

Expanding Knowledge-Flow Theory through Computational Analysis of Knowledge Specialization

by

Mark E. Nissen, Ryan J. Orr and Raymond E. Levitt

Proceedings of the Academy of Management 2007 Conference, Philadelphia, Pennsylvania, August 2007

Abstract

In this article we investigate empirically the theoretical split between emphases upon specialist versus trans-specialist knowledge in the organization—or more generally between exploitation and exploration. This split divides knowledge-flow theory at present, and hence represents an important issue for Knowledge Management (KM). We review the relevant literature, articulate hypotheses, and employ computational experimentation to test them empirically. Our findings provide novel, insightful understanding of the factors that contribute toward understanding the relative balance between specialist versus trans-specialist knowledge in particular, and exploitation versus exploration more generally. We offer three main contributions: 1) we critique extant theory relating to the substitutability of specialist and interspecialist knowledge; 2) we offer new conceptual insight and empirical evidence concerning substitutability of these knowledge types in the organization; and 3) we demonstrate the empirical power of computational experimentation to examine KM questions of both theoretical and practical interest.

INTRODUCTION

Knowledge is central to effective management and performance in the modern organization. In a resource-based view of the firm, knowledge is considered by many scholars (e.g., Cole 1998, Grant 1996, Spender 1996) to represent the most important resource a firm can hold. Indeed, knowledge offers a source of sustainable comparative advantage (Drucker 1995). Toward this end, the field of knowledge management (KM) has undertaken substantial research on organizational knowing (e.g., Cook and Brown 1999) and learning (e.g., Levitt and March 1988), which comprise two key aspects of emerging knowledge-flow theory (e.g., see Nissen 2006).

But the knowledge-flow theory remains divided as to the relative importance of knowing versus learning in the organization. On the one hand, organizational knowing—involving knowledge in action—is key to organizational performance, and focuses generally exploitation of ex-

tant knowledge to perform useful work (i.e., workflow focused). On the other hand, organizational learning—involving knowledge in motion—is key to organizational performance also, but focuses generally on exploration of new knowledge to increase knowledge stocks (i.e., knowledge-flow focused). The balance between exploitation and exploration is noted as critical in the literature (e.g., March 1991), and insight into identifying, achieving and maintaining such balance can elucidate the KM researcher and practitioner alike.

In this article, we examine such balance between exploitation and exploration in the context of emphasis upon specialist versus trans-specialist knowledge in the organization. By developing and maintaining specialist capabilities, the organization realizes deep knowledge that can be exploited directly to work and problem solving within a particular specialist domain. This is the basis for many functional organizations, for instance, where engineering design knowledge is pooled within one functional organization, and manufacturing production knowledge is pooled within another. Development and maintenance of specialist knowledge as such is relatively efficient and economical also (Grant 1996), hence specialization as such is quite attractive to the knowledge (and financial) manager.

However, such specialist knowledge tends to be relatively narrow (i.e., limited to a single domain of application), and when emphasized as such, it can inhibit learning across functional organizations. This is the basis for many cross-functional teams, for instance where engineering and manufacturing people work together to design producible artifacts. Through development of trans-specialist knowledge (e.g., where designers learn about production, and producers learn about design), the combined design-production process can become more effective, hence trans-specialist learning is quite attractive to the knowledge (and financial) manager also.

The question is, how can the knowledge manager determine the extent to which the organization should emphasize specialist versus trans-specialist knowledge in particular, or more generally, how can exploitation and exploration be balanced via KM? In this article, we address such question by building upon recent theory on the topic (Postrel 2002), and employing computational experimentation to examine the underlying factors that contribute toward understanding the relative balance between specialist versus trans-specialist knowledge in particular, and exploitation versus exploration more generally.

The use of computational experimentation to address questions of interest and importance to KM is relatively new (e.g., see Nissen and Levitt 2004), but is demonstrating considerable power—particularly where combined with other research methods. Through the research described in this article, we strive not only to elucidate the specialist versus trans-specialist balance question—and to demonstrate the use and utility of computational experimentation in this KM domain—but to encourage other researchers, using other, complementary research methods, to help us to resolve the existing divide in knowledge-flow theory. The following section presents theoretical background and hypotheses, after which we describe our computational model, present the results, and draw conclusions.

THEORETICAL BACKGROUND & HYPOTHESES

We begin by summarizing key aspects of Postrel's (2002) analytic model, as a basis to craft a null hypothesis. We then summarize relevant theory, and formulate hypotheses that complement, extend and compete with such analytic model.

Islands of Shared Knowledge

In order to investigate the relative contributions of specialist and trans-specialist knowledge to the performance of an organization engaged in product development, Postrel (2002) derived a

two-parameter production function: $M = zh/[zh + (1-z)(1-h)]$. The expected payoff is represented by M , which corresponds to a distribution of design outcomes. Trans-specialist capability is represented by a single parameter h , which specifies the design actor's probability of sending a producible design to its downstream counterpart in manufacturing; as h increases so does the likelihood that a specific design can be produced. Likewise, specialist capability is represented by a single parameter z , which specifies a manufacturing actor's probability of making a successful product; as z rises so does the likelihood that a range of designs can be manufactured successfully.

Notice that the production function returns the value 1, the maximum expected payoff, when *either* parameter z or h is 1; that is, when either knowledge type (i.e., specialist or trans-specialist) is high, maximum payoff is assured for any value of the second knowledge type greater than zero. Notice this production function also returns the value 1 when *both* parameters z and h are 1; that is, nothing is gained when both knowledge types are at high levels. Notice further the production function returns the value 0 (i.e., no payoff) if zero values are allowed for either parameter z or h ; that is, no level of one knowledge type is sufficient to compensate for a complete absence of the other. According to the production function, these payoff relations are guaranteed.

Based on the results of the analytical model, Postrel concludes (pg. 310) that "Specialist capability and trans-specialist understanding are primarily substitutes in the design production function." This constitutes the null hypothesis of our present study,

H0: Specialist and trans-specialist knowledge can be substituted for one another without affecting performance.

Hypothesis Development

As with any nontrivial representation, the analytical model described above necessarily involves numerous assumptions. Some of these assumptions are realistic and helpful to make the mathematics tractable, and their inclusion seems relatively innocuous. Yet several others, which drive the principal results obtained, are questionable in terms of microeconomic and organization theory. Here we discuss seven assumptions, to which we take greatest exception, and articulate hypotheses that complement, extend and compete with results of the analytical model.

1. Unidimensional performance. The analytic model suggests that performance in the design production process can be measured along a single dimension, labeled the “expected payoff” (Postrel, 2002). However, studies of product development suggest instead that *performance* is a multi-faceted, omnibus concept with many trade-offs between sub-dimensions such as *cost*, *time*, *component quality* and *product integration quality* (e.g. Smith and Reinertsen 1991, Bayus 1997).

This assumption of unidimensional performance is problematic, because it obscures finer-grained relationships between the two knowledge types and the individual sub-dimensions of performance. For example, studies of product development indicate that trans-specialist knowledge—often labeled under other rubrics such as “system-wide understanding”, “cross-functional expertise” or “interspecialty capability”—is necessary in order to avoid the kinds of catastrophic interfunctional glitches that lead to product integration failure (Allen and Cohen 1967, Allen 1977, Cooper 1979, Carrol 1998, Hoopes and Postrel 1999). However, organizations like NASA provide evidence that increasing trans-specialist knowledge with the intention of producing a design with greater manufacturability can have the consequence of increasing design cycle time as more constraints are placed on requirements (Vaughan 1990). Finally, with respect to specialist knowledge, studies of human learning reveal that task-repetition and functional-

specialization reduce the likelihood of task-specific errors (Staddon 1983). Based on these insights from theory, we draw our first set of hypotheses:

H1a: Trans-specialist knowledge has benefits in reducing product integration risk, but does not contribute to a reduction in individual component risk;

H1b: Increasing trans-specialist knowledge may actually increase—and not decrease—the schedule duration of the design activity.

H1c: Specialist knowledge has benefits in reducing component risk, but does not contribute to a reduction in overall product integration risk.

