DRUID Working Paper No. 11-16

Specialize or Coordinate?
The impact of organization design on shared mental representations and innovation outcomes in joint search

By

Thorbjørn Knudsen and Kannan Srikanth

Danish Research Unit for Industrial Dynamics

www.druid.dk

Specialize or Coordinate?
The impact of organization design on shared mental representations and innovation outcomes in joint search

Thorbjørn Knudsen
Strategic Organization Design Unit
University of Southern Denmark
DK-5230 Odense, Denmark
+45 6550 3148
E-mail: tok@sam.sdu.dk

Kannan Srikanth
Strategy Area
Indian School of Business
Gachibowli, Hyderabad 500032, India
+91 40 2318 7165
E-mail: Kannan_srikanth@isb.edu

February, 2011

Abstract:
The integration 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. However, apart from the formation of shared mental models, organization also impacts the locus of search and specialization in innovative conditions. This interplay between the emergence of self (specialized) knowledge and shared knowledge is likely to significantly impact joint search outcomes. In this paper we propose a computational model that examines the trade-off in the development of specialized vs. shared knowledge and how it impacts joint search for different innovation landscapes and different levels of agent ability. Our results point to the importance of not generalizing from single searcher problems for joint searcher problems.

Keywords: Organization Design ; Coordination ; Joint Search; Shared Mental Models

Jel codes:

ISBN 978-87-7873-327-6

INTRODUCTION

Innovation as an outcome of recombinant search among technological alternatives is quite well established in the organization and technology management literature (Schumpeter, 1939; Henderson and Clark, 1990; von Hippel, 1988; Iansiti, 1998; Fleming, 2001). However search, typically, is not easy: few combinations yield high value innovations while most combinations are useless. The prevailing imagery is one of landscapes with high peaks and low valleys. Innovators are explorers, who attempt to find the high peaks. However, while most prior work on innovative search characterizes search by an individual, in practice, most innovations and new product development efforts are characterized by multiple specialists jointly searching for the next breakthrough. How to best *organize search processes involving the joint effort of multiple specialists* is one of the most important questions faced by students of organization and innovation alike.

Typical models of innovative search have proposed a boundedly rational, experience-driven model of the search process (Nelson and Winter, 1982; March, 1991). However, others have argued that bounded rationality merely implies imperfect knowledge of available actions and potential outcomes – and does not negate search based on the logic of consequences (March, 1994; Knudsen and Levinthal, 2007). Integrating these thoughts, Gavetti and Levinthal (2000) proposed a model of individual search that took into account both cognitive and experiential bases of search. Drawing on psychological work on individual behavior (Thagard, 1996; Holland et al, 1986), they proposed that agents formed imperfect mental maps of the world and proceeded to search on the basis of these maps, while feedback from experience alters these initial mental maps (Louis and Sutton, 1991; Weick, 1995). In other words, search is a dance between cognition directed action and experiential feedback changing cognition.

The dynamic relationship between cognitive and experiential search, while relatively straight-forward in individual search poses some intriguing challenges in the context of joint search by (multiple) specialists. This is because joint searchers have to solve two problems – (1) determine the locus of search from their own point of view and (2) coordinate their choice with that of the other specialists.

While it is intuitive that shared knowledge among specialists regarding how their actions are interdependent should result in better outcomes (March and Simon, 1958; Huber and Lewis, 2010), developing such shared understandings is usually difficult (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; Bechky, 2003). Organizations employ different kinds of "integration mechanisms" for the purpose of developing such shared knowledge (Lawrence and Lorsch, 1967; Clark and Fujimoto, 1992; Iansiti, 1995). Hoopes and Postrel (1999) show that such integration mechanisms improve product development performance specifically by improving shared knowledge among specialists.

However, different integration devices likely accumulate 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 innovative search in two ways: (1) emergent shared knowledge affects cognition that directs search. (2) The resulting sampling of the search space influences feedback and further refinement of mental models (Denrell and March, 2001). The sequential interaction of environmental variables and internal states (of agents) is a common feature of behavioral and evolutionary models (March and Simon, 1958; Nelson and Winter, 1982). However, the aggregation of these effects across multiple interdependent actors can fundamentally change *joint search* behavior from individual search.

In this paper we investigate the following research question: How do different integration mechanisms influence joint search outcomes in different innovation landscapes? We explore this question in a computational model in the context of understanding the impact of these integration mechanisms on the mental maps (or cognition) of individual specialist agents and its subsequent influence on search behavior. This question is important to investigate because in the economy, most innovations are outcomes of joint search efforts by multiple interdependent specialists. If integration systematically affects the likelihood of finding incremental vs. radical innovations or the timing of innovations, this may have a large impact on how best to organize for innovations.

We find that fruitful joint search involves a tension between adequate scanning for good combinations vs. jointly targeting a specific combination. This tension arises because broader scanning increases the potential target area and thereby delays or even obstructs targeting. Organization design moderates this tension. This has the surprising implication that when agents put forth high effort in coordination rather than in further specialization, they are more likely to find incremental rather than radical innovations. Interestingly, we also find that the impact of organization – the design of integration mechanisms or the design of interdependence - is much greater than the impact of individual agent abilities for joint search problems.

At a broader theoretical level, our results point to some understanding of the need for shared knowledge. Agents in our model simply know (some of) the same things – they do not know that the other knows it as well, but are still able to solve coordination problems. Indeed, we find that the design of coordination mechanism for a team of innovators is just as important as the employment of individuals with useful skills. This has some implication for our understanding of the need for shared vs. mutual vs. common knowledge for coordination. In particular, it appears that standard game-theoretic treatments of coordination problems place unnecessary (and unrealistic) demands on the requisite knowledge conditions to solve joint search problems.

This paper is organized as follows. First we briefly introduce prior models of joint search and define the gap we aim to fill. Next we detail our modeling framework and then report our findings. A discussion of our results and their implications for theory and practice follow.

PRIOR MODELS OF JOINT SEARCH

Search for innovations (and many other activities in the modern economy) typically involves joint contributions from functional specialists with limited understanding of each other's domains (Postrel, 2002). There are two essential features to these search problems: first, the *coordination condition* – where acceptable solutions have to meet the joint feasibility criteria of both agents. Second, the *bounded rationality condition* – specialists do not know the entire universe of solutions from which they select a subset that satisfy their constraints. Rather, agents search in order to identify potential solutions which they then evaluate for fit. We view these two conditions (coordination and bounded rationality) as defining features of joint search models.

Prior models of innovative search have developed in two mutually exclusive strands. The first strand, pioneered by Gavetti and his co-authors considers the importance of cognition or simplified mental models of the search space in order to find high performing configurations (Gavetti and Levinthal, 2000; Gavetti, 2005; Gavetti, Levinthal and Rivkin, 2005). Given our research focus, this stream of work has two limitations. First, and most important, this stream of research does not consider joint search problems.

Second, while it is intuitive that agents with greater or lesser ability would have more or less fine grained mental maps, the impact of agent ability is not modeled. In this work, agents' mental maps of the search space are exogenous – and do not evolve with search. In other words, while they model how cognition impacts learning, they do not consider how experience changes cognitive representations. Indeed, the modeling of how search (and communication) influences the emergence of mental maps – and how these in turn guide search – is where our approach is clearly distinct from prior "cognitive" models of organizational search (Gavetti and Levinthal, 2000; Postrel, 2002).

A second stream of work pioneered by Rivkin, Siggelkow and co-authors considers joint search problems, where they concentrate on understanding how the division of labor in search impact outcomes (Siggelkow and Levinthal, 2003; 2005; Rivkin and Siggelkow, 2003; 2007; 2006; Siggelkow and Rivkin, 2006). However, this stream does not elaborate on the role of cognition in search processes – and how prior representations of the search space guide actors' search effort. While other models of joint search go beyond the NK-formalism (Lounamaa and March, 1989; Puranam and Swamy, 2010), the evolution and alignment of mental representations in joint search has not yet been addressed.

This suggests a gap in our understanding of *joint* search by specialists in novel and complex landscapes where they *construct* mental representations of the solution space as part of their search as well as have a need to coordinate their search activities.

To make the rather abstract idea of joint search 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 with-stand wind power and adverse weather conditions. Good aerodynamics is critical to the efficient conversion of wind energy to electricity. 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 by education and training have little knowledge of each other's domains. How should we best integrate work across these specialists?

If they worked in isolation, it is likely that they end up with individually good, but jointly useless designs. If they communicated their designs to each other very infrequently, they may spend a lot of time exploring solutions that are jointly infeasible, but may come across unsuspected breakthroughs. If they communicate very frequently, they may swap enough knowledge to understand the each others' concerns and search in a jointly feasible area, but on the other hand they may mutually reinforce each other into inferior solutions. We call this attempt to integrate by harmonizing mental models "*cognitive integration*" in this paper.

Instead of trying to communicate their knowledge, the agents could swap designs and each agent could try to improve on the other's effort. This idea is similar to Page's (2007) conception of profiting from diverse knowledge representations. This could potentially help economize on transferring costly knowledge, but on the other hand may lead to an endless loop of swapping equally improbable designs. The actors could agree on limit criteria – for regions to search in – which could cut down on search costs, but inadvertently exclude valuable search regions. We call this attempt to search by building off of each other's efforts "*behavioral integration*" in this paper.

A different example may help starkly differentiate the two mechanisms. In customized application software development, the supplier firm and the client have to achieve some degree of shared understanding of the requirements and contours of the solution. This is done by communicating to each other the business problems and ideas on how to solve them with technology; i.e., by cognitive integration. On the other hand, the software vendor may face multiple technological choices that need to be coordinated across functions such as the GUI and database. These could be solved by authority or iterating on designs; i.e., behavioral integration.

This discussion suggests that predicting the efficiency of joint search is difficult because it depends on a number of criteria, including: (1) ability of agents to search in their own specialties vs. take a more generalist approach; (2) the precise nature of how integration efforts drive change in cognition and efforts at further exploration and (3) interdependence conditions or the ease with which innovations are found in the search landscape. Understanding joint search therefore requires us to carefully trace the relationships between agent ability and integration mechanisms to individual mental models as well as the level of shared mental representations and their subsequent impact on search outcomes. As in much analysis of dynamical systems, a computational model is the best method to trace these feedback-driven interacting relationships.

MODEL DESCRIPTION

To model joint search involving the formation and refinement of mental maps we need to define: 1) a task environment (search space), 2) the mental maps that agents hold of the task environment, 3) the capacity for refinement of mental maps, 4) the abilities of agents to search the task environment conditional on their mental maps, and 5) how the organization of joint search impacts the formation and refinement of agents' mental maps. The relation among these generic components of our model is illustrated in the upper portion of Figure 1, while the lower portion summarizes ideal types of task environment and agent ability. We successively describe how of each of these is implemented in the model – as a reference, Table 1 provides an overview of all parameters (details provided in the following).

FIGURE 1 AND TABLE 1

Task environment

In this paper, we are interested in understanding how the process of generating shared mental models affects specialists' ability to identify the most valuable combinations in the different landscapes. For this purpose, we consider an innovation that involves identifying the most valuable combination of two technologies – such as structural mechanics and aerodynamics in our windmill example. The innovative landscape is defined by a matrix, where each combination of the two technologies yields a payoff. In the baseline model, there are 64 possible decisions to be made in each dimension, leading to a landscape matrix defined by 64x64 payoff values[2]. Our landscapes differ from each other based on the number of peaks (few/many) and the similarity/dissimilarity in the height of peaks. Unlike the NK modeling platform, our technique allows us to generate custom landscapes. Figure 2 shows example landscapes in each of the four cells in our landscape design.

FIGURE 2

Agent's mental model of the landscape

The agent's mental model of the landscape consists of two elements: (1) decision choices available in each dimension and (2) payoffs associated with these choices. Our agents are

[2] The search space is therefore 64*64 =4096 combinations, compared to the 1024 values in the N=10 NK landscape.

boundedly rational and do not see all the 64x64 action choices available to them ex-ante, and as a direct consequence do not know the exact value of the associated payoffs.

At the beginning of the simulation, agents do not have fine-grained partitions of the search space (or do not have sharp vision) – they see only a very limited number of choices for each dimension. In the windmill example, the firm needs to choose a material to make the blades and the shape of the blades. The firm first perceives only a few choices in each dimension: a few candidate materials such as aluminum, steel, wood and composites and curved or straight for shape – i.e., a 4x2 simplified choice set instead of reality (the full 64x64 matrix). Each choice the agent sees is in reality a composite of several possible alternative choices (see Figure 4). For example, what the agent initially sees as wood could later be broken down into finer categories such as pine, oak, etc. The sharpness of the agents' initial vision – 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[3].

