomes. No communication between the agents (autonomous search) results in very poor outcomes (payoff of 0.70), communicating only once in 20 rounds results in quite good outcomes (payoff of 0.97). However, communicating very frequently, i.e. once in two rounds results in comparatively poorer outcomes (payoff of 0.91). This pattern of results is consistent with H2b.

[13] In robustness checks we relax this assumption and find that it does not qualitatively change our results.
[14] We only model perfect communication. The degree to which mental maps overlap is strictly governed by frequency of communication. The autonomous case, which we do model, logically approximates imperfect communication.

We find this pattern because too much communication leads agents to joint myopia, i.e., a pre-mature focus on local peaks. As we argued in the theory section, there is a trade-off between search and coordination. The more the mental maps are aligned, 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 regarding other regions in the landscape. This occurs because communication influences the region in which sharper vision is achieved, since 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 sub-region and neglect exploration.

Pre-mature focus is a powerful detriment to search – it prevents agents from exploring the high value region, which is a natural attractor. Figure 6 shows the search pattern of specialist agents in the two spike landscape – the left column for agents who communicate once in 20 rounds; the right column for agents who communicate once every two rounds. As the granularity figures show (Fig 6, mid-row), in the low communication case, the agents quickly identify an interesting region *from their own view-point* and then slowly try to understand the space from the other point of view. This enables them to explore effectively – their probability of digging is above zero even after 250 time steps (Fig 6, left figure in lower-row). The high communication condition on the other hand stifles exploration and agents converge very quickly – the probability of digging falls to almost zero within 50 time steps (Fig 6, right figure in lower-row).

INSERT TABLE 2, TABLE 3, TABLE 4, FIGURE 6 HERE

----

2. T-shaped skills are useful for unitary search but problematic when agents are interdependent: Table 3 shows the average payoff of an unitary searcher who is a specialist; i.e., increases granularity in only one dimension vs. an agent with T-shaped skills who increases granularity in both dimensions. 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; i.e. on average, 80% of the agents identify the high peak (=1.00) and 20% of agents identify the low peak (=0.50). In Figure 7 (panel-A), which shows the location of agents over time, we see that no agent with T-shaped skills has a final payoff of zero; they reliably identify one of the two valuable peaks.[15]

INSERT FIGURE 7 HERE

We next turn our attention to coordinated exploration achieved by teams of agents. First, we examine the case of specialists. As expected, we see in Table 3 that when *autonomous*, i.e., they make no effort to align their mental models, their payoff of (0.70) is higher than the unitary specialist searcher, but lower than the unitary searcher with T-shaped skills. From Figure 7 (panel-B), we see 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.

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

[15] An individual searcher need not be equally good in both dimensions. In robustness, we find 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.

a flawed mental model of the search space. It is worthwhile to note that in this condition, at steady state, the dig probability of the average specialist agent is almost 0.9; i.e., the agents are still attempting to find the valuable combination despite the negative feedback. Such confounding is impossible in search by a unitary agent. Since agents do not realize when the maximum level of granularity is reached, they try to continually tweak their experiments in the hope of finding better solutions.[16]

Our argument in hypothesis 2b is that agents with T-shaped skills, though possessing superior knowledge of the landscape are still subject to mutual confusion. From Table 3, we see 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. Since these two agents do not have aligned knowledge partitions, they are unable to make choices that are jointly valuable. These poor results are obtained despite that fact that at steady state, the two agents, each have almost as fine-grained partitions of the search space (average of 8.4x8.4) as the unitary searcher at the end of the simulation (8.6x8.4). In other words, the agents achieve very poor outcomes in joint search even though each agent is equipped with a set of traits that give them the same potential to acquire knowledge as the individual searcher. This result is an interesting contrast to joint search models *without* epistemic interdependence. In prior work, joint searchers reach good solutions if they both have very good processing power (i.e., power to consider also the complementary dimension) as long as they do not pre-maturely weed out solutions (Rivkin and Siggelkow, 2003; Siggelkow and Rivkin, 2006).

[16] Situations where 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 geo-centric 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 since the rabies virus did not travel by blood like anthrax and the other pathogens he had worked with previously.

Hierarchy as a coordination device: This problem of mutual confusion in coupled learning problems should reduce when agents are provided with a coordination device, i.e., some ability to align their mental models of the search space. Communication is one means of aligning mental models, as discussed earlier. 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 search autonomously. From table 4, we see that when these agents are provided with an 8x8 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 sub-space 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 are 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 increasing predictability of actions (Puranam et al, 2012; Srikanth and Puranam, 2011; Okhuysen and Bechky, 2009).

3. Impact of exploration on unitary search vs. coordinated exploration: An agent explores when it chooses actions without regard for its immediate payoff. In our model, we investigate the impact of exploration, by varying the curiosity parameter 'τ'. As τ

increases, the more likely the agent switches between sub-spaces without regard to the immediate expected payoff from that sub-space.

The first result we observe is that τ=0 outperforms τ>0 across all conditions; the higher the curiosity (τ), the worse the performance. Whereas the unitary searcher suffers relatively less than the joint searchers for all values of τ tested, the performance of the autonomous specialists is reduced the most. The results further suggest that the negative consequences of increasing mutual confusion dominate the slow learning effect as agents' curiosity increases (starting from low values of τ). Our results also support our intuition that increasing alignment of mental models should help reduce mutual confusion. As agents increase exploration, the top-down condition performs better than the autonomous condition, and agents that communicate more frequently outperform those that communicate less often. These results are consistent with our argument in hypothesis H3b, at least for the 2-spike landscape.

Robustness Checks

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 strengthen our intuition regarding the above results.

1. Different landscapes: Even if we explained our main results with reference to the 2-spikes landscape (figure 1), we did test our results with other landscapes shown in Figure 8 as well as with a few randomly generated landscapes. The random landscapes involved few (5) or many (30) combinations with non-zero value, placed randomly in the landscape. In these landscapes, the global peak was set to a value of 1 and other peaks were drawn from a uniform distribution between 0 and 1. We also generated a landscape with multiple peaks with a gradient that allows for hill climbing, a property that is not clearly present in any of the spiked landscapes.

INSERT FIGURE 8 HERE

In the single spike landscape (figure 8, panel C), autonomous searchers always find the global peak. For profit seeking agents, the one spike acts as an unambiguous attractor as they fine-grain their mental-models. As each agent chooses a good solution for itself, the two agents jointly land on the valuable combination. This is in contrast to the 2-spike landscape, where coordination is important because the shorter peak is an alternative attractor that confuses the agents with incongruent mental models.

The ridge landscape (figure 8, panel D) has several equally valuable solutions, and consequently emphasizes coordination over search when compared with the 2-spike landscape. In the ridge landscape, autonomous search leads to an average payoff of 0.03-0.04, which is about the same as the randomizing payoff expected in a matching game over this sub-space. In the top-down case, our intuition suggests that performance should improve because the initial alignment in partitioning the search space reduces the number of confusing peaks and thereby improves coordination. In line with this expectation, the payoff for the ridge-landscape increases to 0.38. In the coordination condition we find that agents who communicate more frequently perform better than the agents who communicate infrequently. This result is different from the 2-peak case, because here there is no need to search for a tall peak, since all peaks are of equal height. Therefore, coordination is the only problem the agents face, and as expected, higher levels of communication lead to better coordination. Here, in the absence of the search imperative, the prediction from the information processing theory (H1a) holds.