2. Hyper-substitutability of knowledge types. The production function derived in the analytical model from above is not what one would expect *ex-ante* in terms of neoclassical microeconomics (see Pindyck and Rubinfeld 1998, Samuelson 1974). In particular, the production function derived in the analytical model produces many isoquants that are concave to the origin (Postrel 2002: 311). The implication is, the more of one type of knowledge used in a process, the greater its marginal productivity becomes, hence rendering the other type relatively less valuable. In the case of product development and the analytical model, as an organization gains substantial specialist knowledge, for instance, gaining additional trans-specialist knowledge becomes decreasingly valuable, and vice versa. This denotes a condition of *hypersubstitution*: not only do two (or more) input factors of production—*specialist* and *trans-specialist knowledge* in this case—represent economic substitutes for one another, but normative decision rules and cost assumptions drive the level of one (or more) input factor toward zero. This is a central result of the analytical model.

In contrast, a more traditional class of production functions (e.g., the Cobb-Douglas function) reflect isoquants that are convex to the origin. Here levels of one or more input factors are traded off against other factors, depending upon their relative marginal rates of technical substi-

tution (MRTS) and marginal costs (MC). Only in extreme cases would one input factor be driven merely by decision-making rules toward zero level in such arrangements. This represents a condition of *substitution*: input factors such as specialist knowledge and trans-specialist knowledge represent economic substitutes. Normative decision rules would suggest some combination of the two knowledge types, with specific levels determined by MRTS and MC.

Moreover, many, diverse input factors can also represent economic complements: using more (or less) of one input factor benefits from using more (or less) of the other(s) also. To the extent that specialist knowledge and trans-specialist knowledge complement one another, then one can argue, economically, that *more of both* should be used in production; that is, levels of the input factors should covary according to normative decision rules. This can be the case even for classical productions functions such as the Cobb-Douglass (e.g., with nonnegative exponents suggesting supermodularity; see Milgrom and Roberts 1995, p. 183). Indeed, many, common, modern manufacturing approaches (e.g., lean manufacturing, flexible production, mass customization, design for manufacturing) seek to leverage just such complementary relations between input factors (Milgrom and Roberts 1990). Hence arguing microeconomically, we see little reason to assume the kind of hypersubstitution implied by the analytical model. This leads to our second set of hypotheses:

H2a: Specialist and trans-specialist knowledge used in product development may represent economic complements.

H2b: Using predominately specialist knowledge in product development leads to inferior performance than when appreciable levels of trans-specialist knowledge are employed also.

3. “One-size-fits-all” competitive strategy. The analytic model ignores the fact that different firms have different competitive strategies for how they position their products in the market-

place. Yet the strategy literature indicates that firms guided by a cost-leadership strategy maximize their payoff by producing low-cost products; whereas firms with a product-differentiation strategy excel by offering products of superior quality, novelty, technological sophistication, and like differentiating attributes (Porter 1980: 35-37). Porter explains that the effective execution of these generic strategies requires very different skills, resources and organizational capabilities. Competing on product cost alone is usually achieved through tight cost-control and the efficiencies of high-volume production by specialized factory workers (e.g. modular home production). Quite the opposite, product differentiation is attained generally through a creative design process with strong cooperation across specialties (e.g. custom-home production). Thus, a firm's competitive strategy affects the marginal value of the different knowledge types. This leads to our third set of hypotheses:

H3a: The greater the importance of low-cost production, the greater the importance of specialist relative to trans-specialist knowledge.

H3b: The greater the importance of product quality, the greater the importance of trans-specialist versus specialist knowledge.

4. Equal modularity across products. The analytical model assumes that all product designs exist with an equivalent level of modularity between components. But research shows that product modularity is highly variable across products and in different phases of product evolution (Schilling 2000, Baldwin & Clark 2000). Newer products tend to have larger numbers of more quickly evolving, component parts with intricate and changing interfaces, while more mature products tend to have fewer, more standardized interdependencies (Utterback 1996). As product modularity changes, so does organizational modularity (Sanchez and Mahoney 1996, Ethiraj and Levinthal 2004) and the marginal value of trans-specialists in the organization. When modularity is low, interspecialists are needed urgently to resolve negatively interacting sub-goals between

components through the mechanisms of coordination and mutual adjustment (Thompson 1967; March and Simon 1968). But as modularity rises, interspecialists become less important. More formally,

H4: As product modularity rises, the marginal benefit of trans-specialist knowledge falls.

5. Equal complexity across components. The analytic model makes no distinction between the levels of complexity of different component parts in the value chain. Yet complexity is a key factor thought to influence cognitive task performance (Locke and Latham 1990, Simon 1981), and it is commonly believed that as complexity rises, it is necessary to employ increasingly talented specialists to produce components of a given level of quality (Campbell 1988). Thus, we draw a fifth hypothesis,

H5: As product component complexity rises, the marginal benefit of specialist knowledge also rises.

6. Sequential interdependence across work processes. Product development in the analytical model involves only sequential interdependence between the design and manufacturing actors. Although some product development processes in practice do fit this sequential characterization, very few complex products are developed in this “over the wall” fashion. Most are conducted with some level of concurrency between the stages of the value chain (e.g., design-build, prototyping).

From a performance standpoint, while concurrency brings a reduction in overall project duration, it also creates reciprocal interdependence between the design and manufacturing activities (Thompson 1967), which requires additional cost and effort to coordinate (March and Simon 1958). This introduces a time-cost tradeoff (Mustafa and Murphree 1989, Russel and Ranasinghe

1991), for the cost of rework and increased product failure risk can exceed any savings gained through concurrent development (Stacey 1991, Kunz *et al* 1999, Salazar-Kish 2001). With increasing levels of concurrency, spikes in coordination and rework intensify, and thus trans-specialist knowledge becomes increasingly important and beneficial in avoiding and resolving cross-functional glitches (Hoopes and Postrel 1999). All of these effects are ignored in the simple analytical model, and they lead to our sixth set of hypotheses.

H6a: As the level of concurrency between design and manufacturing activities increases, schedule duration decreases, but costs rise, and product quality risks increase.

H6b: As the level of concurrency between design and manufacturing activities increases, the importance of trans-specialist knowledge increases with respect to that of specialist knowledge.

7. Equal centralization across organizations. The analytic model is silent with respect to the impact of different organizational policies of centralization or decentralization on the need for trans-specialist knowledge. However, organizational theory suggests that as centralization increases, and as managers take more responsibility for discretionary decisions, then it becomes less important for lower level subordinates to understand all of the integration requirements of their work with other interdependent teams (Burton & Obel 2004). Alternatively, as centralization decreases, subordinates make more of their own decisions without consulting a supervisor, and thus it becomes more important that they have a system-wide understanding of the overall product architecture in order to avoid product-integration errors. Therefore, we draw a seventh and final hypothesis:

H7: As decision making becomes more highly centralized to management, it becomes less important for subordinates to possess trans-specialist knowledge to achieve a given level of product quality.

COMPUTATIONAL MODELING BACKGROUND

In this section we discuss briefly computational organization theory, and provide an overview of our computational modeling approach. We then describe the computational model developed to test our hypotheses.

Computational Organization Theory Research

Computational organization theory (COT) and computational social science (CSS) comprise an emerging, multidisciplinary field that integrates aspects of artificial intelligence, organization studies and system dynamics/simulation (e.g., see Carley and Prietula 1994). Nearly all research in this developing field involves computational tools, which are employed to support computational experimentation and theorem proving through executable models developed to emulate the behaviors of physical organizations (e.g., see Burton et al. 2002, Carley and Lin 1997, Levitt et al. 1999).

As the field has matured, several distinct classes of models have evolved for particular purposes, including: descriptive models, quasi-realistic models, normative models, and man-machine interaction models for training (Cohen and Cyert 1965, Burton and Obel 1995). More recent models have been used for purposes such as developing theory, testing theory and competing hypotheses, fine-tuning laboratory experiments and field studies, reconstructing historical events, extrapolating and analyzing past trends, exploring basic principles, and reasoning about organizational and social phenomenon (Carley and Hill 2001: 87).

Our research within COT/CSS builds upon the planned accumulation of collaborative research over almost two decades to develop rich theory-based models of organizational processes. Using an agent-based representation (Cohen 1992, Kunz et al. 1999), micro-level organizational behaviors have been researched and formalized to reflect well-accepted organization theory

(Levitt et al. 1999). Extensive empirical validation projects (e.g., Christiansen 1993, Thomsen 1998) have demonstrated the representational fidelity, and shown how the qualitative and quantitative behaviors of our computational models correspond closely with a diversity of enterprise processes in practice.