The agents' limited vision of the choice set also limits their understanding of the performance consequences of the choice set. The payoff the agent associates with each perceived choice is the average of payoffs for the "real" combinations hidden by that choice. For example, for the wood/curved combination – they see the average for all woods and all curved shapes. This treatment is similar to the payoff matrix seen by agents in Gavetti and Levinthal's (2000) model of simplified mental representations. As in Gavetti and Levinthal (2000), the payoff our agents perceive for any given coarse grained partition of the search space is the average payoff that can be achieved from all the latent alternatives that are present in that portion of the

[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 N. In our model, agents have 64 choices for each N.

landscape.[4] The imagery is one of blurred colors perceived from a high point of elevation. Figure 3 shows payoffs in reality and the initial representations perceived by the agent.

FIGURE 3 and FIGURE 4

Refinement of agents' initial mental representations

Agents' initial mental models get more refined with time – they generate more fine-grained knowledge partitions. In the above example, by investing effort the agent can distinguish between different types of woods. 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: choosing the portion of the landscape to further explore and exploring (gaining a sharper vision of) that region. The SWITCH operation chooses the region for further exploration and the DIG operation refines the agents' current mental model. We explain each below.

SWITCH: Switching captures the logic behind how agents' change the focus of their attention from one region in the landscape to another. Our agents are profit seeking and therefore switch to the most promising alternative, given their current vision (mental model). For instance, in Figure 4, the agent may choose to focus on the option Composite/Curved for further investigation instead of the option Aluminum/Curved and switch to it. 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.

[4] In our model, as in Gavetti and Levinthal (2000), agents have correct expectations regarding the attractiveness of each choice they see. This base-line assumption could be refined in future research by adding noise to expectations.
[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 footnote 4).

When agents switch to their chosen sub-space, they are placed in a random combination within that sub-space. Since agents do not have any knowledge of the specific combinations that makeup a sub-space, they have no control of their actual location within the chosen (coarse grained) partition of the search space. Therefore, they may experience a surprising deviation between the payoffs they actually receive and their initial estimate. In order to understand this deviation, the agents investigate further to increase their knowledge of that subspace (or DIG). This tendency to dig deeper and increase granularity highlights the role of surprise in shaping our – *always approximate* – beliefs about what will be best for us. This modeling feature, introduced by Cohen and Axelrod (1984), is a fairly underdeveloped intuition in the behavioral inventory.

DIG: This operation generates more fine-grained knowledge partitions (sharpens vision) of the landscape. The idea is similar to Simon's (1962) conception of refinement (expansion) of the search tree. For example, in the sciences, this refinement typically occurs from sequential increase in experimental control; the discovery of the neutrino is a classic example.

We model surprise-driven refinement of representations in the following way. When an agent decides to examine a cell more closely, a knowledge partition occurs; that cell splits in two at a random position in one of the dimensions. In the above example, if the agent decides to invest in understanding the cell Composite/Curved more minutely, the cell splits into carbon and glass fiber (as shown in Figure 5). Further dig attempts increasingly sharpen the agents' vision.[6]

FIGURE 5

----

[6] The relationship between switch and dig and the equations capturing the switch and dig operations are featured in the section on model dynamics and in Appendix 1.

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. Each dig operation splits a cell in two along a dimension (row or column) that is governed by the agents' inherent ability (as described below). The exact point where the split occurs is chosen at random (since the agent has no prior access to the latent dimensions of the search space).

Agent Abilities

Agents have heterogeneous abilities in refining their mental maps – they are differently able in their attempts at surfacing new knowledge in the two dimensions. These are characterized by two abilities: dig ability and cross-functional ability.

Agent Dig Ability (DA): DA governs the ability of the agent to create knowledge partitions (increase sharpness of vision) in their ***own*** dimension. Some agents are better able than others in acquiring new knowledge: they have superior skills in noting deviations from expectations and devising experiments that yield new knowledge. This is a parameter specified at the beginning and it does not change with time.

Agent Cross-functional Ability (CFA): CFA governs the ability of the agent to create knowledge partitions in the ***other*** dimension. When an agent decides to dig, the dimension in which the cell splits is governed by the ability of the agent to comprehend splits in the different dimensions. In the above example, one agent may be a specialist in material science, but may not understand the subtleties of blade shaping. When this agent digs, she is more likely to comprehend a more promising material than a more promising shape. A generalist agent comprehends new materials and shapes with equal ease.

CFA is specified as the probability with which each dig attempt by the agent results in a split in the *other* dimension. The larger this probability, the less specialized the agent is; when

CFA is 0.5, agents can equally discern alternatives in both dimensions. When CFA is zero, they are able to discern new alternatives only in their specialty. CFA is a parameter specified at the beginning and it does not change with time.

CFA is one of the novel features of our model, and enables an agent to "view the world" from the eyes of the other agent, albeit imperfectly. The higher this ability, the more often the agent tries to find an innovative solution by taking into account a more refined vision of the world both from their own as well as their partner's view point.

Agent Characterization: As shown in Figure 1, we investigate agents of 4 types – that differ in abilities along two knowledge dimensions (own versus other). The general characterization of the specialist is that they are very good at one thing, but not in any other – which in our model translates to high DA and low CFA. The generalist is exactly the opposite: is not very good in any one dimension, but can easily comprehend several dimensions – i.e., has low DA but high CFA. The "genius" agent has both high DA and high CFA, while the "dumb" agent has low ability in both. While we are particularly concerned here with the behavior of the specialist agent, we also include some discussion of the other type of agents. Agent characteristics are described in Table 2.[7]

TABLE 2

Organizing joint search – cognitive or behavioral integration of mental models

Organization for us is a means of aligning the actions of the two actors. In single search – one agent searches both dimensions, and therefore automatically has fully aligned mental maps,

[7] The actual parameter values that we use to characterize the four distinct profiles of agent abilities (i.e. specialist, generalist, genius and dumb agent) were chosen on the basis of numerous numerical analyses and sensitivity tests. These tests ensured that the capacity to develop deep knowledge in the agent's own domain is very high when DA is set to 10 while it is very low when this parameter is set to 1.

as well as aligned actions. In joint search, however, there is division of labor – one agent concentrates on the row dimension (structural mechanics), and the other concentrates on the column dimension (aerodynamics). Different organization designs engender different patterns of interaction among the specialists and therefore determine the rate and the level of mental model alignment and action alignment over time. This, in turn, affects agents' subsequent search strategies. Figure 6 shows a schematic of the different organization designs we consider. Below, we explain how we model different patterns of integration.

FIGURE 6

No Integration: In this case, the payoff to the agents is determined by agents' position in the landscape, where the row-agent specifies the row location and the column agent specifies the column location. In our example, the structural specialist decides which material to experiment with in the next prototype (say move from wood to composite), while the aerodynamics specialist determines which shape to consider in the next round (curved or straight). Both agents are placed in the cell that corresponds to their combined preference.

The key to this design is that agents do not make any attempt to align their mental models. This does not mean that the agents ignore the other dimension. Each agent may (depending on their CFA) try to understand the choices available in the other dimension and use such refined mental maps to make choices in their dimension that will potentially also meet the feasibility criteria of the other agent. This is the "no integration" condition because the agent's mental representations are not aligned with each other (i.e., they do not have the same partitions of the landscape).

Mental maps of two actors can be more or less aligned. Perfect alignment is the condition where two actors' knowledge partitions exactly overlap in all relevant dimensions; in our setup, they have identical vision regarding possible choices of materials and shapes. The level of (imperfect) alignment in the actors' mental maps is the intersection of the two actors' knowledge partitions.[8] The immediate consequence of misaligned mental maps is that the probability of coordination failure increases as a function of incongruence between two actors' mental maps.

Cognitive Integration: 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. This is called cognitive integration because each agent makes an effort to understand the world "precisely" as viewed by the other agent. Cognitive integration occurs when agents attempt to partition the search space in *identical* divisions by communicating their knowledge to each other.

Cognitive Integration 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 agent with one new column (row) partition that the row (column) agent does not already know. The more frequent these requests, the more aligned the knowledge partitions become[9]. The frequency of communication is a parameter in the model and does not change with time. Similar to the previous case, here too the row agent determines the row position and the column agent determines the column position.

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 mental maps -- *differences* between the granularities of two actor's knowledge partitions

[8] See Samuelson (2004) for a comprehensive introduction to the notion of knowledge partitions (and a review of the literature). The general idea is that degree of alignment can be measured as the ratio of the intersection of the agents' knowledge partitions to the union of their knowledge partitions.
[9] There is an (opportunity) cost to communication. Each time period the row agent request the column agent for information, the row agent forgoes the opportunity to search in that time period.

are the only source of misalignment. Though agents in the "no integration" case may have very refined mental maps in both dimensions, these are likely not aligned because each agent independently generates her own mental map by random refinements of the search space. With infrequent communication, cognitive integration approaches the no integration case, while it approaches the opposite, perfect alignment with increasing communication frequency[10].

Behavioral Integration: Instead of attempting to share their mental maps, agents may coordinate by directly constraining their actions in searching for a solution. In this case, unlike the previous cases, each agent searches for a "complete" solution. Each agent's position on both dimensions is determined by them alone, for which they receive a payoff. Agent's are made aware of the position and payoff of the other agent. The agent with the lower payoff shifts to the same position in the landscape as the agent with the higher payoff. Behavioral integration occurs here because the agents' actions – their attempts at locating a favorable subspace and increase their sharpness of vision – are determined by accepting the lead of the high payoff agent.

We could think of this arrangement as a firm that has two teams charged with finding an innovation in a particular area, with the teams approaching the task from different perspectives. Both teams come up with their recommendations regarding plausible promising alternatives, and the firm chooses the one with the highest promise for further investment in understanding better (recall our example of a team trying to design a GUI and database).

Model Dynamics:

Figure 7 shows the generic steps in our model. At t=0, the agents are randomly positioned in the search space and they are endowed with an initial vision of the possible alternatives. First,

[10] We only model perfect communication. The degree to which mental maps overlap is strictly governed by frequency of communication. Since we model the no integration case, there is no logical reason to also model imperfect communication. We have also explicitly modeled sequential search, and as expected, the results are nearly identical to the cognitive integration condition with very high communication.

the agent perceives possible alternatives for action and chooses one that seems best suited to achieve high payoffs (i.e., chooses one of the known branches in Simon (1962)). The probability that an agent switches to any given cell from their current position is determined by the payoff the agent expects from that cell compared to all other cells that they can see, and is determined by a Softmax algorithm[11].

FIGURE 7

Second, having chosen a promising sub-space, the agent samples a specific combination from that space. Since each choice alternative the agent is aware of contains multiple latent combinations of the two technologies, sampling is achieved by placing the agent at random in one of these latent combinations. If the agent chooses *not* to switch from their current cell, they are repositioned at random in one of the combinations within their current cell; i.e., they sample again from their current choice set. [12]

Third, the agent compares received payoff, which is the value of the specific combination sampled, with their expected payoff for this sub-space. A notable difference between expected and received payoffs gives rise to surprise, which in turn may stimulate a more minute investigation of that sub-space (DIG operation). Simon (1962) refers to this process as refinement of the search tree or choice set expansion. In the real world, the dig operation involves things like thinking about the problem, reading about it, etc., which typically yield new knowledge, or a new partition of that sub-space. The agent decides to dig based on the magnitude

[11] Appendix 1 provides further (mathematical) details on the mechanics of the model.

[12] In other words, as long as a cell is larger than 1 x 1, staying in the same cell does not mean they are doing nothing. It means they are unwittingly experimenting with new combinations without becoming aware of qualitative difference among these different combinations. Also see appendix 1.

of difference between expected and received payoffs, and is determined by the agent's dig ability (DA). Higher DA agents dig further even when this difference is very small. Low DA agents do not dig until this difference is fairly large (see appendix 1 for implementation details).

In the fourth step in the model, the agent recalculates expected payoffs for all known choices. If the dig operation in step 3 results in new partitions, the agent now has more choices available. After calculating these new payoffs, the agent retraces the previous steps from step 1.

The simulation ends when agents find the best combination in the landscape or after 500 discrete time steps. The value of 500 time steps was chosen because it ensures that all simulations had approached a steady state.[13] The agents stop searching when they (rightly or wrongly) believe that they have sufficient understanding of the search space.