The ridge-peak landscape (Figure 8, panel B) lies in between the ridge and the 2-spikes landscapes. It simultaneously taxes search and coordination; the higher peaks act as attractors towards that region in the landscape, but since there are several of them close to each other, coordination becomes that much more important. This suggests that the optimal level of communication should lie in-between the 2-peak and the ridge landscape. Our results confirm this intuition. Also, similar to the ridge case, in the top-down condition, we find that higher the initial granularity, the higher the performance.

The random landscapes generated with 5 and 30 peaks respectively also provide evidence in support of the same intuition. As the number of peaks increases, the increasing number of possible combinations challenges coordination. As coordination becomes more challenging, the performance of the autonomous agents declines. This effect can be countered by tighter alignment of mental models, but only to some extent, because even if coordination helps agents in jointly targeting some peak, it is unlikely to be the global peak. We obtain similar results for the multi-peak landscape as well. [17]

These results suggest that the balance between H1a and H1b is likely to depend on the emphasis the landscape lays on search vs. coordination for successful outcomes. In landscapes that emphasize search, a lower level of communication is better. In contrast, in landscapes that emphasize coordination, agents with higher levels of communication perform better.

2. Initial Conditions: We checked to see whether and how the agents' initial granularity of vision affects these results. Specifically, agents could be natural born specialists in the sense that they already have a high level of granularity in their "own" dimension when they begin the search task. Obviously, altering the prior knowledge of the agents may impact the effectiveness of different coordination mechanisms.

[17] These results are available from the authors.

We experimented with initial granularities (own dimension, other dimension) of (2, 1), (4, 1), (8, 1), (16, 1) and (24, 1) for each agent. An agent with very high initial granularity in their own dimension and very little in the other dimension is a specialist. In this condition, by definition the agents' initial granularities are misaligned. Our theory suggests that specialists who do not communicate adequately should perform worse than those that do communicate very often, since communication promotes alignment of mental models in coordinated exploration. The flip-side of alignment, narrowing of search, perhaps may not pose as big a challenge for these agents, since high initial granularity implies an already superior understanding of the landscape from their own dimension. Our results support this intuition – agents with (24,1) initial granularity perform equally well under the high communication condition as under the low communication condition in the 2-spike landscape. In contrast, agents with (2,1) initial granularity, perform better when they have little, but non-zero communication.

3. Different Search Algorithms: As a final robustness check we investigated the effect of a different search algorithm that regulates the DIG behavior of the agents. In the standard model, the agents are surprise-driven: agents increase granularity if they receive a payoff that is different from what they expect (Cohen and Axelrod, 1984). We also modeled a satisficing, aspiration driven model process, where agents continue to increase granularity until they achieve a pre-specified threshold payoff (March, 1988). We experimented with different thresholds, including 0, 0.25, 0.50, 0.75 and 1.0. If the threshold is zero, the agent does not search. If the threshold is 1.0, the agent continues to search until they have identified the global peak.

Our results for all the landscapes except the multi-peak landscape are qualitatively similar to the base-line model. Specifically, as long as the threshold specified is greater than zero, agents achieve approximately the same final payoff in this model as in the baseline model. However, in the multi-peak landscape, agents perform better as the threshold increases. This is because, in this landscape, the peaks have gradients; increasing thresholds ensure that the agents continue to "hill-climb" in pursuit of better performance. Our primary result that "too much communication" is worse than "limited communication", also holds here at all levels of the "threshold".

The fact that our results are robust to changes in this assumption shows that the specific choice of stopping criterion does not transform or eliminate the core problem of coordinated exploration -- where updating of interdependent mental models and confounded feedback (arising from joint actions) challenges adaptive behavior.

IMPLICATIONS AND SUGGESTIONS FOR FUTURE WORK

Organizations exist to manage the trade-off that arises with the division of labor: benefits from increasing specialization vs. losses arising from the need for coordination. Coordinated exploration – the condition where specialist searchers need to coordinate their choices is a significant problem for organizations, but is inadequately addressed by prior work. Prior theories on coordination ignore search. In contrast 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, 1997; Levinthal and March, 1981). Neither of these approaches is helpful to understand coordinated exploration.

Our contention is that predictions from prior theories are incorrect when applied to situations of coordinated exploration, because the simplifying assumptions used in prior work 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 (Tushman and Nadler, 1978; Galbraith, 1977). 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, as 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 coordinating search, i.e., increasing mutual-confusion that again decreases performance in coordinated exploration.

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 action of others, thereby 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. Aligning mental models promote 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.

We argue that 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.

The novel mechanisms explicated in this study throw some light on resolving long-standing empirical contradictions as well as offer some novel predictions. Some empirical work in NPD suggests that high levels of communication improve innovativeness in NPD performance whereas other studies suggest the opposite (Brown and Eisenhardt, 1995; Krishnan and Ulrich, 2001; Tyre and Hauptman, 1992; Montoya-Weiss and Calantone, 1994). This contingency effect is theoretically ill understood.

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 is dependent on the nature of the task environment. The more the landscape emphasizes search over coordination, the more detrimental 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 NPD 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 in their own domains, the lesser the 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". Fleming (2007) from an analysis of patent data suggests 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 (2013) argue that higher levels of common ground are found within firm boundaries than across them. Coupled with the findings from this 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. Understanding 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 achieved typically by weakening the sensitivity of agent actions to its 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, 1989; Puranam and Swamy, 2011), our result is not achieved by reducing agents' sensitivity to payoffs, but come about because search precedes coordination. It is here interesting to note that Lazer and Friedman (2007) in a model of network search (based on the NK framework) reached a similar conclusion as we did 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 regarding 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 their payoffs were independent, i.e. 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 its impact on subsequent exploration. Typical slow learning models do not capture these dynamics.

The contrast with slow learning models does pose the question why the agents are unable to acquire more information and then use this information in a sensible/optimal way to guide both exploration and coordination. The answer to this question lies in the assumptions we make. These include (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 this is the case. (4) Agents act in parallel – 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.

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 payoff, 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 trade-off can be managed. However, this "explore, then coordinate" approach can be fairly time consuming, potentially expensive (some experiments are fairly expensive 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 regarding 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. However, establishing such common knowledge in effect implies the lack of specialization. One way to prevent this may be if the agents established a principal-agent relationship or established rules for sequential search; for example 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, where a superior directs search by aligning mental models upfront, which only has limited impact on improving performance. This is because in order to effectively coordinate, the hierarchy needs to be informed either of the location of the global peak before-hand or of the emergent knowledge (new partitions) of both the specialists. The first condition implies 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 search efforts. Therefore, 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 by sequencing, either of actions or by appointing agents with different skills at different stages of the problem. Examining these options are good avenues for future research.

Limitations: This work is subject to the following limitations. First, the landscape of innovation is exogenous to the model. Since organization of joint search is contingent on the type of innovation landscape, how does a manager know what type of landscape she is 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 from any corner of the landscape to any other, which may not be feasible for reasons 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 performance of agents with asymmetric abilities. More interestingly, we have not explored sequencing – such as first searching with generalists, 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: Despite these limitations, we believe our model makes some important contributions. This is one of the first efforts to model search by considering both 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 agent cognition. We show that joint search is not scaling up of individual search, but is qualitatively different. By employing a richer modeling strategy, we are 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. It should be noted that our mechanism is robust to a number of checks – including the nature of the landscape (Rivkin and Siggelkow, 2007), differences in initial cognition of the agents (Gavetti, 2005), and to the specific assumptions of the search algorithm – surprise driven vs. payoff-driven (March, 1988; Simon, 1962; Cohen and Axelrod, 1982). As a methodological contribution, we also provide an alternative modeling platform where it is possible to understand the consequences of different kinds of common knowledge assumptions. Finally, our work has some very interesting implications for game theory. Prior games have either considered 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 and is likely common in organizational situations. Future work in this area could be very interesting.