Computational Modeling Environment

Our computational modeling environment consists of the elements described in Table 1, and has been developed directly from Galbraith's information processing view of organizations. This view of organizations, described in detail in Jin and Levitt (1996), has two key implications.

Table 1 Model Elements and Element Descriptions

Model Element	Element Description
Tasks	Abstract representations of any work that consumes time, is required for project completion and can generate exceptions.
Actors	A person or a group of persons who perform work and process information.
Exceptions	Simulated situations where an actor needs additional information, requires a decision from a supervisor, or discovers an error that needs correcting.
Milestones	Points in a project where major business objectives are accomplished, but such markers neither represent tasks nor entail effort.
Successor links	Define an order in which tasks and milestones occur in a model, but they do not constrain these events to occur in a strict sequence. Tasks can also occur in parallel. VDT offers three types of successor links: finish-start, start-start and finish-finish.
Rework links	Similar to successor links because they connect one task (called the <i>driver</i> task) with another (called the <i>dependent</i> task). However, rework links also indicate that the dependent task depends on the success of the driver task, and that the project's success is also in some way dependent on this. If the driver fails, some rework time is added to all dependent tasks linked to the driver task by rework links. The volume of rework is then associated with the project error probability settings.
Task assignments	Show which actors are responsible for completing direct and indirect work resulting from a task.
Supervision links	Show which actors supervise which subordinates. In VDT, the supervision structure (also called the <i>exception-handling hierarchy</i>) represents a hierarchy of positions, defining who a subordinate would go to for information or to report an exception.

The first is ontological: we model knowledge work through interactions of *tasks* to be performed; *actors* communicating with one another, and performing tasks; and an *organization structure* that defines actors' roles, and constrains their behaviors. Figure 1 illustrates this view of tasks, actors and organization structure. As suggested by the figure, we model the organization structure as a network of reporting relations, which can capture micro-behaviors such as managerial attention, span of control, and empowerment. We represent the task structure as a separate network of activities, which can capture organizational attributes such as expected duration, complexity and required skills. Within the organization structure, we further model various *roles* (e.g., marketing analyst, design engineer, manager), which can capture organizational attributes such as skills possessed, levels of experience, and task familiarity. Within the task structure, we further model various sequencing constraints, interdependencies, and quality/rework loops, which can capture considerable variety in terms of how knowledge work is organized and performed.

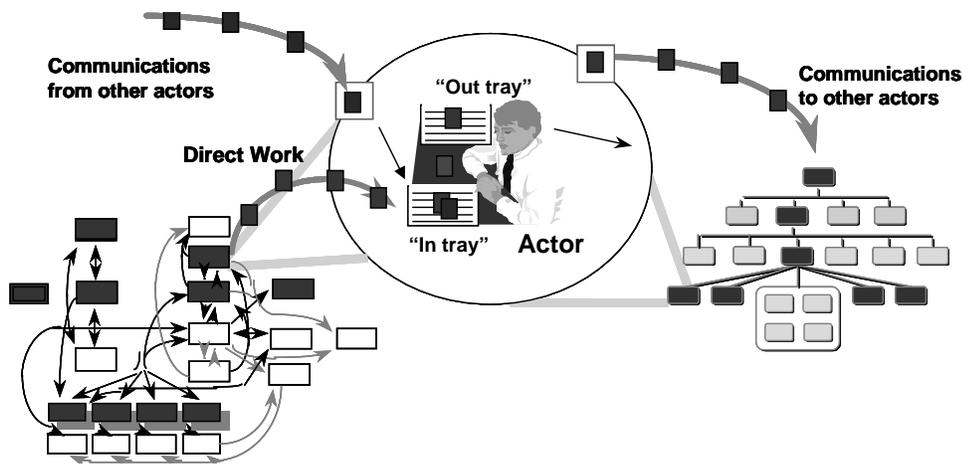


Figure 1. Information Processing View of Knowledge Work

As suggested by the figure also, each actor within the intertwined organization and task structures has a queue of information tasks to be performed (e.g., assigned work activities, mes-

sages from other actors, meetings to attend) and a queue of information outputs (e.g., completed work products, communications to other actors, requests for assistance). Each actor processes such tasks according to how well the actor's skill set matches those required for a given activity, the relative priority of the task, the actor's work backlog (i.e., queue length), and how many interruptions divert the actor's attention from the task at hand.

The second implication is computational: *work volume* is modeled in terms of both *direct work* (e.g., planning, design, manufacturing) and *indirect work* (e.g., decision wait time, rework, coordination work). Measuring indirect work enables the quantitative assessment of (virtual) process performance (e.g., through schedule growth, cost growth, quality).

Computational Model Validation

Our computational model has been validated extensively, over a period spanning almost two decades, by a team of more than 30 researchers. This validation process has involved three primary streams of effort: 1) internal validation against micro-social science research findings and against observed micro-behaviors in real-world organizations, 2) external validation against the predictions of macro-theory and against the observed macro-experience of real-world organizations, and 3) model cross-docking experiments against the predictions of other computational models with the same input data sets. Ours is one of the few extant computational organization models that has been subjected to such a thorough, multi-method trajectory of validation.

COMPUTATIONAL SIMULATION & HYPOTHESIS TESTING

In this section we formulate a computational model of the product development process to replicate and critique Postrel's (2002) analytic model, and to simulate additional scenarios to test our hypotheses. We then discuss our contributions to theory. But first, we describe the model set-up, and define the independent and dependent variables used in the analysis.

Model Set-Up

We use our modeling environment to formulate a model reflecting the basic structure and behavior of the analytical model described in Postrel (2002). Beginning with structural aspects of the model, Figure 2 shows a screenshot of the POW-ER user interface. As with Postrel’s analytic model, this computational model includes only two tasks (i.e., manufacturing and design), each performed by a corresponding organizational unit (i.e., manufacturing actor and design actor), with bi-directional rework links to model Postrel’s requirement (pg. 308) that the model have “an interaction between the two specialists where the choices of one effect the performance of the other.” Including only two tasks is clearly a high-level modeling abstraction; our models of physical organizations in practice typically involve much more detail. However, we specify this minimalist model in conformance with the analytic model, and following Postrel’s advice (pg. 308) that, “the idea is to take the simplest possible situation in which specialized capability and trans-specialist understanding both matter.” Also for consistency with the analytical model, the organization structure reflects no management hierarchy or supervision links.

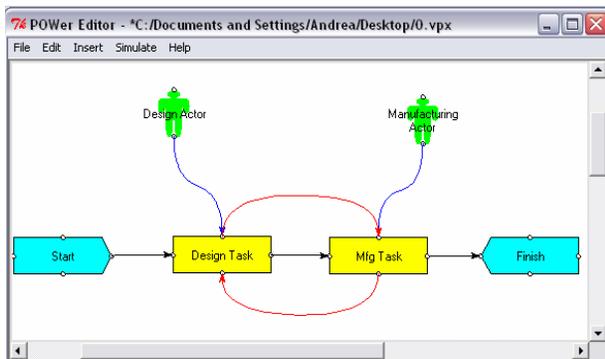


Figure 2, Product Development Model

All model parameters (e.g., *team experience*, *centralization*, *formalization*) are set to empirically determined nominal values for product development. Values for such parameters are held constant for all simulations discussed in this article, except as noted otherwise.

The product development process is specified with a nominal work volume of 200 person-days, which is held constant for all models and simulations discussed in this article. In this model, simulated *total work volume* is determined stochastically (e.g., using Monte Carlo techniques) as a function of nominal work volume (an input) and several empirically determined factors that affect indirect work and productivity (e.g., *actor skill*, *requirement complexity*).

Independent Variables

Table 2 describes our parameters used to represent the set of independent variables included in the model of the product development cycle. To represent behavioral aspects of the analytical model, we disaggregate specialist and trans-specialist knowledge with an approach that is both theoretically consistent and empirically grounded. We manipulate the single parameter *manufacturing skill* to represent different levels of manufacturing specialist knowledge (z), and we manipulate the two parameters *designer role* and *designer application experience* to represent different levels of trans-specialist knowledge (h). The correspondence between our parameter *manufacturing skill* and the analytical model parameter z should be clear: greater manufacturing skill reflects greater specialist knowledge, and vice versa. This representation of z is theoretically consistent with Postrel's argument (pg. 309) that specialist knowledge improves an actor's ability to "hit cost, quality and ramp-up constraints with a high fraction of possible designs." Indeed, the *manufacturing skill* variable within the model offers precisely this outcome—i.e., it reduces the cost, schedule and functional risk associated with completion of the manufacturing task.