Effectively, agents stop searching when they no longer switch or dig. They stop switching when they firmly believe they are located in the most attractive sub-space. They stop digging when their expected payoff more or less equals their actual received payoff. This stopping rule captures search in the context of scientific discovery where it is unclear if continued search will uncover more valuable alternatives, even if it leads to more fine-grained vision of the problem-space.[14] This stopping rule is closely related to March's (1988) classical model where agents with adaptive aspirations gradually learn to expect what they achieve.

---

[13] 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.

[14] We have a "scientific" notion of search, devoid of commercialization interests. For example, engineers who stumble upon a more valuable combination than they expect continue to try and understand the source of the anomaly. They may choose to commercialize this innovation, but continue search, possibly in an attempt to find an even better solution for the next release. Similarly, they may stop searching without identifying any valuable combination in the landscape, because they believe that what appeared to be the most promising region, on further examination, is only as good as or worse than their current solution (the initial combination they started out with). For example, they may conclude that the desired combination is not feasible using existing technology, based on their (imperfect) understanding of the search space. The history of innovation is full of examples where one inventor stopped searching for an innovation (such as phonograph, electric bulb, telephone, vaccines) believing it to be technologically infeasible at that time but another inventor with a different knowledge partition achieved it.

RESULTS

In this section we present results for how organization, agent ability and landscape properties interact to determine search outcomes. Before jumping into the results, we present thoughts on some baseline conditions to help place the results in context.

First, consider the case where each agent had perfect granularity ex-ante that enabled them to see all the possible choices available and their payoff consequences. In this case, each agent would immediately pick the optimal choice for their problem and search is trivial.[15]

Second, consider the case where the agents have constant but less than perfect granularity. In this case, based on their perception of payoff, they choose an option and tinker with it. If agents jointly choose the region that contains the jackpot and continue tinkering within it (i.e. hitting random points within this cell), they are likely to find the optimal combination by chance. The likelihood that they find the correct combination increases with time and decreases with the size of the region they are experimenting in. For example, if both agents have fully aligned mental maps and are located in a region that they perceive as 8 x 8, the probability of hitting the jackpot would only be 0.00024. Any integration mechanism works by making one agent aware of new knowledge that the other agent already knows – in other words, by increasing the agent's granularity. If this is prohibited, shared mental models cannot be created, and therefore coordination is not possible. Note that this reasoning also applies to the case where each agent has perfect granularity in their own dimension and imperfect granularity in the other dimension (as long as the landscape has several peaks).

Third, consider a secular, but random increase in granularity that happens outside the control of the agent. If this leads to perfect granularity, it reduces to the first case. However, all

[15] Note that coordination may still be a problem in matching games – where there are multiple equally attractive solutions. This problem is fairly well-studied (Schelling, 1960) and is not the focus of our efforts in this paper.

models in management assume that the world is complex enough and time short enough for this outcome never to be achieved. Therefore, in all cases of interest, this reduces to the second case. If we allow for the possibility that integration mechanisms can play a useful role (i.e., new knowledge surfaces randomly, but once it does, it can be communicated to the other agent), then over time this reduces to the second case where both agents search in a smaller region of the landscape progressively. The fruitfulness of this effort completely depends on the speed with which new knowledge occurs in both dimensions.

This leads us to the case where the generation of new knowledge is directed by the agent. In this case, we need to both consider the locus of generation of new knowledge and how it is impacted by the efforts of the interdependent agents to build shared mental models. This already suggests that integration mechanisms is going to significantly impact search – since it refines pre-existing mental maps and therefore directs subsequent search. Since we cannot ex-ante predict the nature of generated knowledge or search outcomes, we present simulation results for this case, summarized in Table 3. These results are averages of 300 individual histories where 2 agents are engaged in joint search over 500 discrete time steps.[16]

TABLE 3

We primarily wish to highlight three results. First, that some integration mechanism is necessary for successful search as the number of peaks increases. Second, quick integration of mental models results in ineffective search outcomes. Third, agent ability has little impact on effectiveness of joint search.

Result 1: Integration is usually necessary in multi-peak landscapes:

[16] Averaging over 300 individual histories ensures that all differences across conditions are statistically significant.

Our first result suggests that, contrary to naïve intuition, the larger the number of valuable combinations, the more important is integration. The final payoffs shown in table 3 are average payoff at the end of the simulation for the 300 agents and measures how effective they are in search. Looking at the results in table 3, we see that the no integration condition identifies the best combination in single peak landscapes unfailingly. However, its performance deteriorates in the two-peak landscape, and is nearly zero in the multi-peak landscapes.

This is because in single-peak landscapes, the peak acts as an attractor for each specialist as they refine their mental maps in their own dimension. When each specialist has located the best action choice from their view point, they have simultaneously landed on the best possible combination, without explicitly needing to coordinate. This result is intuitive for peaks with a wide basin of attraction – it is surprising that it also holds for this spike landscape. It suggests that the gradual refinement of mental models serves as a substitute for traditional hill-climbing models.[17] Surprisingly, the logic suggests that on a smooth-hill landscape with a basin of attraction, agents will not reach the absolute peak. This is because agents mis-coordinate near the top where the different combinations have very similar payoffs (more on this below).

However, when another peak is added to the landscape, the no-integration agents perform much worse. This is because the smaller peak (which is half as valuable as the larger peak) confuses some of the agents into choosing it. From table 3 under the column heading "2 spikes", we see that uncoordinated search by specialists on average only leads to a performance of 0.70 while any form of coordination raises the average payoff to 0.90 or higher.

Why does just adding a single peak confuse agents so much? Figure 8a shows the mean footprints of how the specialist agents converged on a solution. These mean footprints are fairly

[17] This result is stunning from a game-theoretic point of view. Common knowledge is unnecessary (or irrelevant) when the structure of the game reveals itself in the course of play.

intuitive in that they represent averages over all locations that the joint searchers have visited throughout the simulation. Looking at the footprints in Figure 8, we note that coordination on the small peak is much lower than mis-coordination on the larger peak. i.e., by chance, a small fraction of the agents end up coordinating on the small peak, based on the successive partitions in their mental maps. For a larger fraction, one agent tries to choose the large peak, while the other chooses the small peak, and jointly they end up with low payoff[18]. They continue in this state because each chooses a valuable region from their own view-point and continue to search further in it since they do not receive the expected payoff. Even small communication attempts provide enough alignment in mental models for the agents to be pushed off this cycle and eventually coordinate.

We may expect that as the specialist agent acquires a higher level of knowledge partition in the other dimension (as they do eventually)[19], this problem should go away. However, in reality, it actually becomes worse. The specialist, having speedily identified the best partition from their own view point (and at least one of the two agents believes the small peak is the most valuable region), they attempt to increase their granularity of vision in the other dimension. Since they continue to act independently, their locus of search remains the individually attractive, but jointly low-value region. With an increasingly fine-grained partition of that region from both dimensions, they (accurately) expect and receive a low payoff. Since they believe that the most valuable region (in their viewpoint) has zero value, they essentially stop searching.

Intuitively, this suggests that if the agents had high CFA they would initially not get trapped in the zero value region and therefore perform better. Actually, from table 3 we see that

[18] If the agents have full initial granularity in their own dimension and none in the other dimension, this mis-coordination would not happen. This example, powerfully illustrates the problem of limited knowledge of available choice sets, which has been assumed away in prior work.

[19] Note that our specialist agents have a very small, but non-zero, ability to identify knowledge partitions from the other dimension. Making this zero will not change our results qualitatively.

the two "genius" agents acting without coordinating their mental maps perform worst of all[20]. From the mean footprints in figure 8b, we see that the genius agents are more likely to fall in positions very close to the large peak but not the peak itself (seen from the yellow markings near the red in figure 8b).

FIGURE 8

This is because a "genius" agent is trying to solve two problems simultaneously, and their way of operating actually makes it more difficult to solve. Each "genius" quickly locates a possible solution in both dimensions. Unfortunately, this exposes both agents to a larger risk of ending up with a useless solution. Unless they both get the location of the superior solution exactly right in both dimensions, and unless this happens simultaneously, they end up with a useless solution. On average, the two "geniuses" mis-coordinate in 37% of the cases.

Now, we ask the question, what is likely to happen if we further increase the number of peaks – connecting the two peaks as shown in the ridge landscape in figure 2d? Does it become easier or more difficult for joint searchers to find a valuable combination? The intuition above that the agents get confused suggests that it becomes more difficult to find even a single valuable combination, even though there is a significant increase in the number of valuable targets. A quick look at table 3 (under the "Ridge" column) shows the abysmal performance of search without any effort at aligning mental models –on average agents achieve a final payoff 0.03-0.04, which is about the same as, the randomizing payoff expected in this matching game (1/32). Similar results are observed for the ridge-peak landscape (shown in figure 2d). Figure 9 shows the respective agent footprints.

[20] In contrast, a single "genius" agent choosing both the row and column dimension performs very well.

FIGURE 9

However, our conclusion about the need for shared mental representations in multi-peak landscapes is tempered by considering the position of the multiple peaks (or level of interdependence). Again, consider the ridge landscape, but now imagine that the ridge is rotated 45 degrees from its current position and lies parallel to one of the axes. In this case coordination and shared mental representations are unnecessary – each actor may choose along their own dimension and still achieve good outcomes, similar to the single peak(results available from authors). This reflects the fact that the position of the peaks in the landscape dictates the degree of interdependence between the agents. The higher the interdependence, the greater is the need for shared mental representations of the search space among the joint searchers. This presents the rather intriguing idea that difficult joint search problems can be made easy by changing the way it is presented to the agents; a designer can design interdependence instead of integration.

A final take-away from table 3 is that in multi-peak landscapes, different organization forms (different levels of cognitive integration and behavioral integration) lead to different outcomes, which leads us to our next finding.

Result 2: Fast integration of mental models leads to inferior search outcomes:

The typical understanding of coordination and shared mental representations is that you cannot have "too much". Discussion of "too much coordination", if any, suggests that it may be expensive, i.e. it is *inefficient*. We however find that too great alignment of mental models in early phases of joint search may reduce the *effectiveness* of search.

We present the data first. Looking at table 3, we find that for the two spikes landscape, communicating once in twenty rounds has a payoff of 0.96 while communicating once in two rounds has a payoff of 0.90; low levels of cognitive integration is better. A similar pattern of results is found for the ridge-peak landscape[21]. However, for the ridge landscape, we find the reverse – agents with high levels of cognitive integration perform better. Why do we observe these differences?

There are two elements to successful joint search: "scanning" for superior combinations and "targeting" of search efforts by the two agents. The key to success in the first element is having a sharper vision of the entire landscape. The key to success in the second element is "coordination" – focusing on identical search regions; this is facilitated by integration. However, integration by selectively increasing knowledge partitions of some choices influences the locus of sharper vision. 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. The tension between these two elements of joint search gives rise to the observed pattern of results.

In the two peak landscape – exhaustive scanning is more important – and therefore we see that heavy efforts at integration lead to inferior solutions. In contrast, in the ridge landscape, since all the peaks are of equal height, alignment is more important – and therefore greater efforts at integrating mental models leads to superior outcomes. The ridge-peak landscape lies between these two extremes, and therefore, medium effort at mental model integration proves to be most effective. This result indicates a fundamental tradeoff between decisions to scan a task environment for promising solutions and decisions to take aim, and target a promising solution.

[21] The difference in final payoffs is statistically significant.

Pre-mature focus is a powerful detriment to search – it prevents agents from exploring the high value region, which is a natural attractor, as we discussed in the single spike case. Figure 10 shows the search pattern of specialist agents in the two spike landscape – the left column for agents who communicate infrequently (1 in 20 rounds); the right column for agents who communicate once every two rounds. As the granularity figures show, in the low communication case, the agents quickly identify an interesting region in their own space and then slowly try to understand the space from the other point of view. This enables them to scan effectively – their probability of digging is above zero even after 250 time steps. The high communication condition on the other hand stifles scanning – agents converge very quickly – the probability of digging falls to almost zero before 50 time steps.

FIGURE 10

Behavioral vs. Cognitive Integration: In general, across all landscapes, behavioral integration is more effective than cognitive integration. The opposing effects of scanning vs. targeting also explain the consistent superior performance of behavioral integration over cognitive integration. In behavioral integration, since each agent acts independently in choosing their location but merely guides the other agent to the high value region, there is no precise alignment of mental maps as in cognitive integration. This allows the agents perform a more exhaustive search of the landscape without premature lock-in to a suboptimal target. Once better combinations are identified by either agent, this integration mechanism ensures that the other agents will also improve their vision of that region, allowing the agents to target specific combinations. However, this effort is also more wasteful, since in landscapes with few peaks there is lots of wasted exploration.

