

A Knowledge-Based Theory of the Firm—The Problem-Solving Perspective

Jack A. Nickerson, Todd R. Zenger

John M. Olin School of Business, Washington University, Campus Box 1133, One Brookings Drive, St. Louis, Missouri 63130-4899
 {nickerson@wustl.edu, zenger@wustl.edu}

In this paper we develop a knowledge-based theory of the firm. While existing knowledge-based theory focuses on the efficiency of hierarchy in economizing on knowledge exchange, we develop a theory of the firm that focuses on the efficiency of alternative organizational forms in *generating* knowledge or capability. Our theory begins with the problem as the basic unit of analysis, arguing that a problem's complexity influences the optimal method of solution search and the optimal means of organizing that search. The distinguishing feature that differentiates among organizational alternatives is the different way each resolves conflict over the selection of solution trials, that is, the way it chooses the path of search. Our theory predicts that efficiency demands that these governance alternatives be matched in a discriminating way to problems based on their associated benefits and costs in governing solution search. Thus, our theory is among the first to simultaneously treat both the boundary choice (i.e., internal versus external) and the choice among alternative internal approaches to organizing.

Key words: knowledge-based view; knowledge management; problem solving; innovation; firm boundaries; organizational boundaries

1. Introduction

In the strategy literature, a key task of the manager is to accumulate and protect valuable knowledge or capability (Rumelt 1984, Barney 1984, Wernerfelt 1984, Teece et al. 1997). Such knowledge or capability defines a firm's capacity to efficiently convert its inputs into valuable outputs (Arrow and Hahn 1971; Debreu 1959; Nelson and Winter 1982, pp. 59–60). Thus, managers enhance the firm's capacity to produce efficiently by updating or advancing knowledge. Therefore, a common assumption and prescription in the strategy literature is that the boundaries of the firm should encompass these valuable competencies and core knowledge (Argyres 1996, Prahalad and Hamel 1990). By internalizing valuable knowledge or keeping this knowledge internal, the firm positions itself to both exploit and protect knowledge. Arguably, however, the key knowledge-based question the manager faces is not how to organize to exploit already developed knowledge or capability, but rather how to organize to efficiently *generate* knowledge and capability.

Recently, scholars within the knowledge- or resource-based perspective have focused their attention on this question. In the process, scholars have sought to develop what some call a knowledge-based view or knowledge-based theory of the firm (Conner 1991; Demsetz 1988; Conner and Prahalad 1996; Kogut and Zander 1992, 1996; Grant 1996; Madhok 1996; Nahapiet and Ghoshal 1998). This work seeks to explain how the choice of organization—particularly the choice of whether

to integrate or outsource an activity—influences the efficient production and protection of valuable knowledge and capabilities. In so doing, this literature seeks a theory of the firm that is independent of transaction cost logic and its behavioral assumption of opportunism. A common argument in much of this literature is that firms as organizational forms exist to economize on the exchange of knowledge rather than to attenuate opportunism. As Conner (1991, p. 139) argues, the integration of an activity within the boundaries of the firm thereby becomes a “creator of a positive” rather than an “avoider of a negative.”

While this existing literature on the knowledge-based view of the literature makes important contributions to our understanding of the firm, it has significant shortcomings. Others have convincingly contested the claim that a knowledge-based theory of the firm can be independent of opportunism (Foss 1996a,b; Heiman and Nickerson 2002; Mahoney 2001; Williamson 1999), but our critique and contribution addresses other shortcomings. First, this literature has primarily focused on the role of firms in providing efficient knowledge exchange rather than their role in efficiently producing knowledge or capabilities. Second, the two fundamental arguments within the literature that support the efficiency of firms in knowledge exchange relative to markets are fully contradictory. One claims that hierarchies exist to essentially *avoid* knowledge transfer (Demsetz 1988, Conner 1991, Conner and Prahalad 1996), emphasizing the firm's capacity to exercise authority in directing

others' actions; the other view claims that hierarchies exist instead to *facilitate* knowledge transfer (Arrow 1974; Kogut and Zander 1992, 1996; Nahapiet and Ghoshal 1998), emphasizing the firm's capacity to support the formation of shared language and identity. Finally, the literature has yet to provide a theory of alignment—a theory that predicts when hierarchies are preferred to markets or when markets prevail, let alone use knowledge-based logic to do so. Indeed, the knowledge-based arguments tend to highlight the virtues of hierarchy, not its limits in forming and transferring knowledge. Hierarchy is argued to be superior in transferring and forming knowledge, leaving other factors to constrain firm boundaries (see Conner and Prahalad 1996). In summary, we do not yet have a knowledge-based theory of the firm.

In this paper, we develop a knowledge-based theory of the firm that addresses these shortcomings. We explain how a firm's prospective objectives for knowledge formation dictate the choice of how to organize. Here the critical question is not whether knowledge should be owned or acquired in the market or how the exchange of knowledge should be facilitated, but rather how a manager should organize individuals to generate knowledge that the firm seeks.

We begin with the assumption that the manager's knowledge-based objective is to create valuable new knowledge. The manager, however, cannot simply choose new knowledge to acquire, because more often than not the desired knowledge does not exist. Instead, the manager must choose valuable problems—those which, if successfully solved, yield desirable knowledge or capability. The value of a particular problem therefore depends on two factors: (1) the values in the array of possible solutions and (2) the costs of discovering a particularly valuable one. Once a problem is selected, the manager then organizes a search for solutions that optimizes the likelihood, speed, and cost with which valuable solutions are discovered. Finally, the manager seeks to appropriate a portion of the solution's value. While we present these steps sequentially for purposes of exposition, in truth the costs of solution search and the capacity to appropriate value must be factored into the initial choice of problems as well. Thus, the unit of analysis in our theory is the problem, and profitable formation of new knowledge (or new capability) is the central goal in choosing a form of governance.

Following previous work, we assume that solutions to complex problems represent unique combinations or syntheses of existing knowledge (Fleming and Sorenson 2000, Henderson and Clark 1990). Recently, scholars have depicted this matrix of unique combinations of knowledge as a landscape. Some solution landscapes are rugged, with many high-value solutions scattered across the terrain. Others are smooth, with a single, high-value solution or with high-value solutions concentrated

in a particular region of the landscape. Consequently, problems differ in the optimal form of solution search. For some, directional search or search based on trial and error leads efficiently to a high-value solution (Gavetti and Levinthal 2000). For others, search guided by a cognitively developed heuristic of the landscape's topography is preferred (Gavetti and Levinthal 2000).

After identifying a problem and assessing the ideal form of search, the manager must decide how to organize this search. In particular, the manager must decide how to access relevant knowledge either within or outside the firm. We explicitly explore the capacity of three alternative governance modes to support these differing forms of search: markets, authority-based hierarchy, and consensus-based hierarchy. The differentiating feature among these governance alternatives is the different way in which each resolves conflict over the selection of solution trials—that is, the way it chooses the path of search. Our theory predicts that efficiency demands that these governance alternatives be matched in a discriminating way to problems based on their associated benefits and costs in governing solution search. Thus, our theory is among the first to simultaneously treat both the boundary choice (i.e., internal versus external) and the choice among alternative internal approaches to organizing.

Our theory development proceeds as follows. In §2, we introduce the problem as the unit of analysis, introducing the concept of solution landscapes that differ based on problem complexity. In that section we match problem characteristics to search strategies. In §3, we discuss knowledge-exchange hazards that hinder efforts to conduct one form of search. In §4, we discuss the capacity of alternative governance structures to conduct differing forms of search. In §5 we finally match problems to alternative governance forms. Section 6 discusses our theory and describes future avenues for research.

2. Capability and Knowledge Formation as Problem Solving

Problem solving and knowledge formation are at the center of our theory. The manager's fundamental knowledge-based objective is to sustain above-normal profits by continually discovering new knowledge or new solutions that form from unique combinations of existing knowledge. Nelson and Winter (1982) explicitly define a firm's knowledge (or capability) as the "input-output combinations achievable with all possible mixes and levels of activities known to the firm" (pp. 63–64). In this literature, the state of a firm's knowledge can be advanced by either absorbing existing knowledge external to the firm or by developing new knowledge by first identifying a problem and then discovering a valuable new solution. Therefore, if a firm is to develop unique

knowledge or a unique new capability through any manner other than luck, it must identify a valuable problem and conduct an efficient solution search. Valuable solutions deliver value to the firm, either through enhancement or development of a product or service or by reducing the cost of production or delivery.