Table 2 Independent Variables and Description of Parameters

Independent Variable	Parameter Description
Specialist expertise	<p><i>Actor skill level</i> represents competence in performing a particular task type and reflects specialist expertise. The default skill is Generic, indicating the abilities of an average worker. Actors can be assigned other skills, such as Design, or Manufacturing, depending on their specific areas of expertise. An actor can have a high, medium, or low level of each skill. An actor's skill level compared with the requirement complexity of a task affects task duration, cost, rework and functional risk, but not product risk.</p>
Trans-specialist expertise	<p>Two parameters are used to represent trans-specialist expertise. 1) <i>Actor role</i> represents competence in integration for a particular type of product. Actor role can be set at three levels, reflecting low, medium, or high levels of product integration expertise. With increasing levels of actor role settings, actors become more likely to rework project exceptions, instead of merely quick-fixing or ignoring them. 2) <i>Application experience</i> represents an actor's program level experience, which transcends specialist knowledge (i.e. is trans-specialist in nature). An actor can have a high, medium, or low level of application experience. An actor's role, combined with application experience, represents trans-specialist capability. Together with solution complexity, these affect task duration, cost, rework and project risk, but not functional risk. Therefore, an increase or decrease in trans-specialist understanding is modeled as a two-by-two increase in the actor role and application experience parameters, both in unison.</p>
Task complexity	<p><i>Requirement complexity</i> is used to represent task complexity. It is defined as the number of internal project requirements that a task must satisfy. Requirement complexity can be set at high, medium, or low levels. A highly optimized design, for example, has many tasks with a high requirement complexity. Increasing requirement complexity increases task duration, cost, rework and functional exception levels.</p>
Product modularity	<p><i>Solution Complexity</i> is used to represent product modularity (inversely). It is defined as the effect that a task has on the tasks that depend on it. Solution complexity can be set at high, medium, or low levels, representing low, medium and high levels of modularity, respectively. Higher solution complexity increases task duration, cost, rework and project exception levels.</p>
Actor salary	<p><i>Actor Salary</i> describes an actor's hourly wage.</p>

Task concurrency *Task Concurrency* is the degree to which tasks are conducted in series or in parallel. When product development schedules are shortened to meet corporate goals, tasks that are usually performed sequentially must be performed in parallel. Coordinating parallel interdependent tasks is more difficult and costly than coordinating the same tasks performed sequentially because there is more rework and coordination. We build on data from cases and analyses that assess the coordination and rework that arise from fast-tracking (Christiansen, 1993).

Centralization *Centralization* reflects the degree to which decisions are made by senior actors or decentralized to individual responsible actors. High centralization means decisions are made by high-level actors. With low centralization, responsible positions tend to make their own decisions and there is thus less communication required. Centralization affects how often information is passed from lower to higher level actors, as well as how high that information goes up the hierarchy.

The correspondence between our parameters *actor role* and *actor application experience* and the analytical model parameter h has a similar basis: greater product-integration expertise and programmatic experience reflect greater trans-specialist understanding. This representation of h corroborates Postrel's stated but unmodeled conviction (pg. 308) that "trans-specialist understanding reduces the likelihood that various functional specialties will create problems for one another (design, prototyping, engineering, manufacturing, marketing, etc.), or that groups assigned to work on one part of a product will create components that do not interact properly." Indeed, the combination of application experience and actor role within the model achieves exactly this result—i.e., it improves the likelihood of successful product integration across functional specialties. Furthermore, these representations of z and h are also grounded empirically in research on operational organizations in practice. Such research was conducted during the development of the model. It indicates that, as manufacturers become more specialized, they tend to produce products of greater quality. Likewise, as design staff becomes more senior and experienced, they tend to exhibit greater sensitivity to project integration requirements. In the next section, these variables are defined more formally, along with other variables that are used in the hypothesis testing exercise.

Dependent Variables

Table 3 notes our parameters used to represent the four dependent variables—*product development time*, *labor cost*, *functional risk* and *product risk*—in the model.

Table 3 Dependent Variables and Description of Parameters

Dependent Variable	Parameter Description
Product development time	<i>Simulated project duration</i> (SPD) is the predicted time to perform a project, in working days, which includes both direct and indirect (i.e. coordination, rework and decision latency) work.
Product development labor cost	<i>Simulated labor cost</i> (SLC) is the predicted cost of labor to perform a project, in dollars, which includes both direct and indirect (i.e. coordination, rework and decision latency) work.
Functional risk	<i>Functional risk index</i> (FRI) (or Component Quality Index), measures the risk to quality arising from functional exceptions. Functional exceptions are problems that affect only the task from which they arise. Any rework incurred applies only to that task. Rework links have no interaction with functional exceptions. In project work terms, FRI represents the likelihood that components produced by this project have defects based on rework and exception handling. Numerically, FRI is calculated as the fraction of effort needed to process ignored and quick-fixed functional exceptions normalized by the total effort to rework all predicted functional exceptions.
Product risk	<i>Project risk index</i> (PRI) (or Project Quality Index) measures the risk to quality arising from project exceptions. Project exceptions are problems that arise in a driver task that may have an effect on work in a dependent task linked to the first task via a rework link. In the absence of rework links, project exceptions have no meaning. In project work terms, PRI represents the likelihood that the components produced by this project will not be integrated at the end of the project, or that the integration will have defects based on rework and exception handling. PRI is thus a measurement of the success of system integration. Numerically, PRI is calculated as the fraction of effort needed to process ignored and quick-fixed project exceptions normalized by the total effort to rework all predicted project exceptions.

Replicating Postrel’s Analytic Model of Product Development

Recall the first hypothesis concerning the substitutability of specialist and trans-specialist knowledge. To address this we examine the production function implicit within our computational model. We say “implicit,” because the computational model is not developed with an explicit production function specified. Nonetheless, we can analyze model outputs to assess the interactions of “micro-behavior” assumptions embedded within the model. Here we examine the basic product development model described above.

Using well-accepted Monte Carlo techniques, each model is simulated 100 times, with means and variances computed from empirically derived statistical distributions (see Jin and Levitt 1996, Levitt et al. 2005 for details). In the data tables below, each cell provides the sample mean of 100 individual simulation trials. Statistical significance is computed using a single factor ANOVA test. Although such statistical inference from simulated performance data and predefined distributions remains one-off from the performance of operational organizations, it provides some sense of variation and hence significance, and it extends the analytical model of Postrel by testing hypotheses statistically (e.g., statistical significance is reported at the 0.05 level).

Table 4a Baseline Product Development Model – Results

	SPD (days)*			SPC (\$K)**		
High h	326	262	210	\$93	\$75	\$59
Med. h	380	316	263	\$109	\$90	\$75
Low h	441	378	325	\$127	\$109	\$94
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences in h and z both significant.

** Differences in h and z both significant.

Table 4b Baseline Product Development Model – Results

	FRI (%)*			PRI (%)**		
High h	0.63	0.39	0.32	0.47	0.47	0.46
Med. h	0.62	0.42	0.31	0.60	0.60	0.62
Low h	0.63	0.40	0.34	0.76	0.77	0.76
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences in h insignificant; differences in z significant.

** Differences in h significant; differences in z insignificant.

Computational results for the baseline model are summarized in Tables 4a and 4b. The values listed in Table 4a reflect simulated project duration (SPD), expressed in workdays, and simulated project cost (SPC), expressed in dollars (thousands: \$K). For instance, notice the result in the middle of the SPD half of the table: a project staffed with actors possessing adequate levels of manufacturing specialist knowledge (i.e., medium z) and adequate levels of trans-specialist knowledge (i.e., medium h) is projected by the model to require 316 total days to complete. This reflects the nominal 200 days of work specified (i.e., work volume), along with 82 non-work days (e.g., weekends), and 34 days of additional problem solving (e.g., internal communication, delay and exception handling associated with noise, uncertainty and errors). The additional 34 days' problem solving time reflects empirically determined relationships between model parameters (e.g., levels of z and h) and organizational performance. Similarly, notice the result in the center of the SPC half of the table. The \$90K project cost consists of \$80K of direct work, \$8K of rework, and \$2K of coordination.