In this condition, contrary to our expectation, genius agents perform only as well as the specialists. We expect the genius agents to perform better because their high CFA enables them to partition the knowledge space from both dimensions and therefore find better solutions. This leads us to a discussion of the relative merits of the different kinds of agents.

Result 3: Genius agents do not add much in joint search.

In a single searcher scenario, it is very intuitive that a "genius" agent – an agent with search skills in both domains (commonly referred to as T-shaped skills) – will outperform agents with skills in only one dimension. Somewhat surprisingly, our results suggest that two specialist agents are as effective as or better than two genius agents in all types of landscapes. In all landscapes, uncoordinated search by genius agents yields worse performance than search by specialists.[22] However, it is surprising that two genius agents with behavioral integration do not outperform other agents.

This result provides a useful note on the potential drawbacks of employing "genius" agents, i.e. agents with T-shaped skills. The successful employment of such agents requires that they work alone or that their efforts are closely aligned in joint search. If not, their superior skills will often lead them to joint solutions that are useless. Ironically, this happens because the power of their cross functional ability expands the space of possible solutions so that it becomes harder to coordinate. The more the peaks and the larger the area they are spread in, the more difficult

[22] In the cognitive integration condition we do not specifically model "genius" agents, because we can predict their behavior based on the other models. When "genius" agents communicate, they align mental maps, but can also independently achieve more fine-grained view of the other dimension. This leads to two effects: lack of better understanding of their own dimension and confusing movement across the landscape on the basis of mis-aligned representations as we see in the no integration condition – both of which are likely to lead to their poor performance.

their problem becomes (for example, compare the 0.63 payoff in the 2-peak condition to the 0.03 payoff in the ridge-peak condition).

In behavioral integration, the high CFA of the genius agent allows them to identify better solutions and lead the other agent to it. The column "total payoff" in table 3 is a measure of how fast the agents found the superior solution. Inspecting this column for the Genius vs. specialist agent, we do see that the genius agent achieves outcomes faster under behavioral integration. However, the slow experimentation of the specialist agent is as effective.[23]

It is ironic that high CFA is more valuable in behavioral integration than in cognitive integration. We find that high CFA is useful mainly in circumstances when agents are *not* dependent on each other to fulfill their knowledge voids – precisely the situation where we intuitively expect high CFA to be useless. The purpose of cognitive integration is to educate the specialist about the other dimension, while the purpose of behavioral integration is to not do so, and yet, CFA is more valuable in the latter. The explanation is that cognitive integration emphasizes targeting over scanning, and leads to premature agreement about inferior solutions, an instance of "groupthink" where both agents see "promise" in exactly the same useless alternative. In all task environments that put a premium on striking a balance between the two by expanding the target area (2 spikes and ridge-peak), behavioral integration is superior.

Interestingly, the generalist and the dumb agent also achieve equally good solutions as the specialist (albeit more slowly).[24] The insight is that the design of coordination mechanism for a team of innovators is more important than ability in joint search problems. Our model helps

---

[23] If the specialist agents had a CFA =0, they would rarely achieve superior outcomes under behavioral integration. This is because each would propose nonsense solutions (except by luck) since they have no way to refine their choices in the other dimension. These results are available from the authors.

[24] The only difference between the specialist and dumb agent is that the latter has low DA and therefore does not dig until the difference between expected and observed payoffs is large enough. However, once they accidentally dig, the nature of our "spiky" landscape amplifies any differences and induces them to dig even more. The results may be different across these agents (and more striking) in more rounded landscapes with very low gradients.

identify distinct classes of coordination mechanisms and assess their usefulness contingent on employee skill and problem type.

Are there any situations when high DA (a specialist) is a liability? The only case where the dumb agent significantly outperforms the specialist agent is in the ridge landscape with high communication. In table 3, under the ridge column, the average payoff for a high communication specialist agent is 0.84, while the corresponding payoff for the dumb agent is 0.92. When both agents have high DA, but lack a natural attractor, they focus faster than they can align their mental maps – leading to joint investigation of a barren part of the landscape. We note that a similar pattern does not occur for the ridge-peak landscape or for the two-spike landscape, where a joint natural attractor exists. This suggests that low DA is valuable when precise alignment of mental models is paramount. We can test this intuition by reducing DA to intermediate levels while maintaining high cognitive integration to achieve even better performance in the two peaks landscape – we tested this and the intuition holds.

This gives us an intuition regarding when a high DA agent (or a specialist) is a liability: This happens when the specialist has a rather fine-grained mental map of a region of the landscape that has low value combinations but only has a very coarse understanding of the regions with the high value combinations (e.g. a pedantic expert that draws on outdated knowledge). In this case, the specialist will concentrate on refining her knowledge of the bad patch – "knowing more and more about trivialities" – while completely ignoring the high value patch, since she does not believe it is even worth cursory exploration. An ex-ante "dumb" agent may perform much better in this case because the lack of knowledge leads them to explore more exhaustively and therefore more likely to come across the jackpot.

DISCUSSION

In this paper we try to understand the effects of organization on joint search for innovations, specifically by understanding how organization influences the development of agents' cognitive representations and thereby impacts search behavior. Three results stand out: first, in most joint search problems, shared mental maps are very important – but surprisingly low levels of sharing is adequate for effective search. Second, for successful joint search there is a need to balance two opposing forces: exhaustive scanning of the landscape and jointly targeting specific valuable combinations. Third, organization is more important than agent ability in joint search efforts.

It is unsurprising that aligning mental maps is important in coordinating interdependent tasks. What is surprising is how little alignment is required in many highly interdependent tasks. This suggests that some of the difficult specialization-coordination problems that firms face are either exceptionally hard problems (Postrel, 2002 conjecture), or they use incorrect organizational structures to coordinate.

While our first result emphasizes a minimal required level of coordination effort in joint search problems, the second result points to potential dysfunctional effects of coordination. Specifically, we find that too much integration (especially at the early stages of joint search) can be harmful. This occurs because joint search needs to balance scanning and targeting. Too much integration upfront reduces scanning of potential targets and is therefore less effective in identifying breakthrough innovations. This result is generated by the feedback loop where joint searchers effectively constrain each other's consideration of the search space. It follows that if the aim is to find the next incremental innovation, fairly constrained and targeted search with thick integration mechanisms should be the preferred organization design. By contrast, organizing to find the next radical innovation requires that the agents search more broadly and are less constrained by the needs for coordination. In other words, coordination restricts further specialization, and the latter is very important in identifying breakthroughs.

While this result challenges the view that more coordination is always better, it is broadly consistent with many formal and empirical studies in the "exploration-exploitation" framework (March, 1991; Fang and Levinthal, 2009; Argote and Greve, 2007; Siggelkow and Rivkin, 2006). Especially, our work is in line with prior work that considers structural solutions to this trade-off; we also find that too much integration favors exploitation over exploration (Benner and Tushman, 2003; O'Reilly and Tushman, 2004; Ethiraj and Levinthal, 2004; Fang, Lee and Schilling, 2009). Our finding that search for radical innovations requires that functional specialists operate in mutual ignorance of one another's knowledge domains resonates well with prior empirical and theoretical work (Fleming, 2007; Postrel, 2002).

Fleming (2007) synthesizes a large body of research concerning the genesis of break-through innovations. He finds that breakthroughs are more likely to happen in brokered rather than dense networks, which reduces targeting in our terminology. Interestingly, he also finds that breakthroughs come when experts with deep knowledge of their fields collaborate, but not in collaborations involving "shallow experts" (Fleming, 2004). Our model clearly explains the intuition for this result. A deep expert already has a very high granularity of her own dimension – and when such experts collaborate they jointly enhance the scanning of the search space. This allows them to identify truly valuable peaks. This is not true of shallow-experts, who are both trying to improve knowledge in their own dimension, while simultaneously trying to improve understanding of the other dimension. They get trapped more quickly in less productive regions.

The graphs of search behavior and attempts to improve granularity by our stylized agents in Figure 10 bear this signature.

Turning to prior theoretical work, Postrel's (2002) analysis is related to ours. He showed that mutual ignorance across specialties is optimal unless there are significant interactions among knowledge domains. It is reassuring that we are able to reproduce this result from Postrel's (2002) analysis even when we use a dynamic model that explicitly captures the emergence of agents' mental maps. Our model also extends Postrel's (2002) model in a number of ways. He did not provide an analysis of the relative benefits of alternative integration mechanisms and explicitly excluded the analysis of shifting emphasis in the locus of problem solving.

We show that the choice of integration mechanism is of critical importance for any joint search problem that occurs in a task environment with multiple interdependent peaks. Behavioral versus cognitive integration are distinct mechanisms that can both lead to successful joint search, but when should one be favored over the other?

Our results establish a natural hierarchy of mechanisms for the organization of joint search (see Figure 6). In a search space where a single alternative outshines any other (the single peak landscape), no integration is necessary. In contrast, if the search space includes even a few interdependent peaks (as in our two spike landscape), some level of integration is necessary. While behavioral integration has an advantage over cognitive integration also in the two spike landscape, it is very slight. Both of these integration mechanisms can be tuned to condition the joint search process so that scanning becomes sufficiently exhaustive while not overriding the ability to target the highest peak. However, successful application of behavioral integration requires rapid and complete information transfer among the agents regarding the actions they tried out and their performance outcomes. In comparison, successful application of cognitive integration is less demanding since it relies on fairly infrequent (and realistic) levels of cross-functional updating.

In landscapes with many peaks, behavioral integration is superior. Since behavioral integration presumes that one specialist takes the lead, it avoids trapping joint searchers in mutual confusion. While cognitive integration can be tuned to perform fairly well even in this task environment, it is inferior to behavioral integration. Interestingly, the best configuration of cognitive integration for this challenge is the employment of agents with low dig ability that are communicating very frequently. This is because greater efforts at integrating mental models lead to superior outcomes.

Given these advantages, we would expect pervasive efforts at behavioral integration and rather sparse use of cognitive integration. However, prior empirical work emphasizes cognitive integration over behavioral integration. What are the costs to behavioral integration? First, the number of experiments tried out is much larger. When experimentation is very expensive – such as testing blade design in wind tunnels for 100000 hours, or performing clinical trials – behavioral integration is no longer attractive. Second, this presumes a very high degree of knowledge transfer across the agents; when that condition does not hold – for reasons of tacitness or secrecy among others, this mechanism is unlikely to work. This mechanism is likely to demand leadership to ensure adequate information transfer among teams – who may well think of themselves as competitors. Trust in the efforts of the other partner and willingness to follow their lead even in territory that is unknown or suspected to be useless is necessary for behavioral integration to work. In this respect, our result is consistent with the findings of Selten and Warglien (2007), who concluded that role asymmetry, is one of the most important features in the emergence of language in coordination games. Some interesting examples of behavioral integration at work may include innovation races, where teams approach the problem from different directions and learn from each other.

Our third result also has fairly interesting implications on choices regarding agent upgrading vs. organization design. While cross-functional knowledge is paramount in landscapes with multiple interdependent peaks – a situation that corresponds to joint search for incremental innovations – one of the surprising findings from our model is how little cross-functional knowledge generally matters to finding innovations. In our stylized model, communication is perfect – and one can argue that in realistic applications, some cross-functional knowledge is necessary for communication to play the role that we have envisaged in our model. Our emphasis is different – we suggest that it is hardly worthwhile to find two people with a PhD in both Physics and Biology to enable joint work in the two sciences. Someone with a PhD in one and a passing understanding of the other may well be enough. There are few innovation instances and specific organization designs where agents with T-shaped skills are genuinely useful. We suggest that while a small amount of cross-understanding is necessary in most interesting situations, it can go a surprisingly long way even in difficult coordination problems. While higher levels of cross-understanding may be beneficial, the process of achieving these high levels may actually hinder superior performance in the joint task. In the latter respect, our results are similar to those of Kretschmer and Puranam (2008), who point out that greater cross-understanding may come at the cost of greater self-understanding.