In our theory, managers choose problems while identifying knowledge sets or existing technology either within or outside the firm that are potentially useful in searching for solutions to that problem (Nelson and Winter 1982). The choice of problems reflects an assessment of the expected value of potential solutions and an assessment of the firm's capacity to profitably reach high-value solutions. A firm's reservoir of knowledge sets and prior experience informs this choice. In choosing problems, managers in essence choose an unknown set of potential solutions, but once a problem is chosen, the task becomes identifying relevant knowledge and then maximizing the probability of discovering a high-value solution. This is achieved by choosing organizational mechanisms that efficiently govern search. To optimally organize search first necessitates an understanding of the solution space under exploration.

2.1. Dimensionalizing Solution Landscapes

Our conceptualization of problems and solution search stems from Simon's (1962) work on complex systems and Kauffman's (1993) broad work on *NK* modeling. Simon (1962) conceptualized a complex system as one "made up of a large number of parts that interact in a nonsimple way" (p. 486). Complex problems represent complex systems in which the value of recombinations of existing technology represents solutions (Fleming 2001, Schumpeter 1939, Henderson and Clark 1990, Weitzman 1996, Fleming and Sorenson 2000). Thus, as Fleming and Sorenson (2000) argue, the invention of the automobile combined technology from "the bicycle, the horse carriage, and the internal combustion engine," while the microprocessor combined knowledge "of a computer's central processing unit with integrated circuit fabrication processes." Continuing with Simon's logic, the complexity of a problem is a function of the degree to which the individual design choices, which define a solution, are either independent or interdependent in their contribution to solution value.

We conceptualize this level of interaction using Kauffman's *NK* modeling. This modeling approach has been used extensively in recent years as a simple mathematical construction of the interaction of elements within systems of varying complexity. In these models *N* represents the number of parts of the system (in our case, design choices) and *K*, which ranges from zero to one, represents the degree of interaction among the parts (or design choices). These two parameters then combine to describe landscapes that are either rugged, reflecting a

high degree of interaction, or smooth, representing a low level of interaction.

Thus, suppose we identify that a particular problem may benefit from five distinct knowledge sets, with each knowledge set informing a distinct binary design choice. Each unique combination of these design choices yields a distinct solution value. The level of interaction among these binary design choices defines the topography of a solution landscape. The greater this interaction, the more rugged the landscape is, the higher the value of the highest valued solutions is, and the lower the spatial autocorrelation among high-value solutions is. The lower the interaction is, the smoother the landscape (and the lower the value of the highest solution).

For any given problem, the set of all possible combinations of relevant existing knowledge (i.e., solutions) can be represented as a landscape the topography of which defines the value associated with any given combination. Peaks on such solution landscapes represent valuable combinations of knowledge sets or technologies that are highly complementary. Valleys on such landscapes represent low-value combinations of existing knowledge. When knowledge sets are highly interdependent, solution landscapes are more rugged and unpredictable (Fleming and Sorenson 2000). Under these conditions, the value of the global maximum rises, but the average height of peaks declines. On these more rugged landscapes, a series of incremental changes in design is unlikely to lead to the discovery of highly valuable solutions (Fleming and Sorenson 2000).

Problems and their corresponding solution landscapes thus can be arrayed according to the level of interaction among knowledge sets. Borrowing from Simon's typology of complex systems, we provide an analogous typology of problems—decomposable or low-interaction problems, nondecomposable or high-interaction problems, and nearly decomposable problems with moderate levels of knowledge interaction.¹

Low-interaction/decomposable problems are problems in which the value of solutions depends very little on the interactions among knowledge sets and design choices. With such problems, groups of individuals possessing rather distinct knowledge sets can independently apply their knowledge to unique design choices with a reasonable expectation that the aggregation of their independent efforts, along with the independent efforts of others with distinctly different knowledge sets, will uncover valuable problem solutions. For instance, the problem of discovering a higher-performing desktop personal computer is currently a rather low-interaction problem in knowledge formation. Within a certain range, computer performance can be increased by actors independently improving any number of subsystems (e.g., disk drive, monitor, Ethernet card) or components (e.g., microprocessor, memory). To use Simon's language, low-interaction problems are "decomposable" in

the sense that these problems can be subdivided into subproblems, each of which draws from rather specialized knowledge sets. Those possessing knowledge set A can independently address the subproblem of selecting from among choices a_i , while those possessing knowledge set B can independently address the problem of optimizing among b_j , where $A \cap B = \{\phi\}$. In Kauffman's NK model, such low-interaction problems correspond to gradual and relatively smooth variations in solution landscapes with few peaks and high spatial autocorrelation, implying that high-value solutions tend to cluster in the same region.

High-interaction problems, by contrast, have solution landscapes in which the value of solutions depends on interactions among design choices. For instance, consider one of the subproblems above, designing a leading-edge microprocessor. To design and build a leading-edge microprocessor circuit demands numerous knowledge sets that extensively interact in determining the value of solutions (Macher 2001). With such high-interaction problems, an actor familiar with a particular technology cannot predictably enhance the value of the product design based solely on the knowledge he or she possesses. In this case, the value of any particular design change will interact with a host of other potential design changes determined by actors possessing distinctly different knowledge sets. In Kauffman's NK logic, the more extensive the interactions among knowledge sets are, the more rugged the landscape with abruptly jagged peaks and valleys. With high-interaction problems, knowledge sets cannot be separated into subproblems. Interactions among distinct knowledge sets are so extensive that any attempt to define subproblems and discover corresponding subordinate solutions offers no predictable improvement over randomly selecting trials. In Simon's language such problems are *nondecomposable*.

Moderate-interaction problems are problems that, in Simon's language, are *nearly decomposable*. The level of interaction among design choices is intermediate in the sense that subproblems associated with distinctive knowledge sets can be defined, but the value of a design choice within one subproblem is not fully independent of the design choices made in another subproblem. Thus, interactions among knowledge sets are not trivial. Nonetheless, near decomposability implies that interactions among knowledge sets within subproblems are greater than among subproblems.²

Staying with our computer illustration, the problem of enhancing the performance of notebook computers is generally regarded as a problem involving significantly more interdependence among design choices than the problem of refining desktop systems (Hoetker 2003). While the overall problem can clearly be separated into subproblems associated with designing the screen, keyboard, motherboard, etc., there are significant trade-offs and interdependencies in design choices for these

components as they relate to performance dimensions of weight, size, and battery life (Hoetker 2003). Thus, the design of a leading-edge notebook computer is currently a nearly decomposable problem in the sense that non-trivial interactions exist among the knowledge sets that support subproblems, but these are significantly less than the interactions among knowledge sets applicable within subproblems. In such cases, Kauffman's NK model yields intermediate levels of ruggedness with such an intermediate level of interaction among knowledge sets.

2.2. Efficiently Searching Solution Landscapes

The search for solutions is necessarily uncertain. The likelihood, speed, and cost of arriving at a valuable solution therefore depend both on luck and on the pattern of trials that actors undertake (Simon 1962), with each trial reflecting a unique combination of knowledge or a unique set of design choices. If trials are chosen randomly, the probability of arriving at a valuable solution to a complex, high-interaction problem is very low. More efficient orderings utilize knowledge to direct this search process. There are two fundamental sources of such knowledge that correspond to two distinct approaches to search: directional and heuristic.³

*Directional or local search*⁴ is search guided solely by feedback or experience from prior trials. New combinations of knowledge are pursued by altering one design element at a time (Gavetti and Levinthal 2000), observing the resulting change in solution value, and then either continuing with the same path of search or—in the wake of declining performance—restoring the design to its original form and adjusting a different element. The key feature of this form of search is that independent actors, either individuals or groups, independently pursue trials and independently observe performance. This form of experiential search or search through trial and error has been described in a wide range of literature (March and Simon 1958, Cyert and March 1963, Nelson and Winter 1982).

Directional search is well suited to low-interaction problems—problems where there is little interaction among design choices in defining solution performance. On such solution landscapes, the only relevant feedback is whether performance is increasing or declining in response to a specific design change. Thus, searching for solutions to such problems is a fully decomposable activity.

Directional search is, of course, quite inefficient when exploring solutions to high-interaction problems. Design choices here are highly interdependent. The rugged, multi-peaked topography of complex solution landscapes ensures that directional search through independent design adjustments is unlikely to discover a valuable peak. Indeed, on this rugged terrain, small changes in design may result in substantial shifts in solution value. Thus, the efficacy of directional search diminishes as

problems become less decomposable and knowledge interactions increase.

Heuristic or cognitive search is a form of solution search in which an actor or a group of actors cognitively evaluate the probable consequences of design choices rather than relying solely on feedback after design choices are made (Simon 1991, March and Simon 1958). Trials are thus selected based on a cognitive map or implicit theory of how knowledge sets and specific design choices relevant to the problem interact to determine solution performance. These heuristics can be thought of as simplified representations of the solution landscapes or, as Gavetti and Levinthal describe, “cognitive representations . . . of lower dimensionality than the actual landscapes” (2000, p. 121). These “theory-driven structures” speed problem, solving by “furnishing a basis for evaluating information” (Walsh 1995). These cognitive representations of the solution landscape are then used to select trials that maximize the probability of quickly discovering a high-value solution. Of course, cognitive maps are not static. As trials are undertaken and knowledge is gleaned through feedback, managers update their heuristics.