The values listed in Table 4b reflect the functional risk index (FRI), expressed as the ratio of functional exceptions that are quick-fixed or ignored versus the total number of exceptions (i.e. the total of those that are quick-fixed, ignored and reworked), and the project risk index (PRI), expressed as the ratio of project exceptions that are quick-fixed or ignored versus the total

number of exceptions. For instance, note the result in the middle of the FRI side of the table: a project staffed with actors with medium h and medium z is projected by the model to quick-fix or ignore 42% of functional exceptions over the course of a project. Likewise, notice the result in the center of the PRI half of the table: 60% of project exceptions are quick-fixed or ignored by the project decision makers.

Tables 4a and 4b report full-factorial results of nine simulation models (each run 100 times), with both the z (i.e., specialist knowledge) and h (i.e., trans-specialist knowledge) parameters varying across three levels: low, medium, and high. Notice in Table 4a that the simulation results vary as expected, and significantly, across the three levels of both z and h. For instance, holding the parameter h constant at the medium level of trans-specialist knowledge, performance in terms of both project duration and cost ranges from 380 days and \$109K dollars when specialist knowledge is low, to 263 days and \$75K when specialist knowledge is high. This mirrors the monotonic relationship between payoff and specialist knowledge described in the analytical model, and indicates the marginal product of such knowledge is positive (i.e., consistent with neoclassical microeconomic theory). This same monotonic relationship is exhibited also at the other levels of trans-specialist knowledge (i.e., low h, high h). Likewise, holding the parameter z constant at the medium level of specialist knowledge, performance in terms of project duration ranges symmetrically from 378 days and \$109K dollars when trans-specialist knowledge is low, to 262 days and \$75K dollars when trans-specialist knowledge is high. This mirrors the monotonic relationship between payoff and trans-specialist knowledge described in the analytical model also, and is evident too at the other levels of specialist knowledge (i.e., low z, high z).

The symmetry reflected in the results of Table 4a corresponds to the microeconomic case of *perfect substitution*: specialist and trans-specialist knowledge can be substituted—unit for unit—to maintain performance at some arbitrary level. For instance from the table, where specialist knowledge (z) is low, but trans-specialist knowledge (h) is medium, performance (380 days, \$109K dollars) is practically the same (378 days, \$109K dollars) as where specialist knowledge (z) is medium, but trans-specialist knowledge (h) is low. Other instances of such substitutability can be identified readily through different combinations of knowledge types z and h (e.g., low z , high h \leftrightarrow high z , low h [326 days, \$93K], high z , medium h \leftrightarrow medium z , high h [263 days, \$75K]). With this our computational model results replicate the basic premise of the analytical model: specialist and trans-specialist knowledge can represent substitutes for one another.

While the data in Table 4a reflect symmetry, the data in Table 4b do not afford such a matched pattern. Instead, we see that project risk is influenced primarily by trans-specialist knowledge h , and that functional risk changes principally in relation to the specialist capability parameter z . Compared to Postrel's finding, this significant, non-symmetric result reveals a more complex and nuanced set of relations between the different kinds of knowledge and multiple dimensions of performance, which we discuss at length in our hypothesis testing below.

In terms of computational experimentation, the significant, simulation results in Table 4a from above support our null hypothesis H_0 (i.e., “specialist and trans-specialist knowledge can be substituted for one another without affecting performance”). With support for the null as such, we note three initial contributions from this work. First, through this computational experiment, we use the model to replicate the basic findings articulated by Postrel's (2002) analytical model; such replication increases our confidence in the analytical model and its implications. Second,

through a model that has been validated empirically, we use our model to show the basic premises of the analytical model have some resemblance to the physical world; such computational validation likewise increases our confidence in the analytical model and its external validity. Third, through a computational model that captures the basic structure and behavior of the analytical model, we can extend the analysis to new cases not yet evaluated, factor in aspects of product development projects that reflect the real world better, and enrich the analysis with a much finer-grained set of measurement variables that represent multiple dimensions of performance including schedule, cost and the risk indices shown in Table 4b.

Hypothesis Testing

Hypothesis testing is organized into seven categories: 1) multidimensional performance, 2) microeconomic complementation, 3) competitive strategies, 4) product modularity, 5) component complexity, 6) concurrency, and 7) centralization. We address each in turn.

Multidimensional performance. The baseline results in Table 4b confirm that specialist knowledge (z) brings a significant reduction in the risk of functional errors (FRI), but that it does not impact the level of project integration risk (PRI). In contrast, while functional risk varies only negligibly across the three levels of trans-specialist knowledge, project integration risk falls substantially as interspecialist capability rises. These findings provide support for hypotheses 1a and 1c. With respect to budget and schedule, Table 4a indicates that both metrics drop sharply as we go from low levels to high levels of specialist knowledge in manufacturing (i.e., this drop is approximately 30%, from 380 to 263 workdays and from \$109K to \$75K). Furthermore, when trans-specialist knowledge is increased in the design phase, both budget and schedule fall also. Thus, the model predicts that greater trans-specialist knowledge brings a reduction in both product integration risk and development time and cost. This finding, which disconfirms hypothesis

1b, also runs contrary to an important conclusion drawn after the Challenger shuttle disaster, which was that a reduction in overall product risks often comes at the cost of a longer product incubation period (Vaughan 1990).

Microeconomic complementation. Recall hypotheses 2a and 2b concerning the expected complementary interaction of the two knowledge types. From Tables 4a and 4b (above) we see that the two knowledge types are close substitutes, not complements, for both the SPD and SPC dimensions of performance. Thus, hypothesis 2a is not supported for the schedule and cost performance dimensions. Furthermore, the model suggests that the two knowledge types play very distinctive roles. While trans-specialist knowledge reduces the risk of product integration failures, specialist knowledge reduces the risk of functional failures; both knowledge types reduce labor cost and schedule duration. Thus, when enriching the analysis beyond cost and schedule variables, to reflect functional and project risks, the two knowledge types reflect complements, not substitutes, in the sense that they both contribute uniquely to an increased likelihood of new product success, or overall “expected payoff.” Thus, hypothesis 2a is supported for the functional- and project-risk performance dimensions.

This result departs from that derived in the analytical model. Recall the analytical model’s production function specifies maximum performance can be attained when *either* knowledge type is high (e.g., $z = 1$, $h = 1$), and nothing is gained when *both* knowledge types are high. The same departure of results applies to low levels of knowledge. In contrast, the computational model reveals specialist and trans-specialist knowledge interacting even when both types are high or low. This provides some evidence to support hypothesis 2b (i.e., “Using predominately specialist knowledge in product development leads to inferior performance than when ap-

preciable levels of trans-specialist knowledge are also employed”). In microeconomic terms, the marginal product of both specialist and trans-specialist knowledge remains positive always.

Competitive strategies. Given that the two knowledge types have different effects on the multiple dimensions of performance, the model suggests that firms will value the two knowledge types in accordance with their competitive strategy. When a firm is competing on price, and is relatively less concerned about product quality, for instance, the symmetric results in Table 4a suggest that the marginal-cost benefit of increased specialist knowledge is equivalent to that of increased trans-specialist knowledge, and that the best performance is achieved through symmetric, high levels of *both* knowledge types. This disconfirms hypothesis 3a. In contrast, if a firm is competing on product quality, and is relatively less concerned about price, for instance, then the asymmetric results in Table 4b suggest that the marginal benefit of increased specialist versus trans-specialist knowledge depends upon the extent to which component quality (i.e., as measured by FRI) or project quality (i.e., as measured by PRI) is stressed. This result elucidates nicely the different contributions of specialist knowledge and trans-specialist knowledge on quality and hence competitive strategy. As such, this result provides mixed support for hypothesis 3b.

These simulated results fit closely with observations in practice of legendary firms like IDEO and Walt Disney, which attract, integrate and retain extremely diverse skill sets in order to build their differentiated product brands. The Walt Disney Imagineering website boasts the integration of more than 140 specialist disciplines! The IDEO corporate website claims that “multi-disciplinary teams are at the heart of the IDEO method,” and lists more than a dozen broad functional areas of expertise. Our simulation results confirm that in this type of product-differentiated firm, integrating interspecialist capabilities to unify diverse skill sets is of great value. Likewise,

our results suggest that for firms competing on cost, such as McDonalds restaurants, it is important to maintain both specialist and trans-specialist capabilities.