CONCLUSIONS

The aim of this paper has been to examine problems of coordinated exploration. We introduce an analytical framework that can help differentiate coordinated exploration from individual search. Epistemic interdependence requires a trade-off between search and coordination that is absent in unitary search. Too much exploration without aligning mental models leads to mutual-confusion. Too much alignment stifles exploration by promoting joint-myopia. We argue that predictions from models of unitary search likely lead to erroneous conclusions when applied to joint search, and explore how organizations can manage this trade-off in joint search. In and of itself the consideration of richer forms of mental representations fills a void in the literature on organizational search and learning that inform our understanding of fundamental activities in firms such as joint problem solving. Our treatment is suggestive rather than exhaustive, but we hope that we have demonstrated why it is important to further our understanding of the relationship between organization and joint search.

REFERENCES

Bettenhausen, K. L.
1991 “Five Years of Groups Research: What We Have Learned and What Needs to Be Addressed.” Journal of Management, 17: 345–381.
Brandenburger, A.
2008 “Epistemic Game Theory: Complete Information”. In L. Blume and S. Durlauf. (eds.), The New Palgrave Dictionary of Economics. Basingstoke, UK: Palgrave MacMillan.
Brown, S. L. and K. M. Eisenhardt
1995 “Product Development: Past Research, Present Findings, and Future Directions.” The Academy of Management Review, 20: 343–378.
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, MA: Harvard Business School Press.
Cohen, M. D. and R. Axelrod
1984 “Coping with Complexity: The Adaptive Value of Changing Utility.” The 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.
Daft, R. L. and R. H. Lengel
1986 “Organizational Information Requirements, Media Richness and Structural Design.” Management Science, 32: 554–571.
Davis, J. P., K. M. Eisenhardt
2011 “Rotating leadership and collaborative innovation: Recombination processes in symbiotic relationships” Administrative Science Quarterly, 56: 159–201.
Donaldson, L.
2005 “Contingency Theory of Organizations.” Thousand Oaks, CA: Sage.
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(5): 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: 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, NY: Wiley.
Gomes, J. F. S., P. C. de Weerd-Nederhof, A. W. Pearson, and M. P. Cunha
2003 “Is more always better? An exploration of the differential effects of functional integration on performance in new product development.” Technovation, 23: 185–191.
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 Vol. 22: 153–193. Greenwich, CT: JAI Press.
Heifetz, A.
2008 “Epistemic Game Theory: Incomplete Information”. In L. Blume and S. Durlauf. (eds.), The New Palgrave Dictionary of Economics. Basingstoke, UK: Palgrave MacMillan.
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.
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, MA: Harvard University 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.
March, J. G.
1988 “Variable risk preferences and adaptive aspirations.” Journal of Economic Behavior & Organization, 9: 5–24.
March, J. G.
1991 “Exploration and Exploitation in Organizational Learning.” Organization Science, 2: 71–87.
March, J. G. and H. A. Simon
1958 “Organizations”. New York, NY: 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.” The 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, XLII: 367–403.
Schelling, T. C.
1960 “The Strategy of Conflict.” Cambridge, MA: Harvard University Press.
Scott, W. R.
2003 “Organizations: Rational, Natural, and Open Systems.” Englewood Cliffs, NJ: Prentice-Hall.
Selten, R. and Warglien, M.
2007 “The emergence of simple languages in an experimental coordination game”. Proceedings of the National Academy of Sciences. 104(18): 7361-7366.
Shenhar, A. J., A. Tishler, D. Dvir, S. Lipovetsky, and T. Lechler

2002 "Refining the search for project success factors: a multivariate, typological approach." R&D Management, 32: 111–126.
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.
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
2013 "Firm as a coordination system: Evidence from software services offshoring". Organization Science (forthcoming).
Sutton, R.S., and A.G. Barto.
1998 "Reinforcement learning." Cambridge, MA: MIT Press.
Thompson, J. D.
1967 "Organizations in Action." New York, NY: McGraw-Hill.
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." The 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.

Table 1: Table of parameters

Parameter	Range	Purpose
Search landscape	64 x 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 x 1; 8x8 2x1, 4x1, 8x1, 16x1, 24x1	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.
Curiosity Parameter in switching	τ: 0 – 0.1	The agents' propensity to engage in explorative activity. This governs the agent's move between choices it 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 (column) agent receives new knowledge provided by the column (row) agent. When probability is zero, agents do not communicate. When 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 the probability with which a row (column) agent's dig attempt results in a partition in the search space in the column (row) dimension.

* In this paper, we do not vary the size of the search landscape.

Table 2 – Impact of level of communication on performance.

Search Regime	Propensity to partition in other dimension	Comm Freq	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	Comm Freq	Final Payoff (at T)
Unitary Search	0	0	0.01
	0.50	0	0.90
Autonomous	0	0	0.70
	0.50	0	0.60

Table 4 – Coordination by Hierarchy.

Search Regime	Propensity to partition in other dimension	Comm Freq	Final Payoff (at T)
Top Down	0	0	0.85
	0.50	0	0.85

In all of the above results, the curiosity parameter τ =0. Initial granularity is set at 1x1 in all results except for the Top-Down search regime, where it is set at 8x8. For specialists, the propensity to partition in the other dimension is 0; for agents with T-Shaped skills it is 0.50. Results obtained at T= 500. Reported results are averages obtained from 300 runs for each condition. Reported differences are statistically significant at conventional levels.

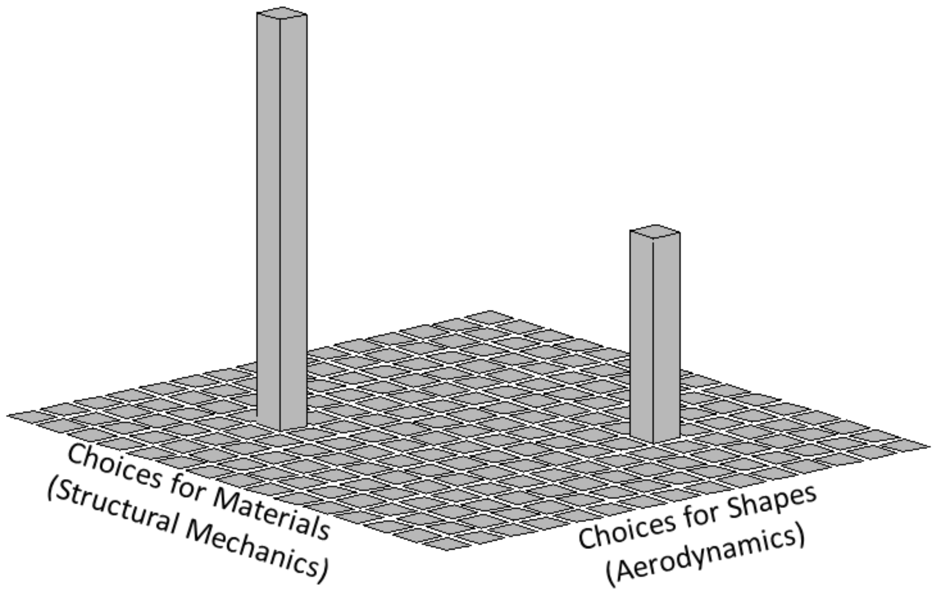

FIGURE 1: Task Environment – 1/4th the size of actual task environment
The search space consists of a 64x64 matrix of possible materials and shapes from which the specialist searchers can choose from. Each combination is associated with a payoff value. In this example, only two combinations are associated with a positive payoff, with one of them being twice as valuable as the other.