Heuristics can be of two essential forms, based either on a single individual’s beliefs and cognitions or on a group’s collective beliefs. If each actor possessed unlimited cognitive capacity, then each actor could quickly absorb all relevant knowledge, develop theories about the structure of interactions, and choose an optimal pattern of search. However, human minds are limited in the rate at which knowledge can be assimilated, accumulated, and applied (Simon 1945, pp. 40–41). Hence, distinct knowledge sets needed for solving complex problems are likely to be widely dispersed and reside in the minds of many agents who have specialized in unique knowledge, which is why we focus our attention on groups of actors engaged in heuristic search.

Thus, consider three actors with different and distinct knowledge sets, all of which are relevant to a particular problem. Assume that these knowledge sets reflect a set of design choices that define a solution landscape to a particular problem and that the landscape is at least somewhat rugged. Each actor can independently develop an ordering for trials based on his or her respective knowledge. However, differences in knowledge are likely to yield differences in cognitive maps and hence the recommended order of trials. Moreover, absent knowledge sharing, these independently developed cognitive maps are unlikely to be well matched to the topography of the solution landscapes. Only by developing heuristics that encompass the knowledge of all actors can the probability of discovering highly valued solutions be enhanced.

While individual beliefs are the basis of heuristics that guide group decisions, developing group heuristics requires the resolution of inherent conflicts in beliefs

(Walsh and Fahey 1986, Fiol 1994). These group heuristics are negotiated beliefs that are shaped not only by the quality of actors’ logic, but also by each actor’s self-interest and political position (Walsh and Fahey 1986). Thus, while individual beliefs form the basis of group beliefs or heuristics, not all individual beliefs are equally important or influential in this process.

Heuristic search first necessitates knowledge transfer to facilitate the development of heuristics that derive from multiple and dispersed knowledge sets. Knowledge transfer in turn necessitates the development of a shared language to support it. Finally, the development of group heuristics requires the reconciliation of the divergent beliefs about the proper shape of the search heuristic. By contrast, directional search requires little more than choosing an initial set of design choices and allowing each agent to independently choose trials.

These two search processes are easily matched to problem types. Heuristic search is necessary when problems are complex, with high levels of interaction among knowledge sets and design choices. Directional search is warranted when problems are decomposable, involving limited knowledge interaction because heuristics offer no clear benefit in this setting. Nearly decomposable, moderate-interaction problems benefit from both directional and heuristic search: Heuristic search essentially defines the area of the solution landscape where directional search will take place. Thus, as problems become more complex and solution landscapes more rugged, efficient search necessitates extensive knowledge exchange to craft search heuristics. However, because these intermediately complex problems yield landscapes with significant autocorrelation and thus locational clustering of high-value solutions, directional search in regions targeted by cognitive maps may prove particularly fruitful. Cognitive maps in this setting aid in discovering regions of the solution landscape particularly attractive for search, while directional search involving independent actors effectively explores these regions. Table 1 summarizes the relationship between problem type and the relative benefit of directional and heuristic search.

Table 1 Relative Benefit of Trial Ordering and Selection by Problem Type

Trial ordering and selection mechanisms	Attributes of knowledge formation		
	Decomposable	Nearly decomposable	Nondecomposable
Directional search	++	+	0
Heuristic search	0	+	++

3. Governance Choices and Knowledge Formation Hazards

If differing problems demand differing search processes, then the manager's task is to identify relevant knowledge sets and to craft a mechanism of governance that supports or enhances the method of search appropriate for the chosen problem. For simple, decomposable problems involving limited interaction among design choices, directional search based on feedback is desired. For nondecomposable problems involving extensive interactions among design choices, heuristic search is desired. However, as noted, heuristic search involving several specialized knowledge sets demands that individuals share the knowledge they possess to develop common cognitive maps. Two conditions, however, impede such knowledge sharing: (1) humans are cognitively constrained in the speed with which they can learn (i.e., assimilate, accumulate, and apply knowledge) and (2) they are prone to self-interest, which may include opportunistic forms of self-interest. The former ensures that knowledge relevant to exploring complex problems is widely distributed among individuals. No one individual has all the knowledge relevant to crafting an effective heuristics or cognitive map to guide search on a complex solution landscape. In and of itself, this wide dissemination of knowledge does not constrain the efficient organization of search because individuals could willingly share knowledge when heuristic search is demanded and collectively agree on an optimal path of search. However, the wide distribution of knowledge in conjunction with self-interest, in particular an individual's desire to capture value from the knowledge he or she accumulates, leads to two knowledge-related exchange hazards. These hazards plague efforts to support knowledge sharing necessary for heuristic search.

Hazard 1: Knowledge Appropriation

Incentives alter an individual's willingness to share knowledge. Unfortunately, with independent actors, incentives for knowledge sharing are plagued by a rather simple paradox: The value of knowledge to its potential acquirer is not known until after the knowledge is revealed; however, once that value is revealed, the potential acquirer has no need to pay for it and can resell it at near zero marginal cost (Arrow 1973, p. 171). Thus, the incentives to simply assimilate others' knowledge and thereby extract its value without payment plague efforts to sell knowledge. While property rights and contracts provide some protection, cognitive limits render contractual protections difficult to identify and costly to draft. Thus, opportunism in knowledge exchange discourages actors from sharing knowledge, and knowledge is therefore not placed into the hands of those who will find it of most value. Instead, self-interest encourages actors to hoard knowledge and

embed it into saleable products. Embedding knowledge into saleable products provides a mechanism for extracting its value. If, however, the value of an actor's knowledge is revealed only through recombination with other actors' knowledge, then knowledge sharing is necessary, which—absent a means of protection—may simply not happen.

The knowledge-appropriation hazard therefore has clear implications for organizing search. While directional search is unaffected because knowledge sharing is not required for this form of search, heuristic search is clearly damaged by this disincentive to share knowledge. Heuristic search requires extensive knowledge sharing, and hence a remedy for the knowledge-appropriation hazard is required.

Hazard 2: Strategic Knowledge Accumulation

Individuals not only possess incentives to hoard knowledge; they also possess incentives to shape its development. Individuals possess incentives to strategically alter the path of search and the shape of the heuristic that guides search. In particular, individual actors may prefer to shape the search heuristic in ways that optimize the value of the knowledge that they personally accumulate as solutions are explored. The search process itself yields new knowledge that may be useful in a host of different applications. The value of knowledge that actors derive through solution search will depend on the complementary knowledge that they already possess. Actors prefer to see solutions explored that draw heavily from their own knowledge base because such solutions heighten the potential value that they can extract personally. If the organization is more dependent on the knowledge possessed by an actor, then that actor is in a stronger position to bargain for a large portion of the value generated by a solution.

Thus, actors have incentives to strategically influence the pattern of trials in ways that enhance their specialized knowledge or complement knowledge that they already possess, while avoiding efforts that require knowledge sharing. Consequently, absent governance remedies, efforts to explore problems requiring heuristic search are likely to lead to attempts to distort cognitive maps, to conflicts regarding the proper ordering of trials, and more generally to an underinvestment in knowledge-sharing activities that facilitate the development of common heuristics.

We assert that both hazards are essentially irrelevant when problems are decomposable. When problems are decomposable, searching for solutions does not benefit from knowledge transfer and heuristic search. However, nondecomposable, high-interaction problems require knowledge transfer to enable heuristic search. Consequently, as problems become more complex and nondecomposable, efficient search demands mechanisms that mitigate knowledge-exchange hazards.

4. Governance Alternatives for Mitigating Knowledge-Formation Hazards

To this point we have argued that problems differ in complexity and in optimal forms of solution search. Complex problems require solution search that involves extensive knowledge transfer, which is discouraged by knowledge-exchange hazards. Efficiency thus dictates the selection of governance forms, which support the appropriate level of knowledge transfer. At the same time, efficiency dictates the maintenance of high-powered incentives that encourage actors to actively seek knowledge through solution search. As we will see, there is a fundamental trade-off between governance forms that powerfully motivate search effort and those that more effectively support heuristic development.

In the discussion below, we explore three prototypical alternative governance forms as mechanisms for governing solution search: markets, authority-based hierarchies, and consensus-based hierarchies.⁵ Following Williamson (1991), we contend that these three governance forms differ fundamentally in their use of three organizational features: (1) decision rights over the path of solution search, (2) communication channels to support knowledge transfer, and (3) incentives to motivate search. In the discussion below, we argue that these organizational features are configured in unique and complementary ways to achieve three organizational forms or prototypes. While scholars have touted the virtues of each form as a solution to the governance of knowledge exchange, these treatments are generally not comparative and provide no basis for discriminating among them.⁶ Thus, we develop here an alignment between governance forms and problem types.