Product modularity. Tables 5a and 5b present ratios to show how project outcome metrics change as product modularity rises, which is represented in the model by a decrease in solution complexity. These ratios reflect increases or decreases in normalized terms relative to the baseline results in Tables 4a and 4b. Notice in Table 5a that as product modularity rises, reflecting a decrease in the complexity of interface requirements between sub-components, then both duration and cost fall by 2% to 3% relative to the baseline. Project risk falls by between 14% and 21%, but functional risk is not nearly as sensitive to product modularity. Notice also that the extent of the drop in project risk is influenced by the level of trans-specialist knowledge. With lower levels, this drop is close to 20%; with higher levels, it is only about 15%. Likewise, for a decrease in product modularity (e.g., a less-mature product with more complex interface requirements) the results generated by the model (not shown) are similar in magnitude but in the opposite direction. This result provides support for hypothesis 4, that the marginal value of trans-specialist knowledge falls as products become increasingly modular.

Table 5a SPD and SPC Ratios – High Product Modularity

	% Change in SPD*			% Change in SPC**		
High h	-2%	-3%	-2%	-2%	-3%	-3%
Med. h	-2%	-2%	-3%	-2%	-2%	-2%
Low h	-2%	-2%	-2%	-2%	-2%	-2%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline significant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Table 5b FRI and PRI Ratios – High Product Modularity

	% Change in FRI*			% Change in PRI**		
High h	2%	3%	0%	-16%	-14%	-14%
Med. h	0%	-5%	6%	-18%	-18%	-19%
Low h	-2%	-2%	0%	-19%	-21%	-19%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline insignificant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h significant; differences in z insignificant.

Component complexity. Tables 6a and 6b present ratios to show how project outcome metrics change as component complexity rises, which is represented in the model by an increase in requirement complexity. These ratios reflect normalized changes relative to the baseline results in Tables 4a and 4b. Notice in Table 6a that as component complexity rises, reflecting more internal requirements that are “hidden” to all but that specific component (Baldwin and Clark 1997), then both duration and cost rise by 2% to 3% relative to the baseline. Functional risk rises by between 14% and 26%, but project risk is not nearly as sensitive to component complexity. Upon inspection, we see that the extent of this rise in functional risk is influenced by the level of specialist knowledge. At higher levels, this rise is near an average of 15%; whereas, with lower levels, it is approximately 22%. Analogously, for a decrease in component complexity the results generated by the model (not shown) are in the opposite direction but of magnitude. This result provides support for hypothesis 5: the marginal value of specialist knowledge rises as product components become increasingly complex.

Table 6a SPD and SPC Ratios – High Component Complexity

	% Change in SPD*			% Change in SPC**		
	Low z	Med. z	High z	Low z	Med. z	High z
High h	3%	3%	2%	2%	2%	2%
Med. h	3%	3%	3%	2%	3%	2%
Low h	3%	3%	2%	3%	3%	2%

* Differences relative to baseline significant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Table 6b FRI and PRI Ratios – High Component Complexity

	% Change in FRI*			% Change in PRI**		
	Low z	Med. z	High z	Low z	Med. z	High z
High h	20%	20%	16%	0%	4%	4%
Med. h	20%	17%	14%	0%	0%	-6%
Low h	26%	21%	14%	1%	0%	1%

* Differences relative to baseline significant; differences in h insignificant; differences in z significant.

** Differences relative to baseline insignificant; differences in h and z both insignificant.

Reciprocal interdependence. To test hypothesis 6, two different models were prepared. The first included a start-start successor relationship, so that the manufacturing task commenced just as the design task was 75% complete, after which the two tasks are performed contemporaneously (i.e., 25% schedule concurrency). The second model followed the same logic, except that the degree of concurrency was set to 50%. Results of the eighteen simulations are summarized in Tables 7a through 7d, and are shown relative to the baseline results depicted in Tables 4a and 4b. Notice first that concurrent development requires less time than does the sequential development specified in the baseline model. For the case of 25% concurrency, this schedule reduction is in the range of 13% to 20%, and for the case of 50% concurrency, it is in the range of 23% to 34%. Notice second that as concurrency rises, the level of project risk (PRI) also rises. For the case of 25% concurrency, PRI rises by approximately 1% relative to the baseline model, and for the case

of 50% concurrency it rises by about 3%. Although these rises appear small, they highlight the classic trade-off between time-to-market and product quality discussed in the product development literature (Bayus 1997). Notice third that as concurrency rises, the project labor costs rise reflecting greater amounts of re-work and coordination. This rise is in the range of 1% in the first case and 2% in the second case. These results confirm support for hypothesis 6a, that with greater concurrency, overall schedule duration decreases, but at the cost of an increase in both product cost and quality risk. While the model can replicate these trade-offs, the results of the experiment do not confirm support for hypothesis 6b. There is no evidence suggesting that trans-specialist knowledge plays an increasingly important role in reducing product risk at higher levels of concurrency.

Table 7a SPD and SPC Ratios – 25% Concurrency

	% Change in SPD*			% Change in SPC**		
High h	-14%	-16%	-20%	0%	0%	1%
Med. h	-13%	-17%	-20%	1%	0%	1%
Low h	-13%	-16%	-19%	1%	0%	1%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline significant; differences in h insignificant; differences in z significant.

** Differences relative to baseline insignificant; differences in h and z both insignificant.

Table 7b FRI and PRI Ratios – 25% Concurrency

	% Change in FRI*			% Change in PRI**		
High h	-2%	3%	-3%	2%	0%	2%
Med. h	-2%	-2%	3%	0%	0%	0%
Low h	-2%	-2%	-3%	1%	0%	1%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline insignificant; differences in h and z both insignificant.

** Differences relative to baseline insignificant; differences in h and z both insignificant.

Table 7c SPD and SPC Ratios – 50% Concurrency

	% Change in SPD			% Change in SPC		
High h	-23%	-27%	-34%	1%	1%	2%
Med. h	-23%	-27%	-33%	2%	1%	2%
Low h	-23%	-27%	-32%	2%	0%	3%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline significant; differences in h insignificant; differences in z significant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Table 7d FRI and PRI Ratios – 50% Concurrency

	% Change in FRI*			% Change in PRI**		
High h	0%	3%	3%	2%	4%	4%
Med. h	2%	-5%	3%	3%	3%	2%
Low h	2%	-2%	0%	3%	1%	4%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline insignificant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Centralization. To test hypothesis 7, we altered the model set-up slightly, by adding a management actor with supervision links to the design and manufacturing actors. Then we performed two, 3x3, full-factorial simulations, one with centralization set to low, and the other with centralization set to high. The results are shown in Tables 8a to 8d, again relative to the baseline model. Upon inspection of Tables 8a and 8b, we see first that adding a management actor increases time and cost slightly, but reduces project risk significantly. This is the case, even when centralization is low, and the management actor is handling relatively few of the exceptions encountered by subordinates. Furthermore, we see in Tables 8c and 8d that at high levels of centralization, the duration and cost continue to rise a bit, but project integration risks fall considerably. Notice that at lower levels of trans-specialist knowledge, this effect is much stronger than it is at higher levels of trans-specialist knowledge. This occurs because, in the model, the management actor plays a larger compensating role in exception handling for subor-

dinates with low interspecialist experience than for subordinates with high interspecialist experience. Therefore, the evidence confirms that in a more centralized organization, it is relatively less important for subordinates to possess trans-specialist knowledge, and vice-versa. Based on this evidence, hypothesis 7 is supported.

Table 8a SPD and SPC Ratios – Management with Low Centralization

	% Change in SPD*			% Change in SPC**		
High h	3%	1%	1%	3%	3%	2%
Med. h	2%	1%	2%	3%	2%	2%
Low h	3%	3%	3%	3%	4%	4%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline significant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Table 8b FRI and PRI Ratios – Management with Low Centralization

	% Change in FRI*			% Change in PRI**		
High h	2%	5%	3%	-23%	-21%	-17%
Med. h	3%	-5%	6%	-23%	-18%	-19%
Low h	0%	-2%	-6%	-26%	-27%	-28%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline insignificant; differences in h insignificant; differences in z significant.