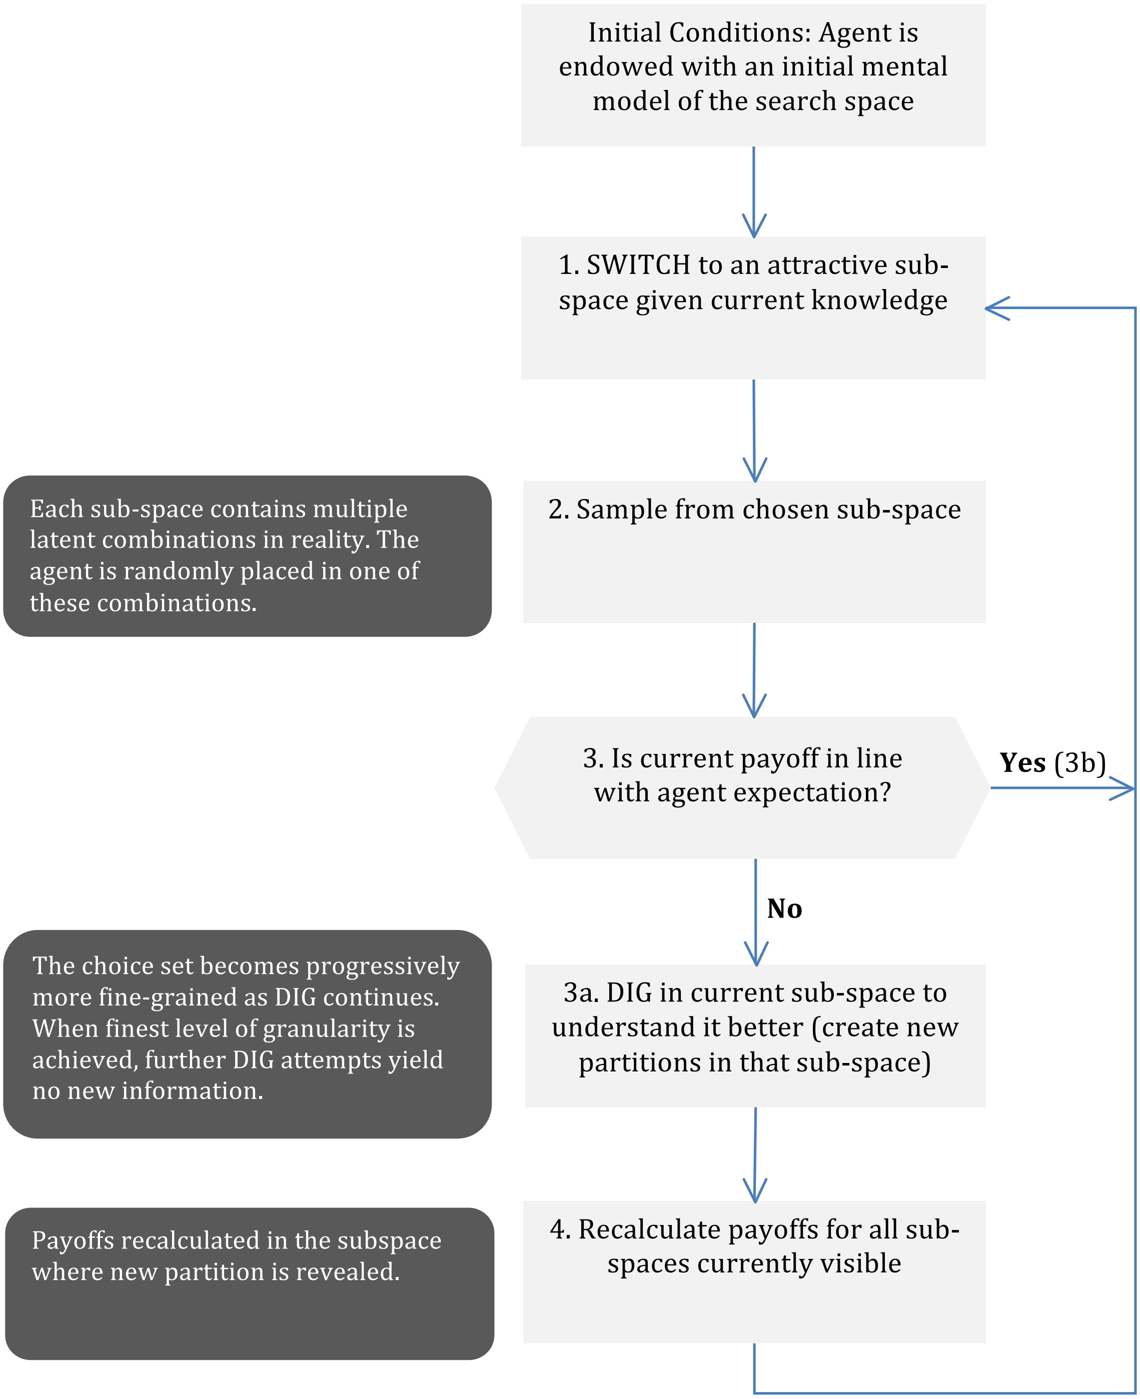

FIGURE 2: Baseline Search Algorithm

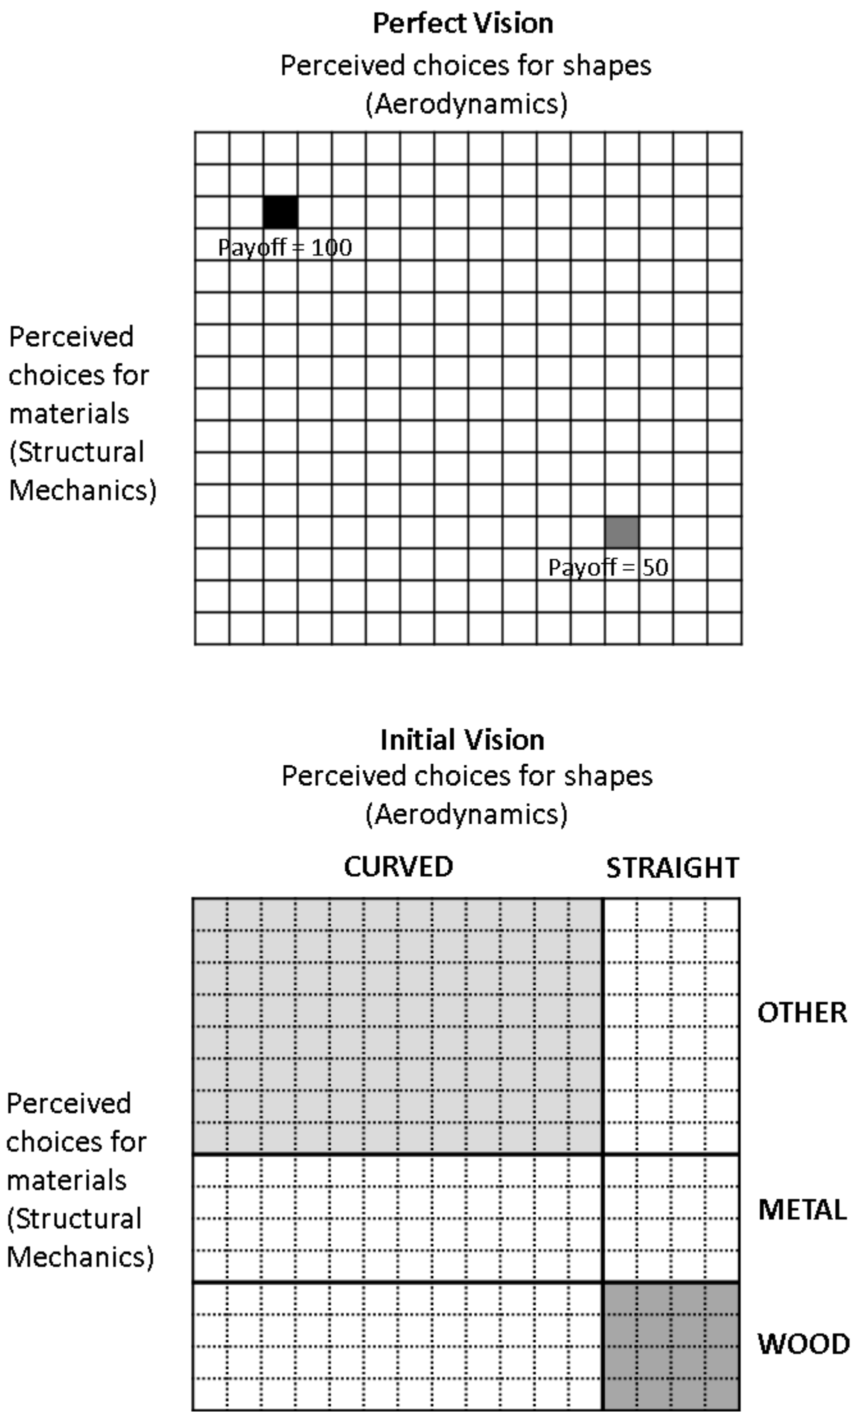

Figure 3: Perfect vs. imperfect initial vision of task environment (1/4 of actual size)

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

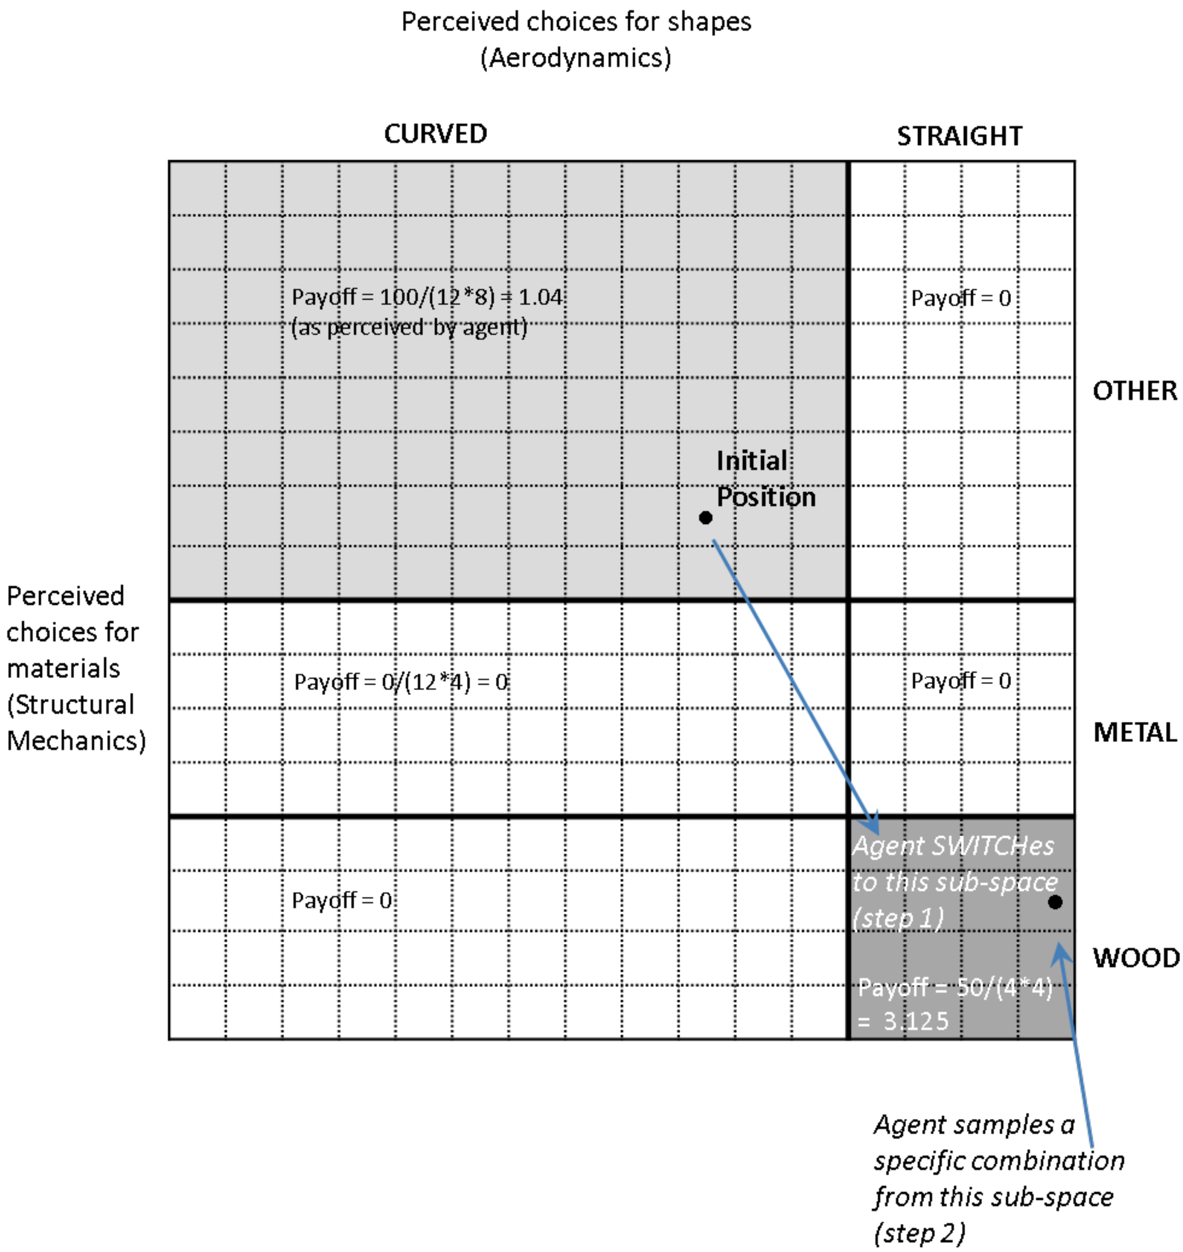

FIGURE 4: Agent Vision and Switching Operation to more promising alternative (1/4th of actual size)

Agent initially sees a 3x2 search space with associated payoff for each sub-space. Agent switches to the most promising alternative, depending on observed payoff (step 1). Agent samples a specific combination (at random) in the chosen sub-space (step 2).