We begin with a discussion of markets, focusing particularly on Hayek's discussion of its virtues as a device for efficiently governing knowledge exchange. Markets, however, have a very low capacity for remedying knowledge-exchange hazards and thus fail as problems become more complex. We argue that hierarchy fundamentally comes in two forms: authority based and consensus based. Each possesses a unique resolution to knowledge exchange hazards. These two governance forms correspond roughly to the distinct and fundamentally competing representations of hierarchy within the knowledge-based theories (see Demsetz 1988, Conner and Prahalad 1996 versus Arrow 1973, Kogut and Zander 1996). Both more efficiently cope with knowledge-formation hazards than markets and in the process support heuristic search more effectively than markets. Nonetheless, these two forms of hierarchy also differ significantly in their approach to supporting heuristic search, with one promoting knowledge transfer and the other essentially economizing on it. In our discussion below, we match these governance solutions and their support for directional or heuristic search to the complexity of problems.

4.1. Markets

Our conceptualization of the knowledge-based governance choice is very consistent with Hayek's (1945) view that the fundamental economic problem is determining the "best way of utilizing knowledge initially dispersed among all people" (p. 520). From this perspective, the manager must decide the best way to organize the creation of knowledge when the existing knowledge is widely dispersed among actors. In addressing this problem of governing knowledge, Hayek espoused the "marvel" of the market, arguing that markets provide "inducements which... make the individuals do the desirable things without anyone having to tell them what to do." Markets, thereby, "dispense with the need of conscious control" (Hayek 1945, p. 527). In our context of problem solving, market governance determines the path of search by decentralizing control to those in possession of valuable, specialized knowledge. In markets, prices provide high-powered incentives that motivate actors to search for solutions that both exploit and enhance their specialized knowledge and that can be encapsulated into saleable products. For a given problem, the path of search is merely an aggregation of the individual search decisions made by a set of actors autonomously choosing paths of search that they anticipate will lead to improved and marketable solutions. Thus, markets achieve unique combinations of knowledge using fully decentralized decision making.

Markets are ideally suited for governing directional search. As previously discussed, effective directional search involves individual actors altering design features based on knowledge they possess and then observing the resulting change in solution performance. Markets also provide high-powered incentives that motivate actors to pursue trials that expand their knowledge or the sale of products and services based on that knowledge. At the same time, however, markets provide weak support for knowledge sharing. They provide little protection against knowledge appropriation and no clear disincentives against strategic manipulation of search heuristics. Further, knowledge sharing requires the formation of a common language by which to communicate knowledge. Markets provide weak incentives to invest in the formation of such language.

However, when problems are decomposable and directional search is desired, knowledge sharing is largely unnecessary. Indeed, it is precisely on knowledge transfer that markets are designed to economize. As Hayek (1945) contends regarding markets: "The most significant fact about this system is the economy of knowledge which it operates, or how little the individual participants need to know [about other actors] in order to be able to take the right action" (p. 527). Thus, the efficiency of markets in governing decomposable problems is that they avoid knowledge-exchange hazards by

severely restricting knowledge exchange. Instead, markets offer powerful incentives for these individual actors to make optimal use of the knowledge. Each possesses incentives to engage in local search and thus to focus on trials that enhance their own accumulation of specialized knowledge.

Markets, however, provide a rather cumbersome and costly mechanism (i.e., the courts and classical contracting law⁷) for developing the heuristics necessary for searching complex solution landscapes. Independent actors governed by the market could in theory contractually agree to particular patterns of search, but disputes over performance in pursuing these search patterns would potentially require the intervention of courts and contract law. When solutions to complex problems are sought and a search heuristic is required to enhance the probability of discovering a valuable solution, conflicts among independent actors in developing this heuristic may be intense because of the knowledge hazards discussed above. Thus, the capacity for markets, even when supported by contracts, to manage heuristic search development is quickly exhausted. However, when markets are matched to decomposable problems where actors possessing independent knowledge sets can make independent design choices, commonly shared heuristics are unnecessary and such disagreements are unlikely to arise.

As governance devices, markets quickly fail as problems become more complex and demand heuristic search. Indeed, markets exacerbate the knowledge-exchange hazards that contaminate efforts to perform heuristic search. Markets' high-powered incentives discourage knowledge sharing and instead promote knowledge hoarding. Further, these incentives encourage individual actors to strategically shape heuristics and subsequent patterns of trials in ways that benefit them individually. Moreover, these incentives actively discourage investments in the formation of common language required for knowledge sharing and heuristic search.

Given the set of organizational features that commonly define markets, we posit markets are poorly suited for governing the process of solution search for a complex, nondecomposable problem. The necessary knowledge sharing is undermined by the absence of both appropriate incentives and shared language. Efforts to manage heuristic search through markets or contracts are limited by the high costs of resolving through legal means the disputes that would emerge as individual actors strategically appropriate knowledge and seek to opportunistically alter the trajectory of search in ways that benefit them individually. Thus, markets are quite efficient in guiding directional search for solutions to decomposable problems but quickly fail as problems become more complex and demand heuristic search. They fail because knowledge-exchange hazards

are poorly managed through contracts, and the high-powered incentives of markets simply promote these hazards.

4.2. Authority-Based Hierarchy

Many argue that the capacity to use centralized authority provides the distinguishing advantage of hierarchy over markets (e.g., Arrow 1974, pp. 68–70; Milgrom and Roberts 1990, p. 72; Simon 1951, p. 294; Williamson 1985). The efficiency properties of authority have been particularly highlighted in the context of problem solving and knowledge formation. As Arrow (1974) argues, "Authority, the centralization of decision making, serves to economize on the transmission and handling of knowledge" (p. 69).⁸ Demsetz (1988) concurs, noting the potential inefficiency of a market in which individual actors must first be convinced of the appropriateness of each action they take. As he contends, the virtue of hierarchy is that "[d]irection substitutes for education (that is, for the transfer of the knowledge itself)." Thus, in solving identified problems, the advantage of authority in hierarchy is the capacity for one actor to identify the precise order of trials, thereby circumventing the need to contractually manage the order of trials or to spontaneously achieve some type of consensus through extensive knowledge sharing.

Because Hayek argued similarly for the efficiency of markets in economizing on knowledge transfer, how do we reconcile the advocated knowledge-based efficiency of both autonomy in markets and authority in hierarchy? The answer, of course, centers on the problem for which efficient governance is sought. Markets are efficient when knowledge transmission is directed at solving decomposable problems where there is little interaction among design choices. When problems are decomposable, knowledge is embedded in products and services and knowledge transmission is largely limited to what can be contained in prices and bundled into products and services. However, bundling knowledge sets within a single firm and exercising authority to direct search becomes efficient when problems become complex and efficient search demands shared heuristics to sequentially guide design choices. In this circumstance, authority economizes on the extensive and costly knowledge sharing and education that would need to occur were the governance of solution search organized through a market interface. As described below, authority-based hierarchy more efficiently resolves the knowledge-based exchange hazards that arise in markets when problems are complex and demand knowledge sharing.

Authority-based hierarchy is composed of a set of complementary features distinct from those that support markets. It is well suited to supporting some degree of heuristic search, though at the expense of dampening incentives for directional search. The key feature of authority-based hierarchy is a central figure who

invests in understanding critical knowledge interactions and then composes suitable heuristics to guide search. For instance, the central figure defines and structures subproblems by imposing constraints or design rules on each subproblem, which greatly narrows the area of the solution landscape to be searched.⁹ Other features of the governance form also support this form of search. Within the boundaries of hierarchy, the courts exercise forbearance. They refuse to hear disputes internal to the firm, including disputes about the proper path of problem search. It is precisely this forbearance that grants authority within the bounds of hierarchy. Because of this forbearance, individual actors have limited capacity and motivation to strategically manipulate the path of search.

Other features of authority-based hierarchy also facilitate centrally directed heuristic search. The high-powered incentives that motivate and reward directional search and knowledge accumulation in markets are limited within hierarchies. Although workers can exit an employer and take with them some portion of the knowledge they accumulate, the incentive structure within the firm is quite different. Employees grant managers authority in exchange for a wage. This act severs the direct linkage between the knowledge they accumulate and the wage they receive. This departure from market incentives also encourages investments in communication channels and codes that lower the cost of a central authority assimilating, accumulating, and applying knowledge. However, authority-based hierarchy does not promote horizontal communication channels that would support broad knowledge sharing among peers. It is precisely such knowledge transmission and broad education for which the authority and direction of this governance form provide a substitute. Authority-based hierarchy is thus superior to markets in supporting heuristic search but inferior in supporting directional search due to incentive damage. It dampens incentives to strategically manipulate the path of search and provides a central authority that economizes on costly knowledge transmission.