** Differences relative to baseline significant; differences in h and z both insignificant.

Table 8c SPD and SPC Ratios – Management with High Centralization

	% Change in SPD*			% Change in SPC**		
High h	3%	2%	1%	4%	4%	3%
Med. h	5%	4%	4%	6%	6%	7%
Low h	8%	8%	6%	9%	10%	10%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline significant; differences in h significant; differences in z insignificant.

** Differences relative to baseline significant; differences in h significant; differences in z insignificant.

Table 8d FRI and PRI Ratios – Management with High Centralization

	% Change in FRI*			% Change in PRI**		
High h	2%	5%	0%	-49%	-45%	-46%
Med. h	2%	-2%	0%	-53%	-52%	-53%
Low h	0%	-5%	3%	-63%	-64%	-67%
	Low z	Med. z	High z	Low z	Med. z	High z

* Differences relative to baseline insignificant; differences in h and z both insignificant.

** Differences relative to baseline significant; differences in h significant; differences in z insignificant.

Summary

For the most part, simulation results from our computational models support extant knowledge-flow theory. This is to be expected, for the micro-behaviors used to develop such models are derived from information-processing views of organization theory. Likewise for the most part, simulation results from our computational models are consistent with KM practice. This is also to be expected, for the micro-behaviors used to develop such models have been validated against practice in operational organizations. Further, our initial basic result—that specialist and trans-specialist knowledge can represent substitutes—supports the central premise of the analytical model. This provides some computational support for the analytical model, and it provides analytical validation for the computational model. Through triangulation between the analytical and computational model results, we gain increased confidence in both approaches.

In contrast, the analytical model's most controversial results are not consistent with those obtained through computational experimentation. Nor do many of its simplistic assumptions appear to generalize well from the tidy realm of mathematical analysis to the messy domain of organization knowing and learning. In particular we show: how multiple performance metrics are necessary to assess the multidimensional and more nuanced influences that specialist and trans-specialist knowledge exert on organizational performance; how the relative costs and benefits of specialist and trans-specialist knowledge vary; how trans-specialist knowledge can actually in-

crease schedule duration as more constraints are placed on requirements; and how the two knowledge types interact, even when both types are high or low, as complements in some cases and as substitutes in others.

In response to Postrel's provocative question, "Under what circumstances is it necessary for specialists to develop mutual understanding?" we offer several theoretical insights. Our results indicate that trans-specialist knowledge becomes increasingly beneficial as competitive strategy becomes more sensitive to product quality, as product modularity declines, and as decision-making becomes less centralized. In contrast to his metaphor *islands of shared knowledge*, our computational experiments suggest how the alternate metaphor *streams of shared knowledge* appears to fit much better: Shared knowledge need not be omnipresent (e.g., consider *oceans of shared knowledge*), but it is important frequently for knowledge to flow across functional specialists.

CONCLUSION

In this article we investigate empirically the theoretical split between emphases upon specialist versus trans-specialist knowledge in the organization—or more generally between exploitation and exploration—a split which divides knowledge-flow theory at present, and hence represents an important issue for Knowledge Management (KM). We review the relevant literature, articulate hypotheses, and employ computational experimentation to test them empirically.

Our findings provide novel, insightful understanding of the factors that contribute toward understanding the relative balance between specialist versus trans-specialist knowledge in particular, and exploitation versus exploration more generally. In particular, we show how specialist and trans-specialist knowledge can represent substitutes in certain circumstances, and hence how exploitation can dominate the KM balance with respect to exploration. Alternatively, we show

also: how multiple performance metrics are necessary to assess the multidimensional and more nuanced influences that specialist and trans-specialist knowledge exert on organizational performance; how the relative costs and benefits of specialist and trans-specialist knowledge vary; how trans-specialist knowledge can actually increase schedule duration as more constraints are placed on requirements; and how the two knowledge types interact, even when both types are high or low, as complements in some cases and as substitutes in others.

Further, our results indicate that trans-specialist knowledge becomes increasingly beneficial as competitive strategy becomes more sensitive to product quality, as product modularity declines, and as decision-making becomes less centralized. Our computational experiments suggest how the alternate metaphor *streams of shared knowledge* appears to fit well: it is important frequently for knowledge to flow across functional specialists.

Our contribution in this article to the KM literature is three-pronged. First, we expose a weakness in the existing body of theory through our critique of Postrel's analytic model. Second, we propose and test a set of hypotheses pertaining to Postrel's central research question, and by doing so we provide a much finer-grained set of theoretical insights than was possible through the analytic model. Third, we apply computational experimentation in our analysis, which represents a relatively recent approach to testing theory related to organizational knowing and learning. We hope that other researchers are stimulated to conduct further research along these lines based on the power and potential of this method.

Further, we can imagine several situations that we were unable to test within our modeling environment, under which trans-specialist knowledge becomes increasingly valued in organizations. When organizational conflict is high, it is likely that trans-specialist capability is important as a means to mediate between divergent "local languages" and "thought worlds."

When managers lack resources to hire sufficient specialists, such as during the early phases of entrepreneurial growth of small businesses, then it is likely that early hires that must “bootstrap” by simultaneously managing multiple roles and functions benefit from trans-specialist understanding. Moreover, when managers establish new organizations or functional-teams, it is likely that their ability to assemble a cooperative, non-redundant group of functional specialists rises with increasing levels of interspecialist knowledge. Finally, and more generally, in any situation where it is important for an organization to explore new product, process or market terrain, trans-specialist knowledge is probably beneficial to evaluate simultaneous disciplinary constraints (e.g., marketing, public relations and manufacturing constraints) that bear on overall decision making, but that are not yet well understood or formalized into codes, rules and procedures.

These situations represent opportunities for future researchers to continue to explore the circumstances under which trans-specialist knowledge is valued most-highly in organizational settings.

REFERENCES

- Allen, T.J. 1977. *Managing the Flow of Information Technology*, Cambridge, MA, MIT Press.
- Allen, T.J., S. Cohen. 1967. Information Flow in R&D Labs. *Admin. Sci. Quarterly* **14** 12-19.
- Arrow, K. 1974. *The Limits of Organization*, New York, Norton.
- Bacharach, S.B. 1989. Organizational Theories: Some Criteria for Evaluation. *Academy of Management Review* **14**(4) 496-515.
- Baldwin, C., K. Clark. 1997. Managing in an age of modularity. *Harv. Bus. Rev.* **75**(5) 84-93.
- Bayus, B.L. 1997. Speed-to-market and new product performance trade-offs. *J. Product Innovation Management* **14**(6) 485-497.
- Birkinshaw, J., R. Nobel, J. Ridderstrale. 2002. Knowledge as a Contingency Variable: Do the Characteristics of Knowledge Predict Organization Structure? *Org. Sci.* **13**(3) 274-289.
- Brown, J.S., P. Duguid. 1991. Organizational Learning and Communities-of-Practice: Toward a