Our results that joint search for innovations requires fairly low levels of integration of mental models, while too much integration often has adverse consequences also help explain some puzzles from prior empirical work. Several studies have found that there are very few boundary spanners even in large projects (Tushman and Scanlan, 1981; Tushman, 1978; Katz and Tushman, 1983), raising the question of how so few people can meet the communication needs of large projects. This is rather puzzling when considering that most managers routinely include “more communication” as one of the requirements of successful product development. Further compounding this puzzle, other studies have found that project performance decreases as the number of boundary spanners is increased (Tushman, 1977; Tushman and Katz, 1980). The need for low levels of shared knowledge and the adverse impact of too much integration explain both of these puzzling findings. Our results are also consistent with recent findings in NPD that shifting lead roles in making decisions at different stages of the process is beneficial to performance (Song, Thieme and Xie, 1998).

The results also indicate a subtle difference between single searcher models that draw on the exploration-exploitation framework and the joint searcher model. In both cases, adequate search of the landscape is important to achieve higher performance. However, in the single searcher models, exploration is typically tantamount to non-greediness; i.e., agents are willing to forego immediate profits to examine what at first glance are unattractive choices. In a joint searcher model – this is never the case. Introducing even very small doses of randomness in agent switching behavior is highly detrimental to joint search. This happens because the agents in their exploration mode confuse the other agent about the consequences of their decisions, the classic confounding effect that other papers in joint search have highlighted (Lounamaa and March, 1991; Puranam and Swamy, 2010).

Limitations: This work is subject to the following limitations. First, the landscape of innovation is exogenous to the model. Since organization 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 work on analogical reasoning and experience (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. Our preliminary tests indicate that our results are robust to combining agents with different abilities in the two dimensions as well as different initial granularities. 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. As a final limitation, we limit our analysis to joint search involving two specialists. There is no reason to believe that including more specialists would qualitatively alter the results, though the analysis would be more complicated when considering higher dimensional search spaces.

Contributions: Despite these limitations, we believe our contribution has some unique strengths. 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. We provide a unique method to model (changing) mental representation and their impact on future search behavior. By considering both organization and cognition we are able to understand how to organize joint search problems in a variety of innovation landscapes. We also provide an alternative modeling platform that overcomes some of the limitations of the NK formalism that is important in some search problems. In this platform it is possible to understand the consequences of different kinds of common knowledge assumptions – such as that of actions, payoffs and rationality. Our findings suggest a contingent theory of organizing for innovation – depending on the type of innovation sought, characteristics of the landscape and agent abilities. These suggest that joint search should be studied in its own right since single searcher models offer limited guidance on organization issues.

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 completely unexplored and is likely common in organizational situations. Some of our findings on the need for shared knowledge are counter to the traditional game-theoretic analysis. Future work in this area could be very interesting.

CONCLUSIONS

We have introduced an analytical framework that can help unravel how functional specialists who have little mutual understanding of each other's knowledge domains are able to engage in successful joint search. Specifically, we show how attempts at integrating mental maps can aid/impede search contingent on the nature of the task environment and agent ability. Our model is very simple and quite easy to replicate. It includes an intuitive representation of the search space, a simple characterization of agent abilities that allows definition of meaningful ideal types, and an easy way to define integration mechanisms. We hope that our analyses have illustrated the usefulness of our approach, by extracting results that are compatible with prior work – and by extending prior work to consider the role of integration mechanisms in developing adequate mutual understanding to solve joint search problems.

The aim of this paper has been to examine joint search for innovations. In and of itself, the consideration of richer forms of mental representations fills a void in the literature on organizational search and learning. But even more important is our finding that the nature of joint search is very different from individual search. Most of our formal models of "organizational search" are notably non-organizational (Knudsen and Levinthal, 2007) and in effect ignore the essential features of joint search. While this is probably justified when top management teams (a group of generalists) search for new policies (a common topic in the literature on organizational search), we point to very different conclusions when distinct functional specialists engage in joint problem solving, as is typical for new product development and innovation teams. We are aware that our treatment is suggestive rather than exhaustive, but hope that we have demonstrated why it is important to further our understanding of the relationship between organization and joint search.

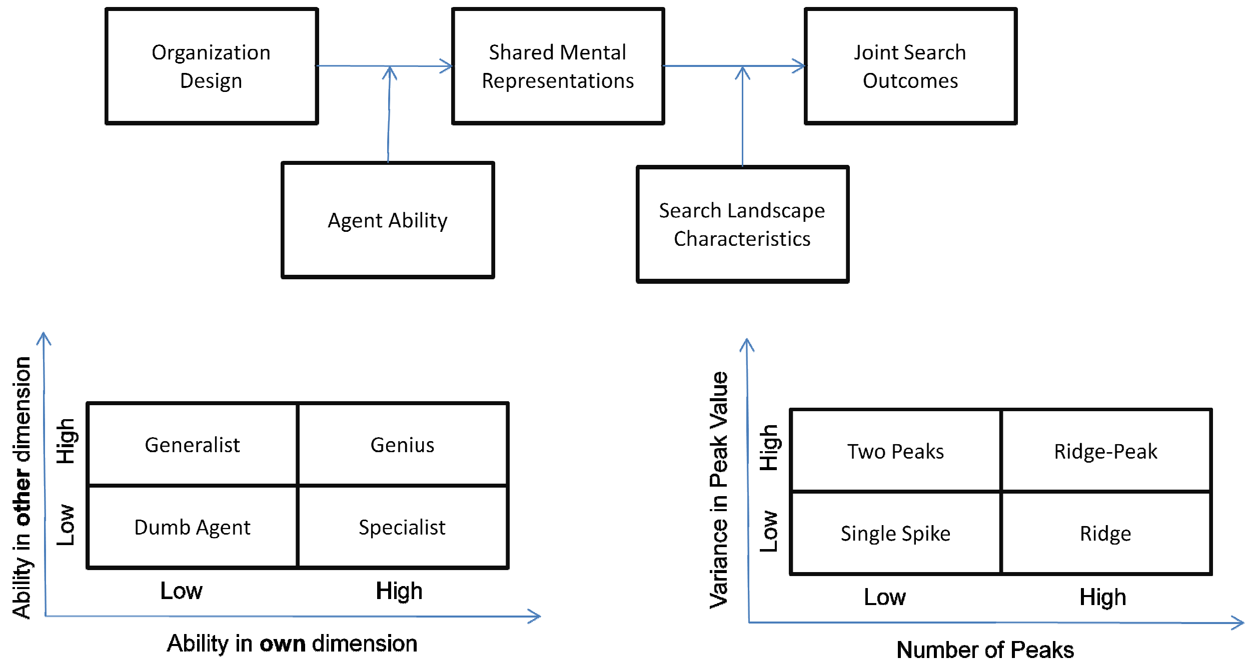

Figure 1: Exploring the relation between organization and cognition in joint search

Table 1: Table of parameters

Parameter	Range	Purpose
Search landscape	64 x 64	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	1 x 1	The number of initial choices the agent sees in each domain and can vary between 1 and 64 in each domain. **
Switching	τ: 0 – 0.1	The ability of the agent to move from a less attractive to a more attractive choice. The movement is governed by a Softmax algorithm. The higher the switch parameter (τ), the more uncorrelated the actual movement of the agent with payoff differences.
Dig Ability	λ: 1– 10	The ability of the agent to bring forth new knowledge. The agent searches for new information depending on the difference between current payoff and expected payoff for a given choice. The lower the dig parameter (λ), even fairly large differences are ignored and the agent does not dig further. A high ability agent has a high λ and further investigates even very small differences between expected and actual payoffs.
Cross-functional ability	0 – 0.5	The probability with which the agent surfaces new knowledge in the other dimension at each dig attempt – i.e., a row agent becomes aware of a new column. When the probability is zero, the row agent never becomes aware of new columns. When the probability is 0.5, the row agent has equal chance of observing a new row or a new column with each dig attempt.
Communication frequency	0 – 0.5	The probability with which the row (column) agent requests 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. ***

* In this paper, we do not vary the size of the search landscape.
** Initial granularity was varied between 1 and 64 in the robustness tests.
*** Note that in the model involving communication, the agents are strictly specialists – their cross-functional ability is held at zero.

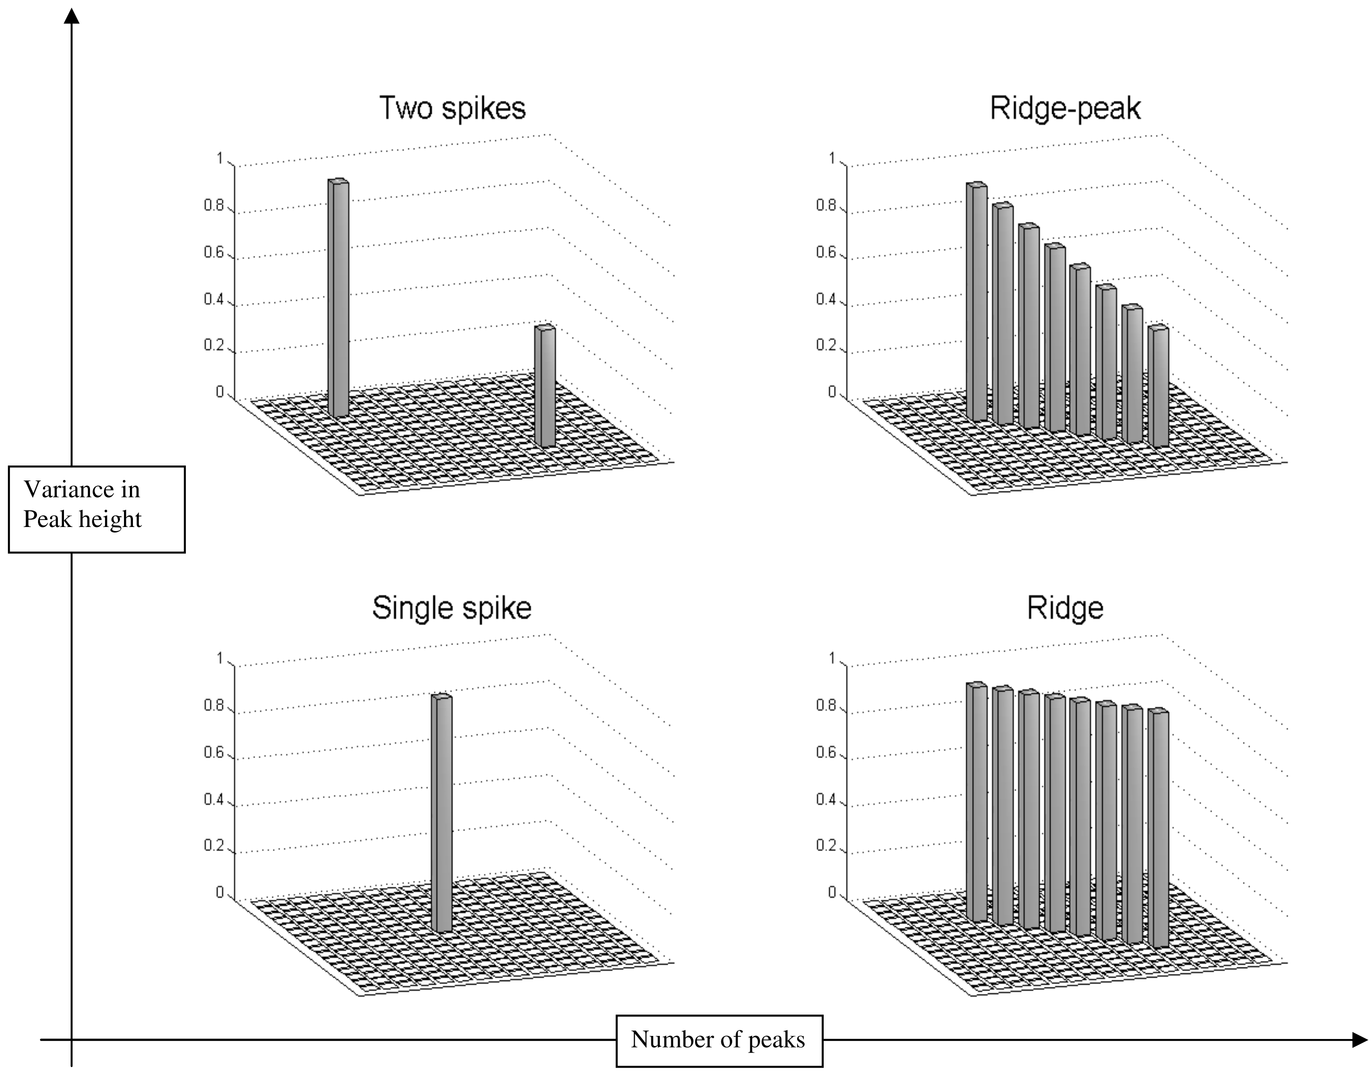

Figure 2 – Task environments (1/4 of actual size)