The use of authority within hierarchy to guide search is ideally suited to a range of problems that are moderate in their complexity. Authority has significant limits that become apparent as problems become either too complex and nondecomposable or too simple and decomposable. An authority's heuristic is only valuable when it more efficiently guides search than actors searching independently, each basing their own search on their own specialized knowledge. However, managers have a propensity to meddle in subordinates' decisions (Williamson 1985), in part due to overconfidence in their own judgment (Bazerman 1994). As problems become more complex and nondecomposable, the cognitive limits of managers to develop useful search heuristics, combined with a manager's propensity to meddle, contaminates the efficiency of search. As Weber and other organizational theorists remind us, "Authority does not imply

expertise" (see Hammond and Miller 1985, p. 2). When authority is exercised in the absence of knowledge, it contaminates rather than accelerates search. Direction effectively substitutes for knowledge transfer only when managers have valuable knowledge with which to direct subordinates. When search dictates extensive recombining of knowledge, it is quite unlikely that authority is well suited to governing solution search. If many actors possess the required knowledge to define an optimal search heuristic and a central authority is incapable of absorbing that knowledge in a timely manner, then the development of a commonly shared search heuristic is required. Hayek noted precisely this limit to authority when he commented, "We cannot expect that [the problem of coordinating knowledge] will be solved by first communicating all this knowledge to a central board which, after integrating *all* knowledge, issues its orders" (1945, p. 524, emphasis in original).

If firms could constrain managers to intervene only when they possess knowledge that improves the direction of search, then the contributions of authority could extend to searching highly complex problems. However, managers appear to have great difficulty constraining their use of authority within hierarchies. In part, this reflects managers' inability to discern whether they possess information useful in directing search, and a strong bias to assume that they do. Managers are simply prone to meddle, wielding authority where they should delegate. Thus, we conclude that authority-based hierarchy provides an important advantage over markets in directing the search for solutions the more nondecomposable the problem. However, the limits to this governance solution are reached as the level of knowledge interactions escalate and the cognitive capacity of a single individual to assemble the required specialized knowledge reaches its limits.

The effectiveness of authority-based hierarchy relative to market-based control also diminishes rapidly as problems become more decomposable. When problems involve low interaction among design choices, using a manager's authority to dictate search may again damage search performance. An authority is simply unlikely to develop the necessary specialized knowledge required to direct solution search for all relevant subproblems and design choices. Again, managers are prone to meddling, directing the path of search when they lack the knowledge to effectively do so. The low-powered incentives that accompany hierarchy are also detrimental when autonomous searching by independent actors is desired. The weak incentives of hierarchy limit specialized knowledge formation and thereby constrain the efficiency of search. Further, the use of hierarchy is simply more costly than markets due to the costs of paying and supporting the central authority. If no clear value is added, then the added costs discourage its use. Thus, as displayed in Table 2, the combination of greater cost

Table 2 Comparative Analysis of Alternative Governance Mechanisms

	Market	Hierarchy	
		Authority based	Consensus based
Instruments			
Incentive intensity	+	0	0
Communication channels and codes	0	Vertical	Horizontal
Dispute resolution regime	Classical contract law	Forbearance-authority	Forbearance-relational
Performance			
Directional search	++	+	0
Heuristic search	0	+	++

and contaminated search renders authority-based hierarchy inefficient for decomposable problems.

4.3. Consensus-Based Hierarchy

Scholars also have argued for the knowledge-based efficiency of an alternative hierarchical form of governance, one that emphasizes extensive knowledge sharing and commonly shared search heuristics to guide decision making. It is this form of governance that Kogut and Zander (1992, 1996) emphasize when they describe the firm as providing a specialized social community that creates and transfers knowledge more quickly and efficiently than markets. In this conception of the firm, hierarchy is a device that supports knowledge transfer. However, this conception contrasts sharply with authority-based hierarchy, which minimizes knowledge transfer. Rather than hierarchy substituting direction for education as Demsetz (1988) and Conner and Prahalad (1996) suggest, hierarchy in this case substitutes education (i.e., knowledge transfer) for direction. Actors must first educate one another regarding knowledge relevant to defining a collective search heuristic.

This form of hierarchy achieves extensive knowledge transfer by enhancing the efficiency with which knowledge transfer occurs among actors within the firm. Extensive information sharing potentially permits actors within the firm to collectively agree on a path of search that is a consensus reflection of the specialized knowledge sets housed within the firm. However, as Arrow notes, consensus-based decision making, herein referred to as consensus-based hierarchy, is the near polar opposite of authority. Authority in this form of governance is used only to select a project, not the path of search. How then does a manager choose between the efficiency of these alternative internal governance choices of authority- and consensus-based hierarchy? The answer, of course, lies in the nature of the problem to which a manager seeks to find a solution. Consensus-based hierarchy is a potential solution to the failure of authority in governing heuristic search as problems increase in complexity.

Consensus-based hierarchy is supported by a distinct configuration of features that support active knowledge

exchange. As Arrow (1974) argues, consensus can substitute for authority so long as knowledge transfer is inexpensive and actors within the group have “a sufficiently overriding commonly valued purpose . . .” (pp. 69–70). Under these circumstances, each agent within the firm perceives the correct decision based on his or her knowledge and interests. Because interests and knowledge are commonly shared, consensus is spontaneous.

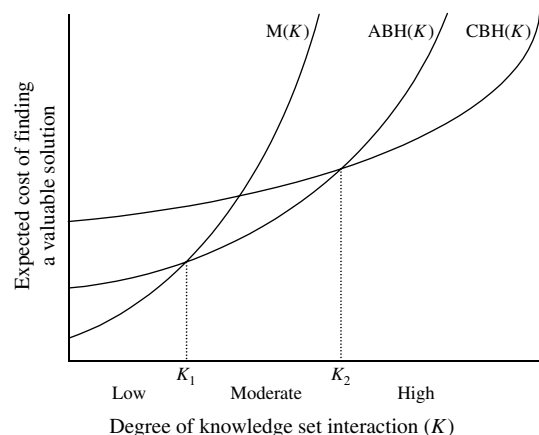
Achieving consensus, however, first demands extensive knowledge sharing, and knowledge sharing requires a commonly shared language. Hierarchy possesses a distinct advantage over markets in promoting the formation of shared language. Arrow, in his treatise *The Limits of Organization*, recognizes the distinct advantages of hierarchy in facilitating knowledge transfer through firm-specific language and identity. He contends that a primary component of firm-specific capital formation is “learning the information channels within a firm and the codes for transmitting information through them” (Arrow 1974, p. 56). He further notes that such investments are not only individually specific, but that they in aggregate represent “irreversible capital accumulation for the organization” that leads to organizations having “distinct identities” (pp. 55–56).¹⁰ Kogut and Zander similarly argue that communication costs are lower within the firm due to the presence of “higher order organizing principles” that “[establish] the context of discourse and coordination among individuals with disparate expertise” (Kogut and Zander 1996, p. 503). In their view, the boundaries of the firm define qualitative differences in the scope of social knowledge available to individuals. The shared identity that exists within a firm lowers the cost of communication and establishes “rules of coordination and influences the direction of search and learning” (Kogut and Zander 1996, p. 503).¹¹

Incentives and dispute resolution mechanisms within consensus-based hierarchy are also configured to support knowledge transfer and consensus decision making. Very low-powered incentives are essential to consensus, because such incentives encourage (or more accurately do not discourage) knowledge sharing.¹² High-powered incentives within the firm would encourage knowledge hoarding and strategic manipulation of search.