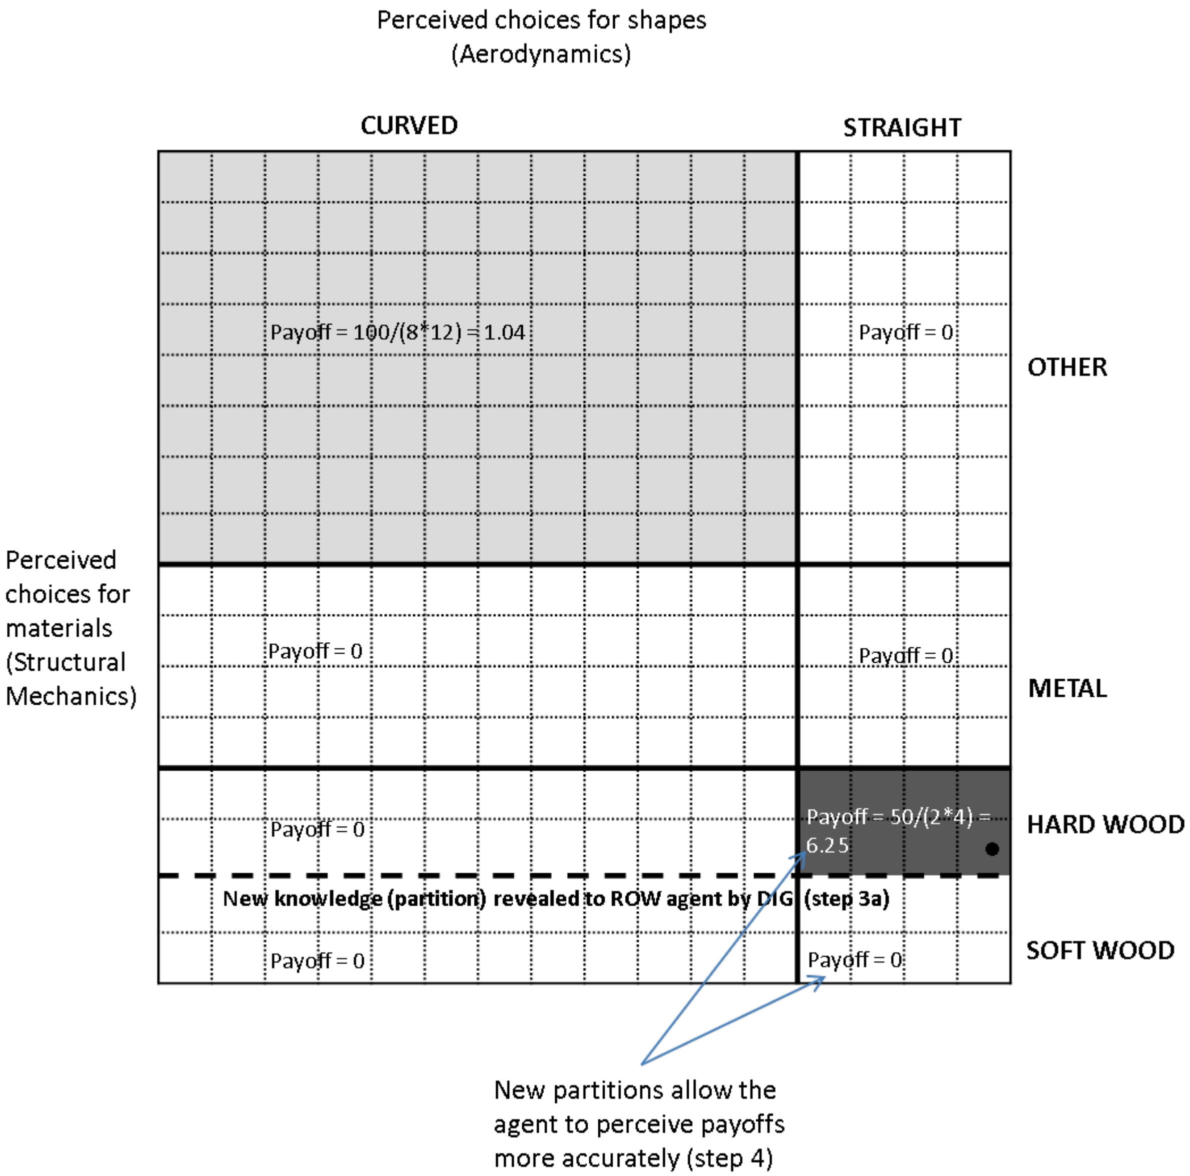

FIGURE 5: Dig Operation: New Information revealed to row agent
A new knowledge partition occurs (step 3a). The agent now sees four sub-spaces (hardwood-straight, hardwood-curved, softwood-straight, softwood-curved), when earlier it perceived only two (wood-straight, wood-curved). The agent now calculates expected payoffs for each of these four sub-spaces (step 4).

Note that when the column agent digs a new knowledge partition occurs in the column dimension. In figure 4, both the row and column agent perceive a 3x2 matrix. After dig, the row agent perceives a 4x2 matrix (as above), the column agent would perceive a 3x3 matrix.