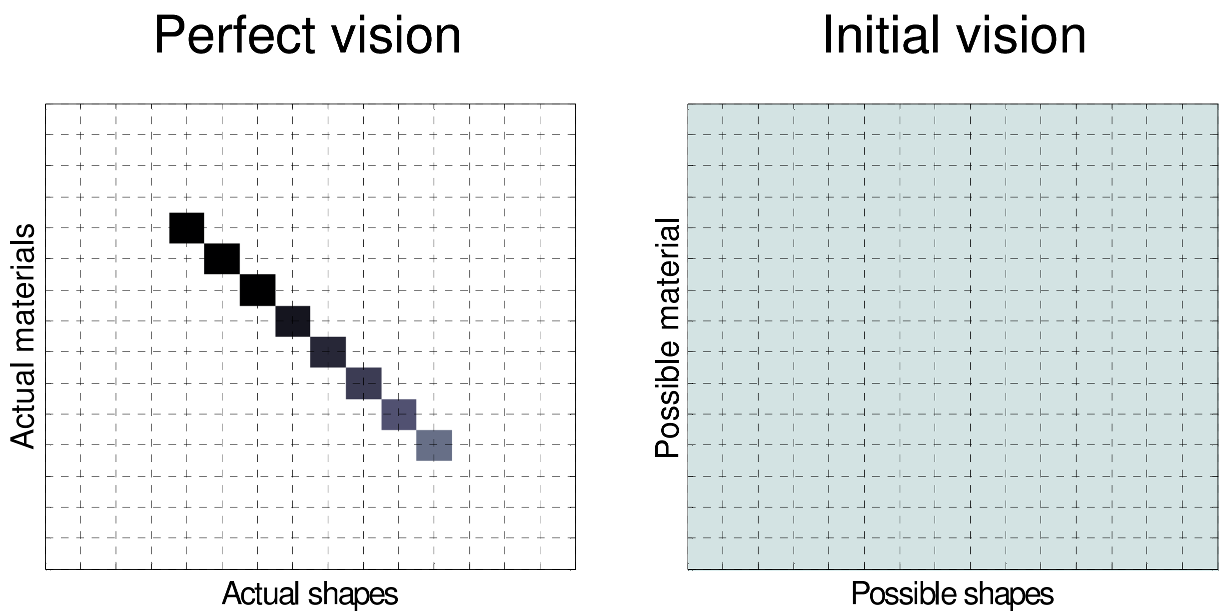

Figure 3: Perfect versus imperfect vision of ridge-spike landscape (1/4 of actual size)

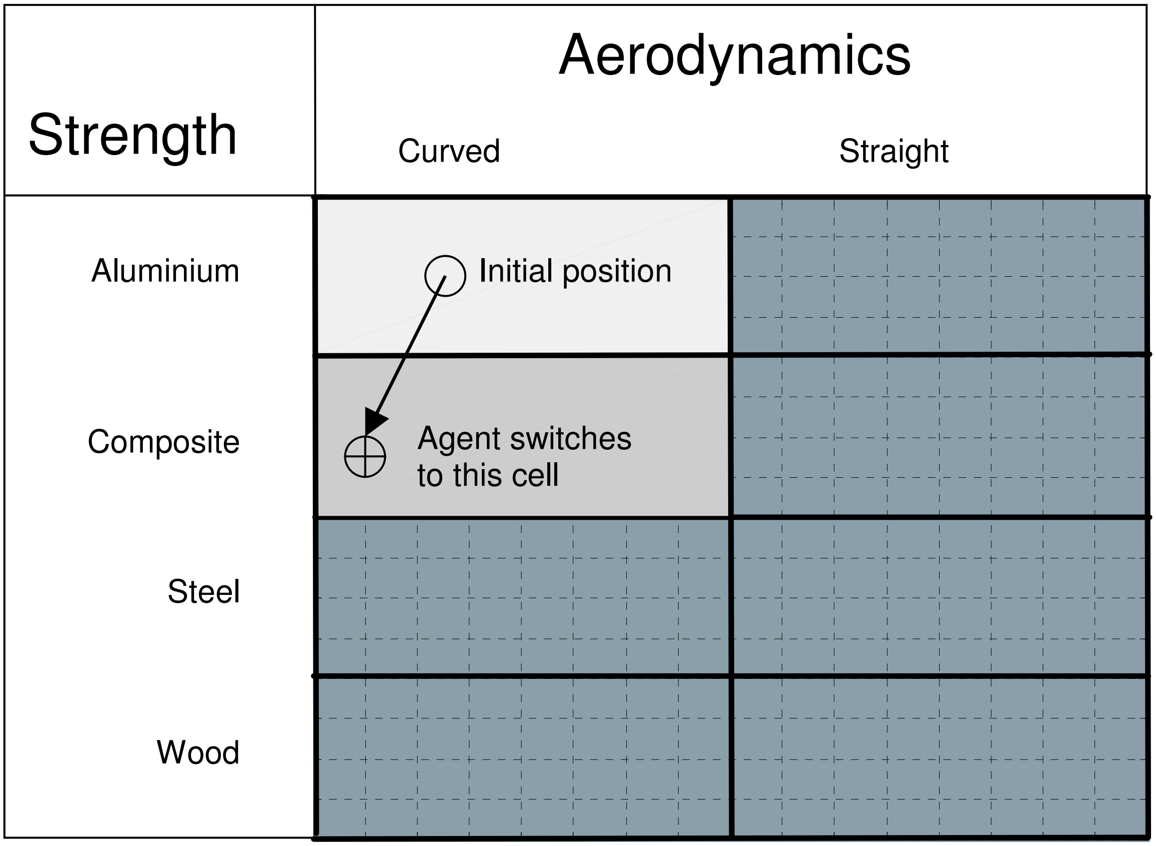

Figure 4: Agent Vision and Switching Operation to more promising alternative

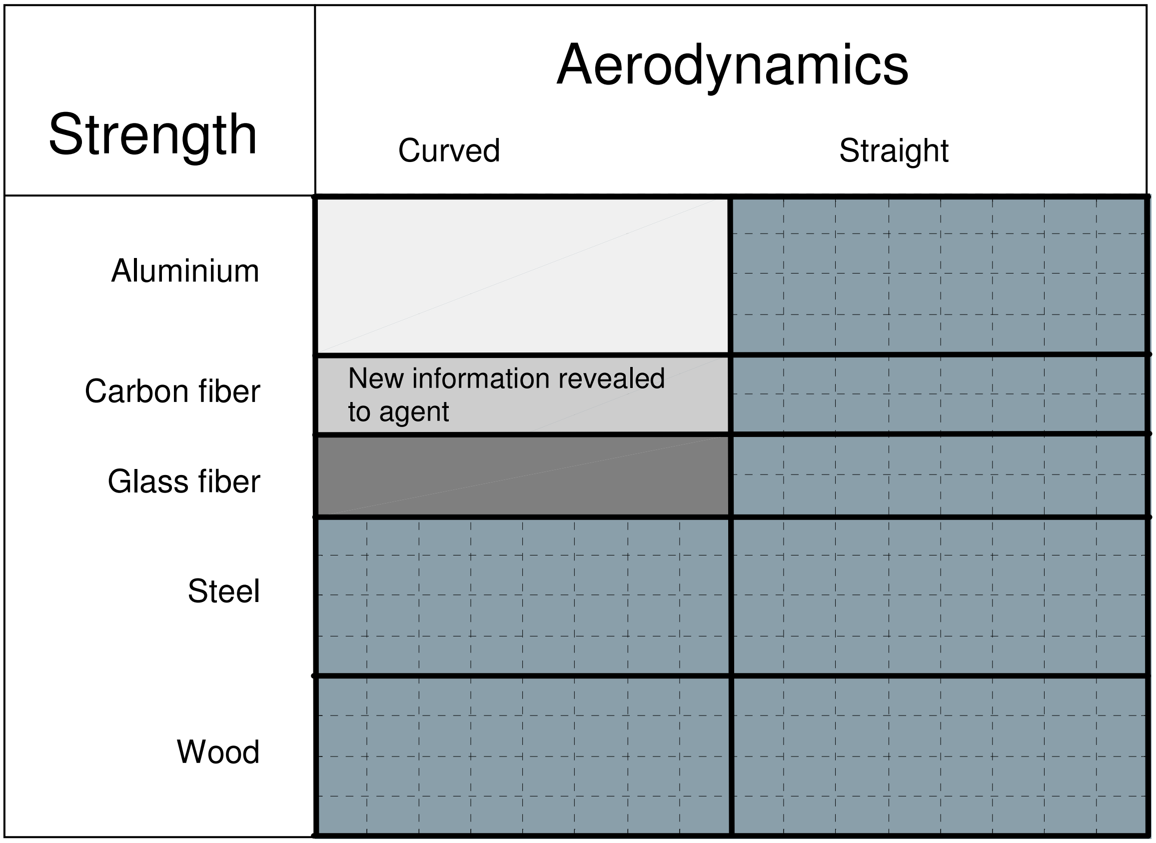

Figure 5: Dig Operation: New Information revealed to agent

Table 2: Agent characterization.

	Dig Ability, (DA)	Cross Functional Ability (CFA)
Dumb Agent	λ =1	0.05
Specialist Agent	λ =10	0.05
Generalist Agent	λ =1	0.5
Genius Agent (or T-shaped Skills)	λ =10	0.5