Dispute resolution in consensus-based hierarchies, of course, also differs from authority-based hierarchy. While forbearance by the courts remains the central characteristic of hierarchy compared to markets, dispute resolution in consensus-based hierarchy involves individual actors collectively deciding this path. Such group decision processes equate to Williamson's (1985, pp. 246–247) relational team and approximate Ouchi's (1980) clan form of organization. In this case, the firm "will engage in considerable social conditioning to help insure that employees understand and are dedicated to the purposes of the firm and employees be provided with considerable job security, which gives them assurance against exploitation" (Williamson 1985, p. 247). Such efforts build relationships among agents and facilitate the formation of a common identity. These relationships and shared identity in turn ease knowledge transfer, facilitate agreement, and discourage the exploitation by other agents of knowledge transferred into the firm (Allen 1967, Tushman 1978).¹³

While consensus-based hierarchies possess distinct advantages over authority-based hierarchies in facilitating heuristic search to highly complex problems, consensus fails in comparison to authority as problems diminish in complexity. The costs associated with supporting extreme levels of knowledge transfer are substantial and become unwarranted as problems diminish in complexity. Further, low-powered incentives constrain the motivation to develop specialized knowledge and dampen incentives for solution search. The scope of investment in shared language and socialization and the efforts involved in the transfer of knowledge can be excessive when problems are only moderately complex. Such investments and effort slow the accumulation of specialized knowledge necessary for effective directional search. Moreover, social attachments and idiosyncratic language that accompany consensus-based hierarchy can increase the cost of search by generating search heuristics that are limited in the knowledge sets that they incorporate. While firm-specific language and social attachments lower the cost of communication among coworkers, they encourage actors to oversearch their channels for knowledge rather than searching out knowledge not contained within the firm. Thus, not only are such social attachments costly to maintain (Hanson et al. 1999), but they also may misguide the process of search. Such oversocialization may reduce the infusion of new ideas and result in "parochialism and inertia" (Adler and Kwon 2002), or as Powell and Smith-Doerr (1994, p. 393) argue, the "ties that bind may become the ties that blind." Individual choices of search trajectory may become increasingly determined by the knowledge that workers and their close friends possess. Thus, social attachments may bias decisions toward continuing patterns of search that extensively utilize existing knowledge sets within

Figure 1 Discriminating Alignment Between Governance Alternatives and Problem Complexity



the firm and may limit the firm's capacity to search and absorb new forms of knowledge.¹⁴

This discussion suggests that the organizational costs of generating consensus are high. Hence, consensus-based hierarchy should only be adopted when the benefits of consensus are high—when problems are highly complex. Thus, for decomposable problems, consensus-based hierarchy is more costly than either market governance or authority-based hierarchy, as it is likely to take more time and trials and hence cost to identify a valuable solution. These costs increase, but slowly, as problems become increasingly nondecomposable, leaving consensus-based hierarchy as a more efficient governance solution for complex, nondecomposable problems.¹⁵

5. Discriminating Alignment

Our theory suggests an alignment between the search needs of problems and the search costs and performance of governance alternatives. We summarize the proposed alignment in Table 3 and Figure 1. Table 3 summarizes the set of relationships from problem type to governance choice. Figure 1 compares the stylized costs we described above for each governance alternative over a range of problems that vary in the degree to which knowledge sets interact to define the solution landscape. For simplicity of presentation, we represent the horizontal axis as a continuous measure of this degree of knowledge set interaction, holding the actual number of knowledge sets constant (this assumption equates to holding constant the N in an NK model and varying K). The vertical axis represents the expected cost of finding a valuable solution. Also, while the cost curves are all upward sloping, which might suggest that the absolute lowest cost (e.g., the cost of the market for low-interaction problems) is most desirable, it is important to remember that the highest peak of the landscape, and hence the value of the optimal solution, increases with

Table 3 Summary of Theoretical Relationships in the Problem-Solving Approach to Governance Choice

Type of problems	Degree of knowledge set interaction (K) among subproblems	Solution landscape characteristics	Appropriate search mode	Need for knowledge transfer	Knowledge hazards	Predicted governance form
Decomposable (A problem can be subdivided into subproblems, each of which draws from rather specialized knowledge sets.)	Low	Few peaks, gradual and smooth variations	Directional	Low	Low threat	Market
Nearly decomposable (A problem can be subdivided into subproblems that draw on specialized knowledge sets, but interaction among subproblems is sufficient so that value of design choices within one subproblem is not fully independent of the design choices made in another subproblem.)	Moderate	Intermediate levels of ruggedness	Directional and heuristic	Moderate	Moderate threat	ABH
Nondecomposable (Interactions among distinct knowledge sets are so extensive that any attempt to define subproblems and discover corresponding subordinate solutions offers no predictable improvement over randomly selecting trials.)	High	Rugged with abruptly jagged peaks and valleys	Heuristic	High	High threat	CBH

knowledge set interaction. Thus, the relevant issue is finding the lowest governance cost for a given problem.

Our hypothesis is that markets efficiently govern low-interaction problems, authority-based hierarchy efficiently governs moderate-interaction problems, and consensus-based hierarchy efficiently governs high-interaction problems. With low-interaction problems, markets promote specialization and directional search. However, knowledge-exchange hazards ensure that the cost of the market (indicated by $M(K)$) accelerates rapidly as the degree of knowledge set interaction increases. For low-interaction problems, the cost of authority-based hierarchy (indicated by $ABH(K)$) is greater than markets, reflecting the inability of authority to efficiently motivate the intense specialization required to facilitate directional search or for those in authority to constrain their predilection to meddle.¹⁶ As problems become increasingly complex, the costs of markets accelerate rapidly, reflecting their inability to cope with knowledge-exchange hazards. By contrast, the cost of authority-based hierarchy accelerates less quickly, leaving authority as the efficient governance choice for problems with moderate levels of interactions, which are between the points K_1 and K_2 . Nonetheless, authority-

based hierarchy's costs continue to accelerate, eventually leading to organizational failure as problems display high levels of interaction and the central authority lacks the relevant knowledge to guide search. The cost of consensus-based hierarchy (indicated by $CBH(K)$) is higher than the cost of authority-based hierarchy for problems with low and moderate levels of interaction. Thus, the organization provided by consensus-based hierarchy fails with these types of problems. Nonetheless, the costs of consensus do not accelerate until it is used to organize problems with high levels of interaction, which makes it the economic choice for problems with levels of interaction greater than K_2 .

6. Discussion and Conclusion

Our theory addresses several deficiencies in the existing literature on the knowledge-based view of the firm. In particular, while more recent work (Conner and Prahalad 1996, Kogut and Zander 1996) explains from the knowledge-based view why firms exist by articulating the knowledge-based advantages of hierarchy, our theory focuses on *when* firms exist and in what form. Our theory articulates the knowledge-based advan-

tages and disadvantages of both markets and hierarchy and matches these to desired types of knowledge formation. In addition to answering the question of what the knowledge-based reasons for a firm are, our theory informs the boundary choice between and among firms. Equally important, our theory provides some basis for reconciling and choosing between the two distinct and contradictory knowledge-based explanations for the existence of firms. While one theory claims that firms exist to economize on knowledge transfer (Demsetz 1988, Conner 1991, Conner and Prahalad 1996) and the other argues firms facilitate knowledge transfer (Arrow 1974; Kogut and Zander 1992, 1996; Nahapiet and Ghoshal 1998), we contend that a manager has three distinct governance choices in supporting knowledge formation: markets, authority-based hierarchies, and consensus-based hierarchies. These latter two forms correspond roughly to the competing conceptions of hierarchy within the literature with a knowledge-based view of the firm.

It is important to note that our theory is inherently based on a probabilistic assessment of solution discovery. Correctly aligning governance with problem complexity does not guarantee the rapid discovery of a valuable solution. Nonetheless, we expect that over a large number of solution searches, efficient alignment yields superior performance by enhancing the probability of discovering a valuable solution. Our theory, of course, assumes that there are multiple valuable solutions. On reaching a valuable solution, the firm must weigh the probable gain from further search against the added cost of further search.

While our theory aligns governance forms to problems, our theory was developed by implicitly assuming that successively chosen problems are independent. Clearly they are not. Interdependence of problems has important implications for governance choice. Managers likely choose new problems reflecting upon the composition of knowledge sets already in the firm's possession. Therefore, for any given knowledge set, managers are concerned not only with its proper governance, but also with its development and protection. These latter concerns about development and protection place constraints on governance choices not addressed in our theory. Thus, selling off a knowledge set to accommodate its appropriate governance for a highly decomposable problem may leave the firm poorly positioned to govern a subsequent or concurrent nondecomposable problem that requires this knowledge set. Nonetheless, for a knowledge set housed within the boundaries of the firm with wide application to a set of problems pursued by the firm, important governance differences across problems can still occur. Thus, those possessing a particular knowledge set can be tightly integrated with a team in pursuing a nondecomposable problem, receive simple design instructions from an authority in pursuing a

nearly decomposable problem, and be granted enormous latitude in contributing to a decomposable problem.

While our theory proposes a rather static alignment between problem choice and governance choice, the impetus for dynamic changes in the composition of the firm are inherent in the model. In our knowledge-based theory of the firm, firms shift their boundaries in response to changes in the problems that they address. Moreover, problems have life cycles. A problem may begin as a nondecomposable task and necessitate heuristic search using consensus-based hierarchy. However, once a valuable solution is achieved and the resulting product widely sold, refined solutions to the original problem may be a decomposable problem that warrants governance through authority-based hierarchy or, with further time, the market. For instance, this may be the case with the emergence of so-called dominant designs. Such evolutionary changes in the nature of the problem to be solved imply dynamic changes in organization structure when viewed through the lens of the problem-solving perspective. A shift from nondecomposable to decomposable problems implies a shift in organizational form from consensus-based hierarchy to markets. Yet, as described above, management may instead take advantage of knowledge and capabilities within the firm and seek solutions to more complex problems. Such strategic and dynamic choices represent a substantial arena for further inquiry.