LOW Communication (1/20) **HIGH Communication (1/2)**

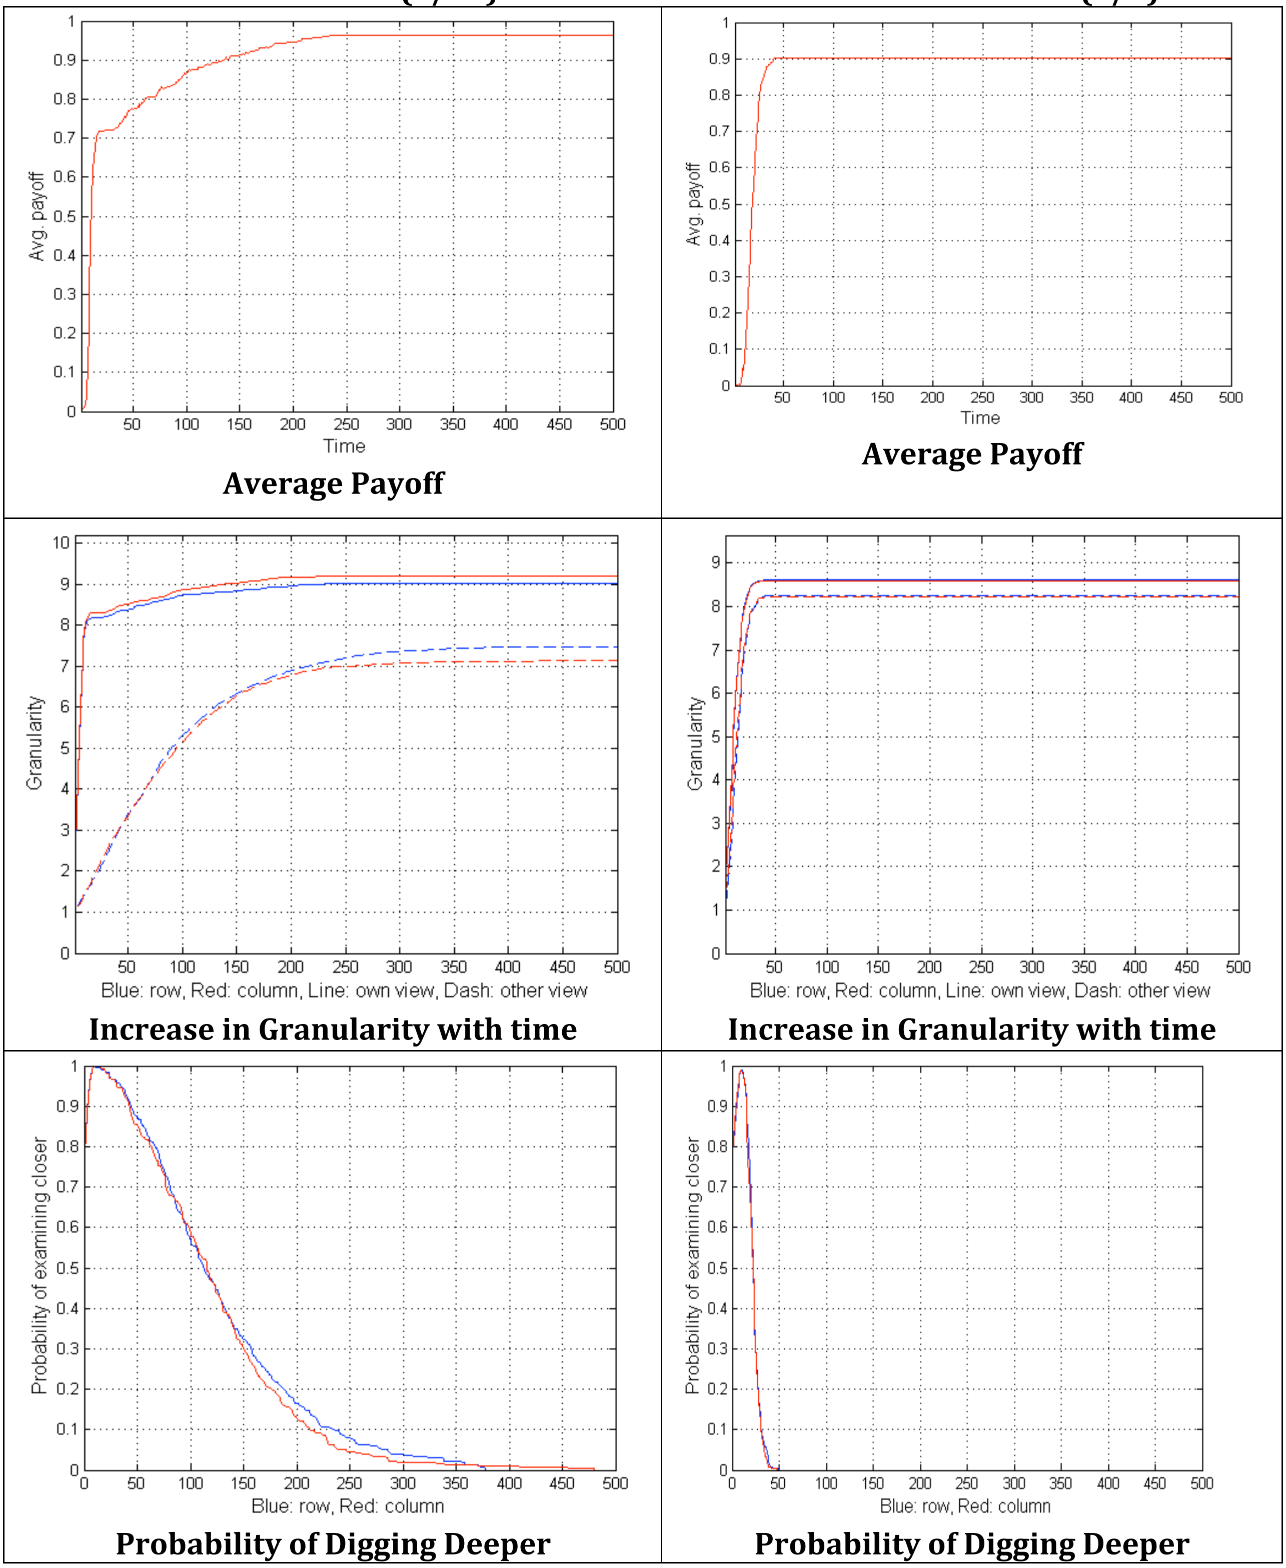

Figure 6: Search by Specialist Agent in 2-Spike landscape – Impact of frequency of communication on payoffs, granularity and dig probability.
In all of the above results, curiosity parameter, τ =0; initial granularity is 1x1; propensity to partition in other dimension is 0.

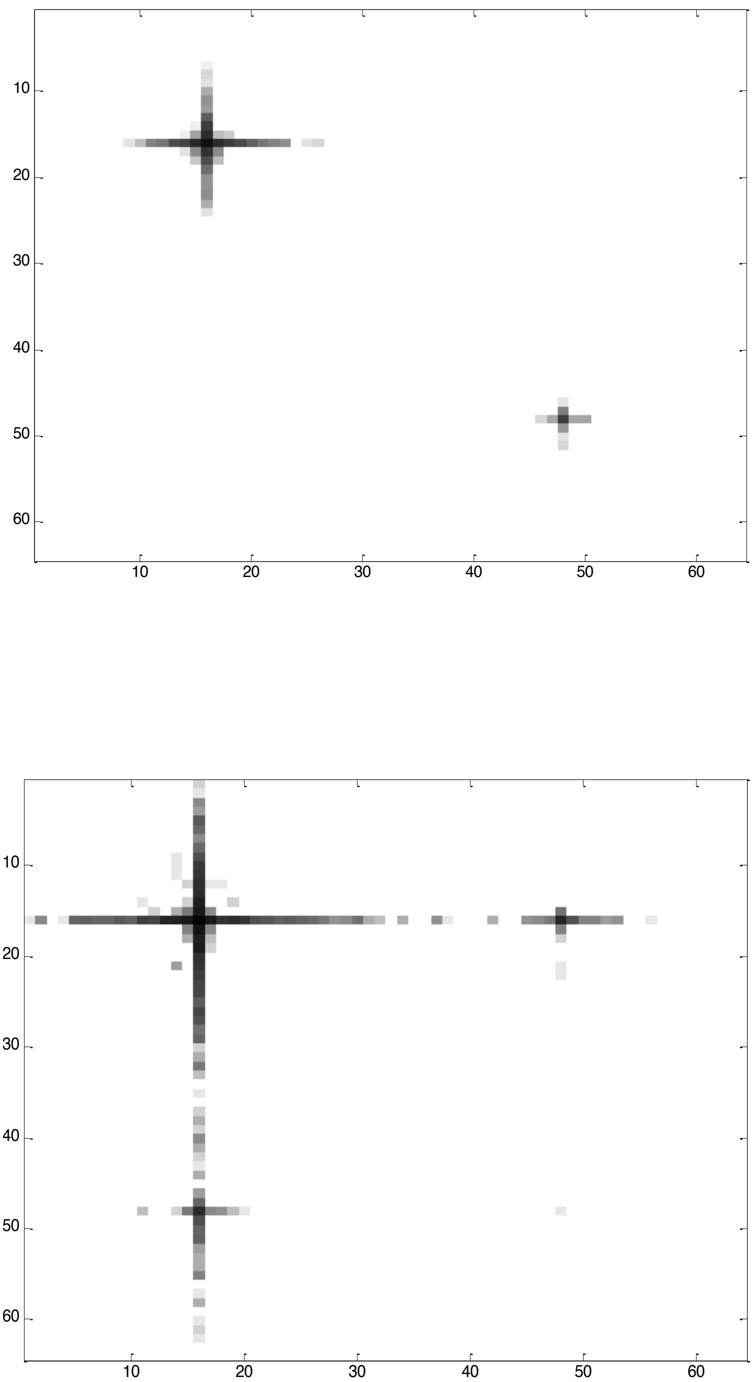

PANEL A:
Unitary Search;
agent propensity to partition in other dimension =0.5.