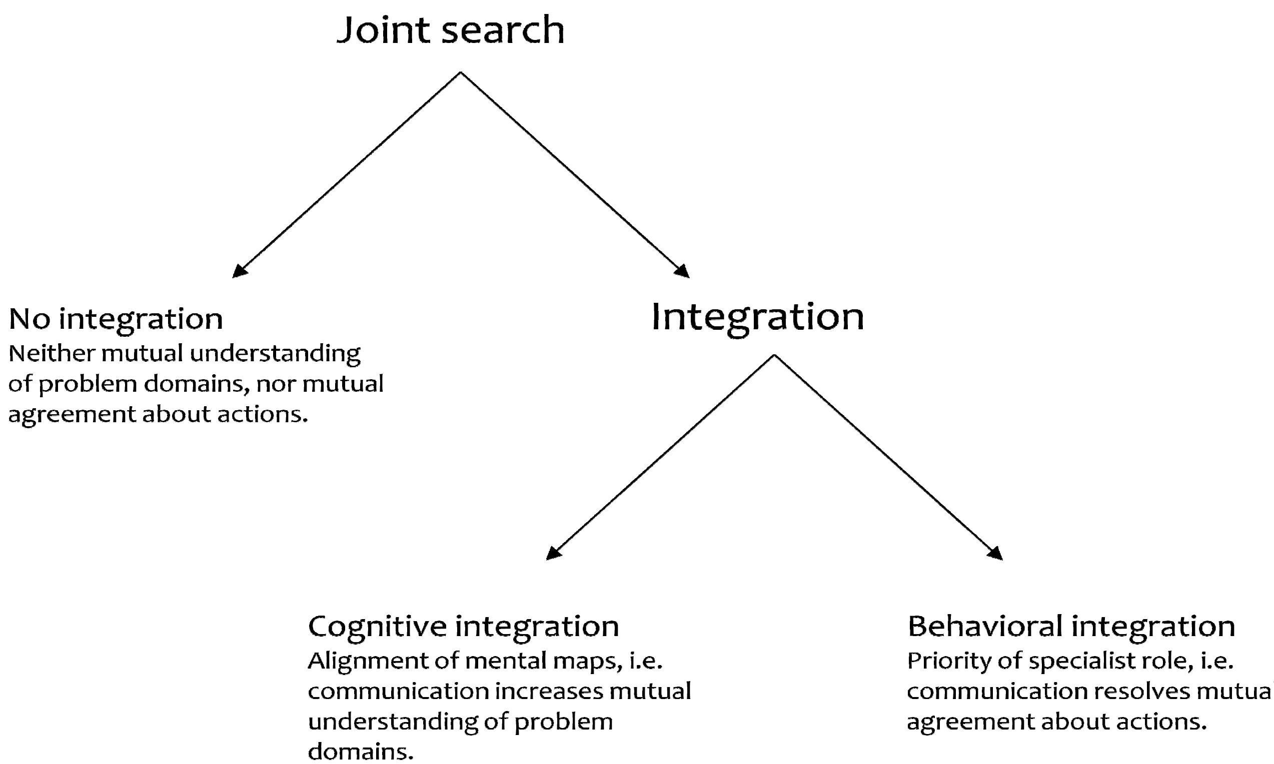

Figure 6: Different Organizations (Integration Mechanisms) investigated

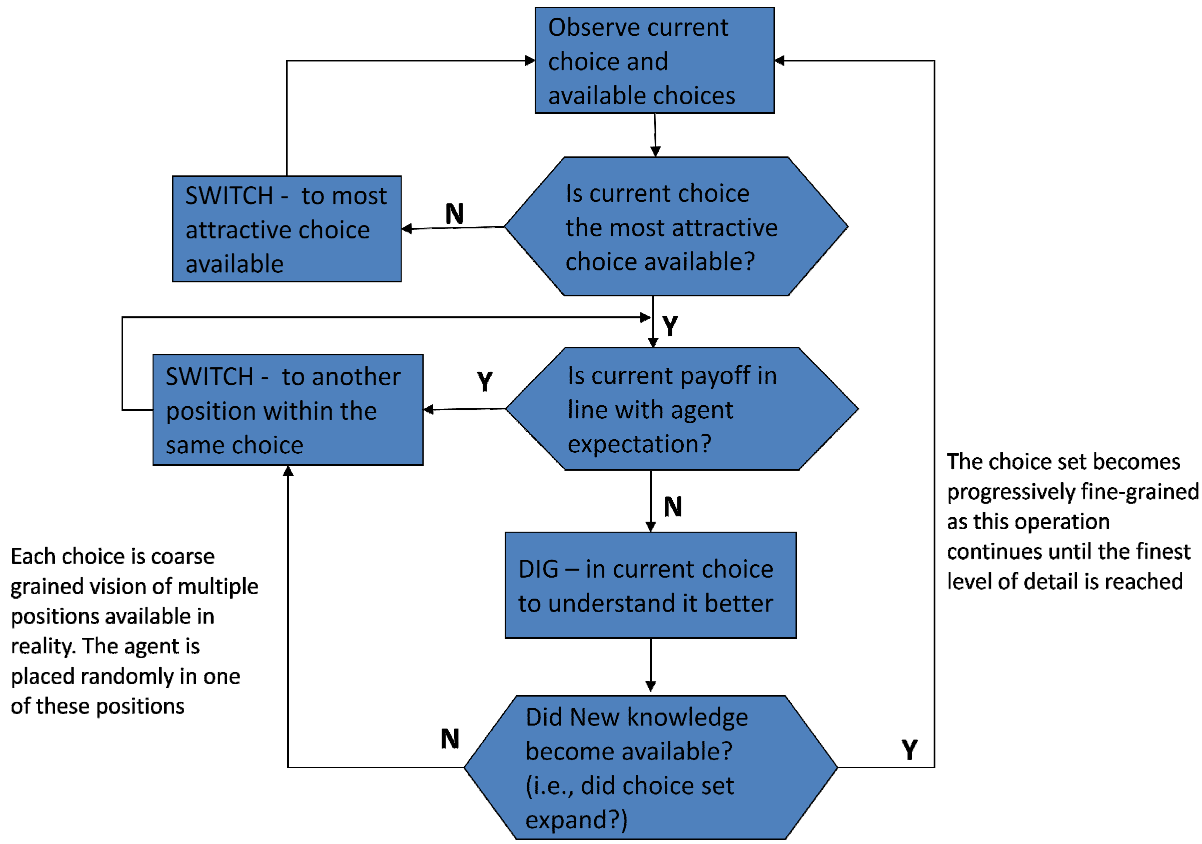

Figure7: Search Algorithm

Table 3: Results.

	Single spike		2 Spikes		Ridge		Ridge-Peak	
	Final Payoff	Total Payoff	Final Payoff	Total Payoff	Final Payoff	Total Payoff	Final Payoff	Total Payoff
No Integration								
Specialist	1.0	489	0.70	346	0.04	20	0.04	22
Generalist	1.0	463	0.65	295	0.03	16	0.03	16
T-Shaped Skills	1.0	481	0.63	304	0.04	20	0.03	13
Dumb Agent	1.0	455	0.65	303	0.04	20	0.05	24
Behavioral Integration								
Specialist	1.0	398	0.98	387	1.0	426	0.90	375
Generalist	1.0	453	0.97	449	1.0	480	0.89	424
T-Shaped Skills	1.0	483	0.98	473	1.0	487	0.90	439
Dumb Agent	1.0	316	0.97	315	1.0	386	0.90	314
Cognitive Integration – SPECIALISTS								
Low Frequency (1/20)	1.0	489	0.96	451	0.52	219	0.86	331
Medium Frequency (1/10)	1.0	489	0.96	459	0.58	257	0.85	377
High Frequency (1/2)	1.0	480	0.90	434	0.84	413	0.82	396
Cognitive Integration – Dumb Agent								
Low Frequency (1/20)	1.0	454	0.96	401	0.52	181	0.83*	251
Medium Frequency (1/10)	1.0	455	0.95	413	0.57	248	0.85	340
High Frequency (1/2)	1.0	434	0.89	400	0.92	415	0.82	382

Each value is computed from averages of 300 individual histories where 2 agents jointly search over a time-span comprising 500 discrete time steps. While different agents are characterized by distinct combinations of parameter DA and CFA as described in Table 2, the switch parameter (τ) is set to zero in all simulations shown here.

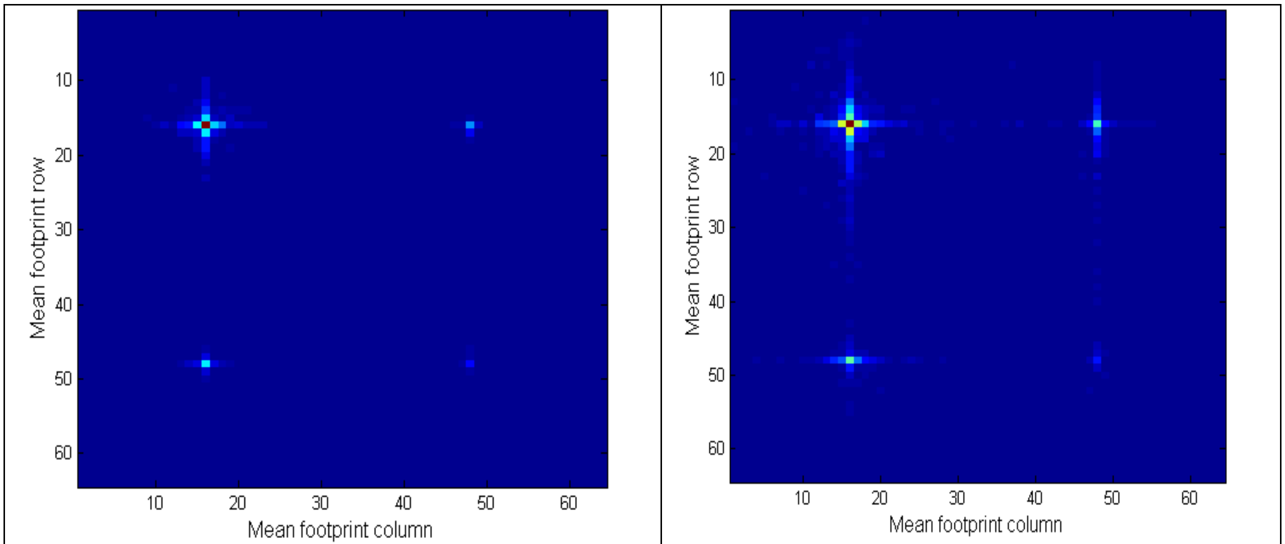

Panel a: Specialist Agent **Panel b: Genius Agent**

Figure 8: Mean Footprint of movement of specialist and genius agent in 2-Spike landscape

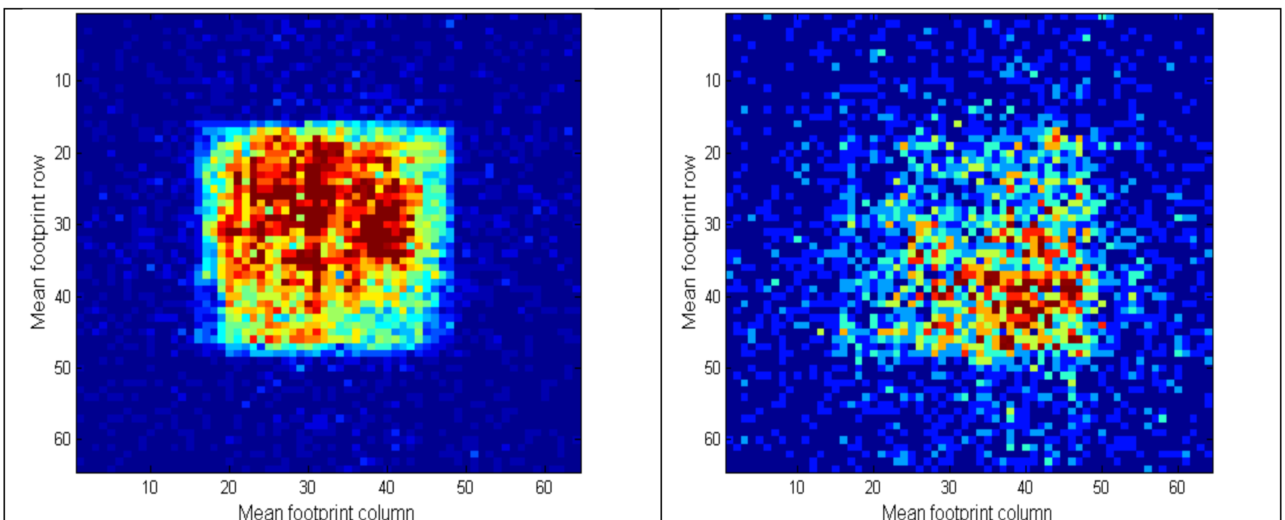

Panel a: Ridge landscape **Panel b: Ridge-Peak landscape**

Figure 9: Mean Footprint of movement of specialist agent in ridge and ridge-peak landscapes