Obvious questions that arises from our theory are how the theory can account for large organizations with many levels of hierarchy or how it can account for complex linkages among multiple firms. Our theory is developed by discussing problems that may be separable into subproblems if complexity is not too severe. While not fully unpacked here, we anticipate that our theory can be used in a recursive way in the sense that problems with a large N might be decomposed into subproblems, which in turn might be decomposable into more subproblems, etc., until the problems are no longer decomposable. Then our theory can be applied to each level of subproblem to predict not only the organization of the lowest level of subproblems, but also entire branches of subproblems. If this line of thinking is further developed, our theory might provide deeper insights into the structure of firms as well as into the complex set of interorganizational relationships that characterize many industries.

Our theory highlights both the perils and virtues of authority in managing knowledge formation. Managers must recognize that the domain of problems for which their direction is of value is bounded on two sides. Effective knowledge management requires managers to recognize both limits. On the one hand, when problems are quite decomposable, managers must resist the temptation to assume value in their direction by integrating relevant knowledge sets. Instead, they must trust

the market's intense motivation to specialize and guide directional search when problems are decomposable. Managers must recognize that cognitive constraints limit their capacity to possess the specialized knowledge to guide directional search and integration will simply dull the incentives of those specialists who do possess the required knowledge. On the other hand, when problems are nondecomposable, managers must also resist the temptation to assume value in their directing the path of search. Again, the manager must recognize that cognitive constraints limit his or her capacity to understand the wide-ranging knowledge interactions required to develop heuristics useful in guiding search. Instead, the manager must trust that a culture of widespread knowledge sharing and consensus decision making is the organizational approach most likely to yield a valuable solution. Therefore, depending on the set of problems pursued by the firm, those gaining from direction (substituting for education) may be quite limited.

Acknowledgments

The authors would like to thank Carliss Baldwin, Michael Leiblein, Glenn MacDonald, Jeff Macher, Laura Poppo, Lee Fleming, Brian Silverman, and seminar participants at Berkeley, Dartmouth, Copenhagen Business School, Harvard Business School, and Ohio State for helpful comments. They also thank Anne Miner, Jim Wade, Claudia Bird Schoonhoven, and anonymous referees at *Organization Science*.

Endnotes

¹By a complex system, Simon roughly meant a system “made up of a large number of parts that interact in a non-simple way. In such systems, the whole is more than the sum of the parts, not in an ultimate, metaphysical sense, but in the important pragmatic sense that, given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole” (1962, p. 486). Kauffman's *NK* model provides a simple mathematical construction that arguably provides one parameterization of Simon's characterization. Like Levinthal (1997), these two parameters in our framework combine to describe a solution landscape. We develop our discussion implicitly assuming that *N* is nontrivial and constant (although we relax this view toward the end of our theory development). Thus, our focus is on variations in *K*. Following Simon, we develop our logic using a typology of complexity instead of continuous variables because the typology captures the central issues. Identifying the precise degree of interaction *ex ante* is more likely to confront cognitive limits than a typology when moving toward operationalization. Nonetheless, also we resort to employing the continuous variable *K* when describing how the costs of governance alternatives vary with complexity to make clear the relationships we predict.

²*Ex ante* identification that problems are either decomposable or nondecomposable may be relatively straightforward. Yet the degree to which a problem is nearly decomposable raises a question of operationalization. While this operationalization question deserves more attention than can be given here, we will point out that others have grappled with this issue. From a

mathematical perspective, there is now a substantial literature on nearly decomposable systems and matrices. This mathematical definition forms the basis of implementable measures of near decomposability. For instance, McCord and Eppinger (1993) develop a task structure matrix for the design process of a laptop computer. This system, which is nearly decomposable, forms the basis for much on Baldwin and Clark's (2000) discussion on degrees of modularity. While surely it is the case that additional work is needed to provide refined measures of near decomposability, this prior work suggests that implementable measures can be constructed.

³Simon similarly identified two basic kinds of trials selectivity. One kind of selectivity, feedback, is where “various paths are tried out, the consequences of following them are noted, and this information is used to guide further search” (1962, p. 473). The second kind of trial selectivity is based on previous experience, which greatly reduces search based on trial and error. While this classification resonates with our own, Simon's terminology is less descriptive because feedback ultimately is used for both kinds of selectivity and because it does not relate to the theory of algorithms, which provides mathematical approaches for solving complex problems. In contrast, the terminology we introduce assumes feedback in both kinds of trial selectivity and parallels two broad classifications of solution algorithms. Directional search parallels hill-climbing or gradient methods of maximization, whereas heuristic search parallels combinatorial methods and probabilistic methods such as TABU, simulated annealing, and genetic algorithms for solving complex problems. Of course, we do not anticipate that such heuristics yield tractable algorithmic solutions for problems in which we are interested.

⁴We thank an anonymous referee for pointing out that the terms “directional search” and “heuristic search” are used here in a more narrow sense than is found in the broader organizational literature. The key difference between our usage and that in the broader literature is our unit of analysis. In particular, these terms are used here in reference to searching for a solution to a particular problem.

⁵It is well known that many hybrid organizational structures fall between the polar modes of market and hierarchy and that they may offer a variety of intermediate governance mechanisms. We believe that considering these alternatives here complicates theory development without the promise of adding deep insights to our main hypothesis. Thus, we anticipate that the theory development offered herein will be useful in future research that considers the knowledge-creation costs and competencies of hybrid organizational structures.

⁶Arrow and Demsetz are exceptions here.

⁷Classical contracting law refers to—

“thick” markets [that] are ones in which individual buyers and sellers bear no dependency relation to each other. Instead, each party can go its own way at negligible cost to another Such transactions are monetized in extreme degree; contract law is interpreted in a very legalistic way; more formal terms supersede less formal should disputes arise between formal and less formal features (e.g., written agreements vs. oral amendments), and hard bargaining, to which the rules of contract law are strictly applied, characterizes these transactions (Williamson 1995; 1991, p. 95).

⁸Demsetz (1988) and Conner and Prahalad (1996) make the identical point: For instance, authority provides a “low cost

method of communicating between specialists and the large number of persons who either are nonspecialists or who are specialists in other fields” (Demsetz 1988, p. 172).

⁹Some have referred to such heuristics as “design rules.” Baldwin and Clark (2000) provide an overview of this literature and describe the role of design rules set by a central authority to coordinate design of complex products. More generally, their analysis discusses the relationship between modularity and organizational structure.

¹⁰Hierarchy may not be the only institution able to support investments in communication channels and codes. For instance, channels and codes may emerge for groups of agents such as professional associations, regional cultures, industrial districts, etc. Such communication channels and codes may shift the relative costs and benefits of alternative governance modes. Nonetheless, we maintain that hierarchy remains unique in its ability to choose channels and codes and that the codes have the ability to provide greater richness for knowledge transfer than those that arise endogenously from other institutions.

¹¹Monteverde (1995, p. 1,628) similarly contends that organizations are distinguished by their unique communication codes, by the “specific business dialect[s] their members use.” He contends that firm-specific language is the “defining essence of what we recognize as organizations.” Thus, firm boundaries, according to this logic, “should congeal around transactions rich in such technically necessary, unstructured dialog” (Monteverde 1995, p. 1,629).

¹²While both authority-based and consensus-based hierarchy rely on incentive structures that are much lower powered than those available in the market, we anticipate that they are not identical. Incentives associated with consensus-based hierarchy are more likely than those associated with authority-based hierarchy to be structured to encourage collective action.

¹³Strong social ties that emerge within hierarchies facilitate such knowledge transfer and consensus. Research on technical communication highlights the importance of close social ties in facilitating knowledge flows within organizations (Allen 1967, Tushman 1978). Similarly, the product innovation literature suggests that close, frequent interactions are necessary for timely integration of knowledge (Clark and Fujimoto 1991, Henderson and Cockburn 1994). In a study of knowledge transfer around best practices, Szulanski (1996) finds that knowledge transfer is facilitated by positive personal relationships. Hansen et al. (2001) find that rich networks of social ties are instrumental in pursuing exploratory tasks, which demand extensive knowledge transfer.

¹⁴Discontinuing patterns of local search often involves shutting down a particular research or development stream. Such disruption requires severing existing social commitments and damaging social attachments. However, loyalty to such social commitments then clouds the process of search. Thus, an inability to sever or ignore social attachments may lead to inefficient rather than efficient search.

¹⁵This does not mean that the net benefit of consensus-based hierarchy is positive for all nondecomposable problems. Indeed, for any problem where $K > 0$ the number of interactions and the number of relevant knowledge sets may be sufficiently large to increase the cost of search beyond the value of a potential solution (i.e., in terms of an NK model, assuming $K > 0$, then the expected value of finding a solution decreases

with N). Moreover, with consensus-based hierarchy, the larger the group of specialists, the more difficult it will become to maintain a commonality of purpose, thereby increasing the cost of arriving at consensual decisions.