PANEL B:
Autonomous Joint Searchers
Agent propensity to partition in other dimension =0; communication frequency =0.

Figure 7: Average position of agents in the 2-peaks landscape.
In the above results, initial granularity is 1x1, and curiosity parameter, τ =0. Results obtained at T= 500. Figure represents average over 300 agents.

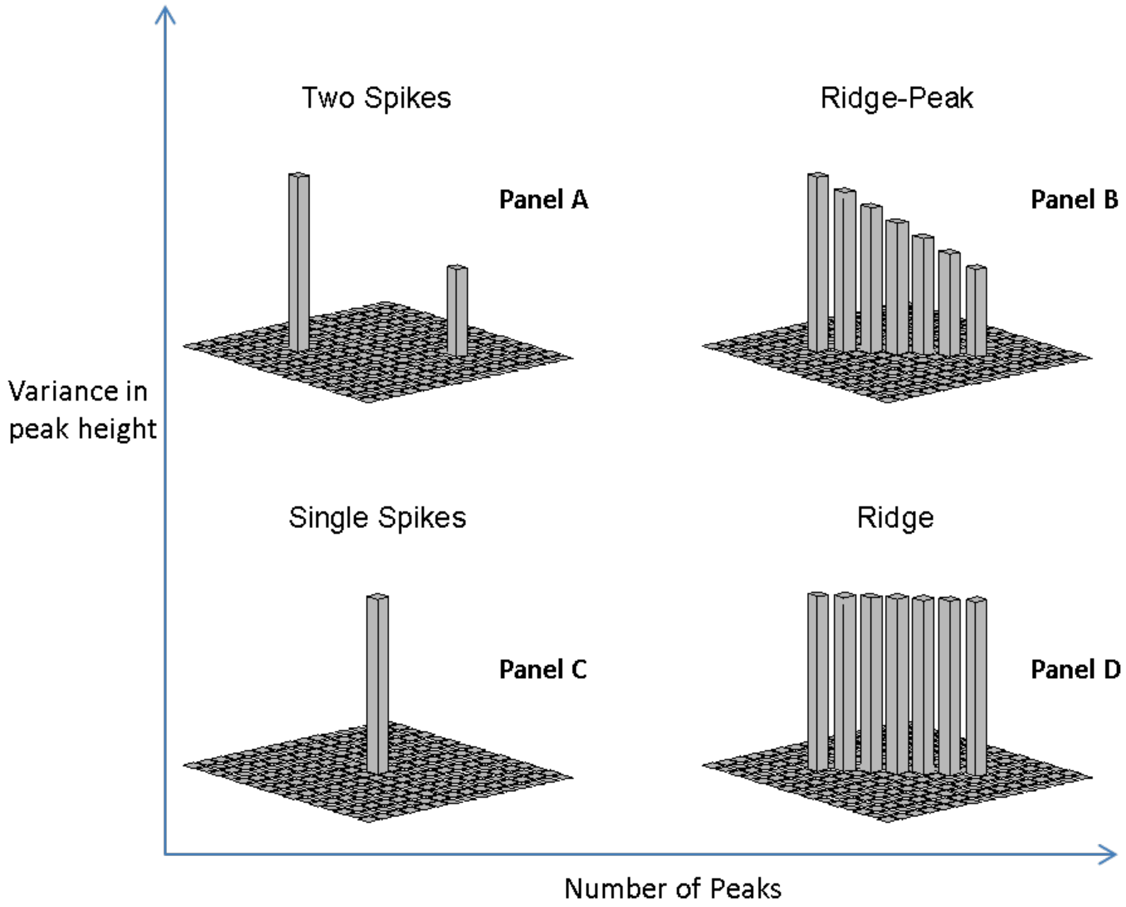

Figure 8: Task environments (1/4 of actual size)

Appendix 1 – Model Mechanics

This appendix provides the mathematical details of agent behavior in our model. The switch and dig operations of the agents are governed by the softmax algorithm.

<u>*Switch*</u>: The Softmax algorithm allows tuning of the agent's propensity to apply the *switch* from the current cell to another cell. More formally, according to the Softmax algorithm, the agents' probability of sampling a particular cell in the landscape *i* at time step *t* is dependent on that cell's observed mean performance $\overline{x_i}$, when compared to the mean performance across all perceived cells *d*.

$$p_{it} = \frac{e^{\frac{\overline{x_i}}{\tau}}}{\sum_{i=1}^{d} e^{\frac{\overline{x_i}}{\tau}}}$$

This is the standard version of the Softmax algorithm (Luce, 1959; Sutton & Barto, 1998). The "temperature" τ (0<τ≤1) is the agent's exploration parameter – the higher the τ, the more curious the agent is and the more likely that the agent switches among the known cells in a random manner (ignoring expected payoffs). It is important to note that the Softmax algorithm operates on divisions of the search space as defined by the agent's emergent mental maps (see Figures 4 and 5 for examples). If an agent holds a mental map as illustrated in Figure 4, this agent perceives a total of *d*=8 cells. The Softmax algorithm then computes the probability of locating in each of these eight cells and thereby switching from the current cell to a different portion of the search space.

In each cell, the agent is positioned at a random combination. In consequence of not being able to see the latent combinations underlying a cell, the agent in effect performs a random walk among them – the probability distribution over each of the underlying positions in the reality matrix is uniform. The agents receive a payoff for the position in the reality matrix where they are actually situated.

<u>*Dig*</u>: The probability of digging into a particular cell in the landscape *i* at time step *t* depends on the difference between that cells observed mean performance $\overline{x_i}$ and the

actual payoff x_i received from the agent's current the position in the reality matrix – and on the propensity λ to apply the dig operation:

$$q_{it} = 1 - \frac{1}{e^{\frac{\lambda|x_i - \bar{x}_i|}{\bar{x}_i}}}$$

This algorithm defines a probability distribution over differences between expected and actual performance. Higher values of λ makes the agent much more sensitive to deviations from expected performance, which results in application of the dig operation and an increase of granularity in the cell where the agent is currently active (see Figure 5 for an example).[18]

We also performed a threshold-driven dig operation as a robustness check instead of the above surprise-driven formulation. This is given by

$$q_{it} = 1 - \frac{1}{e^{\frac{\lambda(\rho - x_i)}{\bar{x}_i}}} \quad if\ \rho > x_i$$

$$q_{it} = 0 \quad if\ \rho \leq x_i$$

where ρ is the threshold payoff above which the agent stops searching, and the actual payoff received is x_i .

[18] Based on extensive sensitivity tests, we chose actual values of λ in the interval [1,10] as shown in Table 1. A value of λ= 10 corresponds to a very high propensity to dig . By contrast, a value of λ= 1 corresponds to a rather low propensity to dig. Our main results are robust the value of λ. We report results for λ= 1.