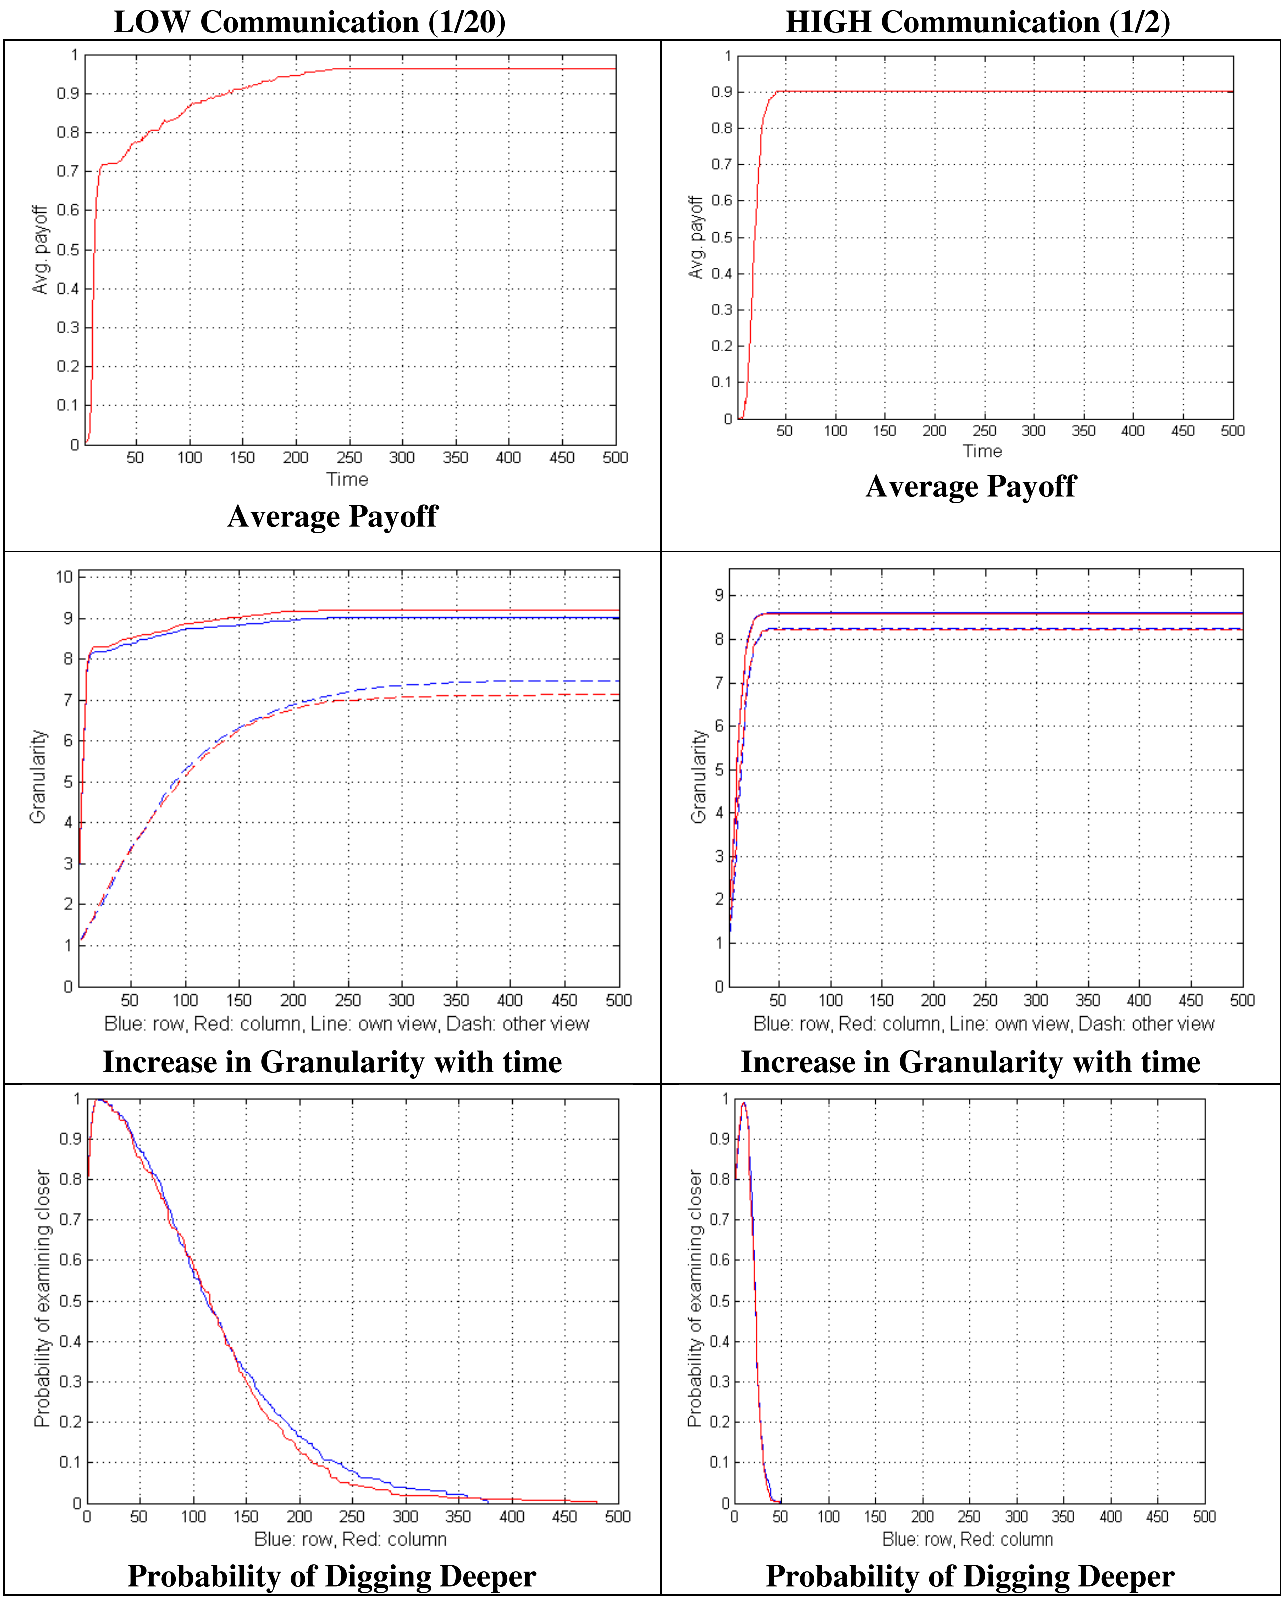

Figure 10: Search by Specialist Agent in 2-Spike landscape.

REFERENCES

Argote, L., and H. R. Greve
2007 "A Behavioral Theory of the Firm--40 Years and Counting: Introduction and Impact." Organization Science, 18: 337-349.

Bechky, B.
2003 "Sharing meaning across occupational communities: the transformation of understanding on a production floor." Organization Science, 14(3): 312-330.

Benner, M. J., and M. L. Tushman
2003 "Exploitation, exploration, and process management: the productivity dilemma revisited." Academy of Management Review, 28: 238-256.

Clark, K., 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 Utilities." American Economic Review, 74(1): 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.

Denrell, J. and J. G. March
2001 "Adaptation as Information Restriction: The Hot Stove Effect," Organization Science, 12 (5): 523-538.

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

Ethiraj, S. K., and D. Levinthal
2004 "Modularity and Innovation in Complex Systems." Management Science, 50: 159-173.

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.

2001 “Technological uncertainty in recombinant search.” Management Science, 47: 117-132.

Fleming, L.
2004 "Perfecting Cross-Pollination." Harvard Business Review, 82(9): 22-24.

Fleming, L.
2007 “Breakthroughs and the "Long Tail" of Innovation.” MIT Sloan Management Review, 49: 69-74.

Gavetti, G.
2005 “Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities' Development.” Organization Science, 16: 599-617

Gavetti, G., and D. Levinthal
2000 “Looking Forward and Looking Backward: Cognitive and Experiential Search.” Administrative Science Quarterly, 45: 113-137.

Gavetti, G., D. A. Levinthal, and J. W. Rivkin
2005 “Strategy making in novel and complex worlds: the power of analogy.” Strategic Management Journal, 26: 691-712.

Gavetti, G., and J. W. Rivkin
2007 “On the Origin of Strategy: Action and Cognition over Time.” Organization Science, 18: 420-439.

Heath, C., and N. Staudenmayer
2000 “Coordination Neglect: How Lay theories of organizing complicate coordination in organizations.” In R. L. Sutton, and B. M. Staw (Eds.), Research in Organization Behaviour, 22: 153-193, Greenwich, CT: JAI Press.

Henderson, R. and Clark, K.B.
1990 Architectural innovation: The reconfiguration of existing product technologies and failure of established firms. Administrative Science Quarterly 35 9-30.

Holland, J., K. Holyaok, R. Nisbett and P. Thagard
1986 Induction: Processes of inference, learning and discovery. Cambridge, MA: MIT Press.

Hoopes, D. G., and S. Postrel
1999 “Shared knowledge, "glitches," and product development performance.” Strategic Management Journal, 20: 837-865.

Huber, G. P., and K. Lewis
2010 “Cross-understanding: implications for group cognition and performance.” Academy of Management Review, 35: 6-26.

Iansiti, M.

1995 Technology integration: Managing technological evolution in a complex environment. Research policy, 24(4): 521-542.

Iansiti, M.
1998 Technology Integration. Boston, MA: Harvard Business School Press.

Katz, R. and M. Tushman
1983 "A longitudinal study of the effects of boundary spanning supervision on turnover and promotion in research and development." Academy of Management Journal, 26: 437-456.

Knudsen, T. and D.A. Levinthal
2007 Two faces of search: Alternative generation and alternative evaluation. Organization Science 18: 39-54.

Lawrence, P. and J. Lorsch
1967 Organization and environment: Managing differentiation and integration. Boston, MA: Harvard University Press.

Louis, M. and R. Sutton
1991 Switching cognitive gears: From habits of mind to active thinking. Human Relations. 44: 55-75.

Lounamaa, P. H. and J. G. March
1987 "Adaptive coordination of a learning team." Management Science, 33: 107-123.

Luce, R.D.
1959 *Individual Choice Behavior*. New York: Wiley.

March, J.G.
1988 Variable risk preferences and adaptive aspirations. Journal of Economic Behavior and Organization. 9(1): 5-24.

March, J. G.
1991 "Exploration and exploitation in organizational learning." Organization Science, 2: 71-87.

March, J.G.
1994 A primer on decision making. New York: Free Press.

March J. G., and H. A. Simon
1958 Organizations. Wiley Press: New York.

Nelson, R., and S. Winter
1982 An evolutionary theory of economic change. Cambridge, MA: Harvard University Press.

O'Reilly C. A., and M. L. Tushman

2004 “The ambidextrous organization.” Harvard Business Review, April: 74-81.

Page, S.E.
2007 The Difference. Princeton, NJ: Princeton University Press.

Postrel, S.E.
2002 “Islands of Shared Knowledge: Specialization and Mutual Understanding in Problem-Solving Teams.” Organization Science, 13: 303-320.

Puranam, P., and Swamy, M.
2008 Expeditions without Maps: Learning to Coordinate under Structural Uncertainty (April 2008). http://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
2006 “Organizing to Strategize in the Face of Interactions: Preventing Premature Lock-in.” Long Range Planning, 39: 591-614.

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, 57: 367-403.

Schelling, T. C.
1960 Strategies of conflict. Cambridge, MA: Harvard University Press.

Schumpeter, J.A.
1939 Business Cycles. McGraw Hill, NY.

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.

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 D. A. Levinthal
2005 “Escaping real (non-benign) competency traps: linking the dynamics of organizational structure to the dynamics of search.” Strategic Organization, 3: 85-115.

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.

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.

Sutton, R.S., and A.G. Barto.
1998 *Reinforcement learning*. Cambridge, Ma.: MIT Press.

Thagard, P.
1996 Mind: Introduction to cognitive science. Cambridge, MA: MIT Press.

Tushman, M. L.
1977 "Special Boundary Roles in the Innovation Process." Administrative Science Quarterly, Administrative Science Quarterly, 22: 587-605.

Tushman, M. L.
1978 "Technical Communication in R & D Laboratories: The Impact of Project Work Characteristics." Academy of Management Journal, 21: 624-645.

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 T. J. Scanlan
1981 "Boundary Spanning Individuals: Their Role in Information Transfer and Their Antecedents." Academy of Management Journal, 24: 289-305.

Von Hippel, E.
1988 The sources of innovation. Oxford University Press, New York.

Weick, K.E.
1995 Sensemaking in organizations. Thousand Oaks, CA: Sage.

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*.

$$p_{i,t} = e^{\overline{x_i}/\tau} \Big/ \sum_{i=1}^{d} e^{\overline{x_i}/\tau}$$

This is the standard version of the Softmax algorithm (Luce, 1959; Sutton & Barto, 1998). The "temperature" τ 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. As the exploration parameter approaches zero at the limit, the Softmax algorithm becomes an instance of greedy search.[25]

In each cell, the agent is positioned at a random combination. When the agent chooses not to switch cells, the agent is randomly assigned to a combination within the current cell.

[25] Somewhat surprisingly, it turned out that greedy search was far superior to even a little exploration. Therefore, we limit our presentation of results (as shown in Table 1) to very low values of τ =[0, 0.1]. Since this finding point to an essential feature that sets joint search apart from individual search, we retain the description in our presentation.

Figure 4 can illustrate. Suppose the agent is currently located in the cell labeled "composite materials" and "curved blades". While this is the cell that the agent currently perceives, the underlying reality matrix of that single cell is actually composed of 32 fine-grained choices. However, the agent is not (yet) able to see any differences at the level of such fine granularity. In consequence of not being able to see these finer possibilities, 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. For this reason, the agent will often experience a difference between the expected performance (average payoff across 32 locations in our example) and the received performance (actual payoff from one out the 32 possible locations). It is precisely this difference that motivates the agent to investigate further and improve the granularity of their vision.

<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:

$$q_{i,t} = 1 - 1 \Big/ e^{\lambda \left|\overline{x_i} - x_i\right| / \overline{x_i}}$$

This algorithm defines a probability distribution over differences between expected and actual performance. The dig parameter λ works very much like the exploration parameter for the switch operation, and is our operationalization of Dig Ability (DA)[26]. Higher values of λ makes the agent much more sensitive to deviations from expected performance, which results in

[26] We can recast this in terms of agent's ability – high ability agents are more sensitive to small payoff differences.

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).[27]

[27] Based on extensive sensitivity tests, we chose actual values of λ in the interval [1,10] as shown in Tables 1 and 2. A value of λ= 10 corresponds to a very high propensity to dig – this value is used to represent the specialist and genius agent. By contrast, a value of λ= 1 corresponds to a rather low propensity to dig – it is used to represent the generalist and the dumb agent (as shown in Table 2, the difference between these two agents is that the generalist has high cross-functional ability).