¹⁶By efficiency we mean that the economic benefit from the speed and likelihood of finding a valuable solution, less the cost of doing so, will be the greatest, on average, when problem types are paired with governance modes as hypothesized. Pairings that differ from these predictions will be less efficient, thereby increasing the expected cost and the expected time, or decreasing the likelihood of finding a valuable solution. Because knowledge formation is an uncertain activity, our notion of efficiency is a probabilistic assessment of net benefits including both technical and behavioral costs of knowledge formation.

References

- Adler, P. S., S.-W. Kwon. 2002. Social capital: Prospects for a new concept. *Acad. Management Rev.* **27**(1) 17–40.
- Allen, T. J. 1967. Communications in the research and development laboratory. *Tech. Rev.* **70**(1) 31–37.
- Argyres, N. 1996. Evidence on the role of firm capabilities in vertical integration decisions. *Strategic Management J.* **17** 129–150.
- Arrow, K. J. 1973. Information and economic behavior. *Stockholm, Federation Swedish Indust.* **28**. Reprint of Nobel Prize paper.
- Arrow, K. J. 1974. *The Limits of Organization*. Norton and Co., New York.
- Arrow, K. J., F. H. Hahn. 1971. *General Competitive Analysis*. Holden-Day, San Francisco, CA.
- Baldwin, C. Y., K. B. Clark. 2000. *Design Rules: The Power of Modularity*. MIT Press, Boston, MA.
- Barney, J. B. 1984. Strategic factor markets: Expectations, luck, and business strategy. *Management Sci.* **32**(10) 1231–1241.
- Bazerman, M. H. 1994. *Judgment in Managerial Decision Making*. Wiley, New York.
- Clark, K. B., T. Fujimoto. 1991. *Product Development Performance: Strategy, Organization, and Management in the World Auto Industry*. Harvard Business School Press, Boston, MA.
- Conner, K. R. 1991. A historical comparison of resource-based theory and five schools of thought within industrial organization economics: Do we have a new theory of the firm? *J. Management* **17**(1) 121–154.
- Conner, K. R., C. K. Prahalad. 1996. A resource-based theory of the firm: Knowledge versus opportunism. *Organ. Sci.* **7**(5) 477–501.
- Cyert, R. M., J. G. March. 1963. *A Behavioral Theory of the Firm*. Blackwell Business, Cambridge, MA.
- Debreu, G. 1959. *Theory of Value*. Wiley, New York.
- Demsetz, H. 1988. The theory of the firm revisited. *J. Law Econom. Organ.* **4**(1) 141–162.
- Fiol, C. M. 1994. Consensus, diversity, and learning in organizations. *Organ. Sci.* **5**(3) 403–420.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47**(6) 117–132.
- Fleming, L., O. Sorenson. 2000. Science as a map in technological search. Manuscript, Harvard Business School, Boston, MA.
- Foss, Nicolai J. 1996a. Knowledge-based approaches to the theory of the firm: Some critical comments. *Organ. Sci.* **7**(5) 470–476.

- Foss, Nicolai J. 1996b. More critical comments on knowledge-based theories of the firm. *Organ. Sci.* 7(5) 519–523.
- Gavetti, G., D. E. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* 45 113–137.
- Grant, R. M. 1996. Toward a knowledge-based theory of the firm. *Strategic Management J.* 17 109–122.
- Hammond, T., G. Miller. 1985. A social choice perspective on expertise and authority in bureaucracy. *Amer. J. Political Sci.* 29 1–28.
- Hansen, M. T., J. Podolny, J. Pfeffer. 2001. So many ties, so little time: A task contingency perspective or social capital. *Res. Sociology Organ.*
- Hayek, F. 1945. The use of knowledge in society. *Amer. Econom. Rev.* 35(September) 519–530.
- Heiman, B., J. A. Nickerson. 2002. Towards reconciling transaction cost economics and the knowledge-based view of the firm: The context of interfirm collaborations. *Internat. J. Econom. Bus.* 9(1) 97–116.
- Henderson, R. M., K. B. Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* 35 9–30.
- Henderson, R. M., I. Cockburn. 1994. Measuring competence? Exploring firm effects in pharmaceutical research. *Strategic Management J.* 15(Special Issue) 63–84.
- Hoetker, G. 2003. Do modular products lead to modular organizations? Manuscript, Department of Business Administration, University of Illinois, Urbana-Champaign, IL.
- Kauffman, S. 1993. *The Origins of Order*. Oxford University Press, New York.
- Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities and the replication of technology. *Organ. Sci.* 3(August) 383–397.
- Kogut, B., U. Zander. 1996. What firms do: Coordination, identity, and learning. *Organ. Sci.* 7(5) 502–518.
- Levinthal, D. 1997. Adaptation on rugged landscapes. *Management Sci.* 43 934–950.
- Macher, J. 2001. Vertical disintegration and process innovation in semiconductor manufacturing: Foundries v. integrated device manufacturers. Manuscript, McDonough School of Business, Georgetown University, Washington, D.C.
- Madhok, A. 1996. The organization of economic activity: Transaction costs, firm capabilities, and the nature of governance. *Organ. Sci.* 7(5) 577–590.
- Mahoney, J. 2001. A resource-based theory of sustainable rents. *J. Management* 27(6) 651–666.
- March, J. G., H. Simon. 1958. *Organizations*. John Wiley, New York.
- McCord, K., S. Eppinger. 1993. Managing the integration problem in concurrent engineering. Working Paper 3594-93-MSA, Sloan School of Management, MIT, Cambridge, MA.
- Milgrom, P., J. Roberts. 1990. Bargaining costs, influence costs, and the organization of economic activity. J. Alt, K. Shepsle, eds. *Perspectives on Positive Political Economy*. Cambridge University Press, New York, 57–89.
- Monteverde, K. 1995. Technical dialog as an incentive for vertical integration in the semiconductor industry. *Management Sci.* 41(10) 1624–1638.
- Nahapiet, J., S. Ghoshal. 1998. Social capital, intellectual capital, and the organizational advantage. *Acad. Management Rev.* 23 242–266.
- Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Belknap Press, Cambridge, U.K.
- Ouchi, W. 1980. Markets, bureaucracies, and clans. *Admin. Sci. Quart.* 25 129–141.
- Powell, W. W., L. Smith-Doerr. 1994. Networks and economic life. N. J. Smelser, R. Swedberg, eds. *The Handbook of Economic Sociology*. Princeton University Press, Princeton, NJ, 368–402.
- Prahalad, C. K., G. Hamel. 1990. The core competence of the corporation. *Harvard Bus. Rev.* (May–June) 79–91.
- Rumelt, R. 1984. Toward a strategic theory of the firm. *Competitive Strategic Management* 556–570.
- Schumpeter, J. 1939. *Business Cycles*. McGraw-Hill Book Company, Inc., New York.
- Simon, H. A. 1945. *Administrative Behavior*. The Free Press, New York.
- Simon, H. A. 1951. A formal theory of the employment relationship. *Econometrica* 19(3) 293–305.
- Simon, H. A. 1962. The architecture of complexity. *Proc. Amer. Philos. Soc.* 106(6) 467–482.
- Simon, H. A. 1991. Bounded rationality and organizational learning. *Organ. Sci.* 2(1) 125–134.
- Szulanski, G. 1996. Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management J.* 17 27–43.
- Teece, D., G. Pisano, A. Shuen. 1997. Dynamic capabilities and strategic management. *Strategic Management J.* 18(7) 509–533.
- Tushman, M. L. 1978. Technical communication in R&D laboratories: The impact of project work characteristics. *Acad. Management J.* 21 624–644.
- Walsh, J. P. 1995. Managerial and organizational cognition: Notes from a trip down memory lane. *Organ. Sci.* 6(3) 280–321.
- Walsh, J. P., L. Fahey. 1986. The role of negotiated belief structures in strategy making. *J. Management* 12(3) 325–338.
- Weitzman, M. 1996. Hybridizing growth theory. *Proc. Amer. Econom. Association* (May) 207–212.
- Wernerfelt, B. 1984. A resource-based view of the firm. *Strategic Management J.* 12 75–94.
- Williamson, O. E. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. The Free Press, New York.
- Williamson, O. E. 1991. Comparative economic organization: The analysis of discrete structural alternatives. *Admin. Sci. Quart.* 36(June) 269–296.
- Williamson, O. E. 1996. *The Mechanisms of Governance*. Oxford University Press, New York.
- Williamson, O. E. 1999. Strategy research: Governance and competence perspectives. *Strategic Management J.* 20(12) 1087–1103.