- Unified View of Working, Learning, and Innovation. *Organization Science* **2**(1) 40-57.
- Burton, R.M., J. Lauridsen, B. Obel. 2002. Return on Assets Loss from Situational and Contingency Misfits. *Management Science* **48**(11) 1461-1485.
- Burton, R.M., B. Obel . 2004. *Strategic Organizational Diagnosis and Design: The Dynamics of Fit*, 3rd Edition New York, Kluwer.
- Burton, R., B. Obel. 1995. The Validity of Computational Models in Organization Science: From Model Realism to Purpose of the Model. *Comp. and Math. Org. Theory* **1**(1) 57-71.
- Carley, K.M. 1992. Organizational Learning and Personnel Turnover. *Org. Science* **3**(1).
- Carley, K.M., Z. Lin. 1997. A Theoretical Study of Organizational Performance under Information Distortion. *Management Science* **43**(7) 976-997.
- Carley, K.M., M.J. Prietula . 1994. *Computational Organization Theory*, Hillsdale, NJ, Lawrence Erlbaum Associates.
- Carley, K., V. Hill . 2001. Structural Change and Learning Within Organizations. In: Lomi, A. and Larsen, E.R., (Eds.) *Dynamics of Organizational Computational Modeling and Organization Theories*, Cambridge, MA, The MIT Press.
- Campbell, D. 1988. Task Complexity: A Review & Analysis, *Acad. of Mgmt Rev.*, 13 (1) 40-52.
- Carroll, J.S. 1998. Organizational Learning Activities in High-Hazard Industries: The Logics Underlying Self-Analysis. *Journal of Management Studies* **35**(6) 699-717.
- Cheng, Carol H. F. and Levitt, Raymond E. 2001. Contextually changing behavior in medical organizations. 2001 Annual Symposium of the American Medical Informatics Association. Washington, D.C.
- Christiansen, T.R. 1993. Modeling Efficiency and Effectiveness of Coordination in Engineering Design Teams. Unpublished Ph.D Dissertation. Stanford University.
- Cohen, G.P. 1992. The Virtual Design Team: An Information Processing Model of Design Team Management. Unpublished Ph.D Dissertation. Stanford University.
- Cohen, K.J., R.M. Cyert . 1965. Simulation of Organizational Behavior. In: March, J.G., (Ed.) *Handbook of Organizations*, Rand McNally.
- Cohen, W., D. Levinthal. 1990. Absorptive Capacity: A new perspective on learning and innovation. *Administrative Science Quarterly* **35** 128-152.
- Cole, R.E. 1998. Introduction. *California Management Review* **45**(3) 15-21.
- Cook, S.D.N., J.S. Brown. 1999. Bridging Epistemologies: The Generative Dance between Or-

- ganizational Knowledge and Organizational Knowing. *Org. Science* **10**(4) 381-400.
- Cooper, R.G. 1979. The dimensions of industrial new product success and failure. *Journal of Marketing* **43** 93-103.
- Cyert, R.M., J.G. March . 1963. *A Behavioral Theory of the Firm*, Cambridge, MA, Blackwell.
- Dougherty, D. 1992. Interpretive barriers to successful product innovation in large firms. *Organization Science* **3**(2) 179-202.
- Drucker, P.F. 1995. *Managing in a Time of Great Change*, New York, NY, Truman Talley.
- Ethiraj, S., D. Levinthal. 2004. Modularity and Innovation in Complex Systems. *Management Science* **50**(2) 159-173.
- Galbraith, J.R. 1977. *Organizational Design*, Addison-Wesley.
- Grant, R.M. 1996. Toward a Knowledge-Based Theory of the Firm. *Strat. Mgmt J.* **17** 109-122.
- Hoopes, D.G., S. Postrel. 1999. Shared Knowledge, "Glitches," and Product Development Performance. *Strat. Mgmt. J.* **20** 837-865.
- Jin, Y., R.E. Levitt. 1996. The Virtual Design Team: A Computational Model of Project Organizations. *Journal of Computational and Mathematical Organizational Theory* **2**(3) 171-195.
- Kogut, B., U. Zander. 1992. Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology. *Organization Science* **3**(3) 383-397.
- Kunz, J.C., R.E. Levitt, Jin Y. 1999. The Virtual Design Team: A computational model of project organizations. *Comm. of the Association for Comp. Machinery* **41**(11) 84-92.
- Law, A.M., D. Kelton . 1991. *Simulation Modeling and Analysis 2nd Ed.*, New York, NY, McGraw-Hill.
- Lawrence, P.R., J.W. Lorsch . 1967. *Organization and Environment: Managing Differentiation and Integration*, Boston, Division of Research, Harvard Business School.
- Levitt, R.E., J. Thomsen, T.R. Christiansen, J.C. Kunz, Y. Jin, C. Nass. 1999. Simulating Project Work Processes and Organizations: Toward a Micro-Contingency Theory of Organizational Design. *Management Science* **45**(11):1479-1495.
- Levitt, R.E. 2004. Computational Modeling of Organizations Comes of Age. *Journal of Computational and Mathematical Organization Theory* **10**(2) 127-145.
- Levitt, R.E., R.J. Orr, M. Nissen. 2005. Validating the Virtual Design Team (VDT) Computational Modeling Environment. The Collaboratory for Research on Global Projects, Working Paper #25, 1-15. Available at:
http://crp.stanford.edu/publications/working_papers/WP25.pdf

- Locke, E., G.P. Latham. 1990. *A Theory of Goal-Setting & Task Perf.*, Eng. Cliffs, NJ: Pren.-Hall.
- Lomi, A., E.R. Larsen . 2001. *Dynamics of Organizations: Computational Modeling and Organization Theories*, Menlo Park, CA, American Association of Artificial Intelligence.
- March, J.G. 1991. Exploration and exploitation in organizational learning. *Org. Sci.* **2**(si) 71-87.
- March, J.G., H.A. Simon 1958. *Organizations*, New York: Wiley.
- Milgrom, P., J. Roberts. 1995. Complementarities and fit strategy, structure, and organizational change in manufacturing. *Journal of Accounting and Economics* **19**(2-3) 179-208.
- Milgrom, P., J. Roberts. 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review* **80**(3) 511-528.
- Morgan, L.O., R.M. Morgan, W.L. Moore. 2001. Quality and Time-to-Market Trade-offs when There Are Multiple Product Generations. *Manufacturing & Service Operations Management* **3**(2) 89-95.
- Mustafa, M.A., E.L. Murphree. 1989. A Multicriteria Decision Support Approach for Project Compression. *Project Management Journal* (June).
- Nissen, M.E., R.E. Levitt. 2004. Agent-Based Modeling of Knowledge Dynamics. *Knowledge Management Research & Practice* **2**(3) 169-183.
- Nissen, M. 2005. Hypothesis Testing of Edge Organizations: Specifying Computational C2 Models for Experimentation. In: Anonymous *Proceedings International Command & Control Research Symposium*, McLean, VA.
- Nissen, M. 2006. *Harnessing Knowledge Dynamics: Principled Organizational Knowing & Learning*, Hershey, PA: IRM Press.
- Nogueira, J.C. 2000. *A Formal Model for Risk Assessment in Software Projects*, Doctoral Dissertation, Dep. of Computer Sci., Monterey, CA, Naval Postgraduate School.
- Nonaka, I. 1994. A Dynamic Theory of Organization Knowledge Creation. *Org. Science* **5** 14-37.
- Pindyck, R.S., D.L. Rubinfeld . 1998. *Microeconomics (Fourth Edition)*, Upper Saddle River, NJ, Prentice-Hall.
- Postrel, S. 2002. Islands of Shared Knowledge: Specialization and Mutual Understanding in Problem-Solving Teams. *Organization Science* **13**(3) 303-320.
- Rulke, D.L., J. Galaskiewicz. 2000. Distribution of knowledge, group network structure, and group performance. *Management Science* **46**(5) 612-625.

- Russel, A.D., M. Ranasinghe. 1991. Decision Framework for Fast-Track Construction: A Deterministic Analysis. *Construction Management and Economics* **9** 467-479.
- Salazar-Kish, J.M. 2001. *Modeling Concurrency Tradeoffs and their Effects on Project Duration and Rework, Doctoral Dissertation, Department of CEE, Stanford University.*
- Samuelson, P.A. 1974. Complementarity: An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory. *Journal of Economic Literature* **12**(4).
- Schein, E.H. 1996. The three cultures of management. *Sloan Management Review* **38** 9-20.
- Schilling, M. 2000. Toward A General Modular Systems Theory and Its Application to Interfirm Modularity. **25**(2) 312-334.
- Scott, W.R. 2003. *Organizations: Rational, Natural, and Open Systems, 5th ed.*, Englewood Cliffs, NJ, Prentice Hall.
- Simon, H.A., 1981. *The Sciences of the Artificial*, 2nd ed., Cambridge, Mass.: The MIT Press.
- Smith, P., D.D. Reinersten . 1991. *Developing Products in Half the Time*, New York, NY, Van Nostrand Reinhold.
- Spender, J.C. 1996. Making Knowledge the Basis of a Dynamic Theory of the Firm. *Strategic Management Journal* **17** 45-62.
- Staddon, J.E.R. 1983. *Adaptive Behavior and Learning*, Internet Edition Cambridge, MA, First Edition Published by Cambridge University Press.
- Szulanski, G. 2000. The Process of Knowledge Transfer: A Diachronic Analysis of Stickiness. *Organizational Behavior and Human Decision Processes* **82**(1) 9-27.
- Thompson, J.D. 1967. *Organizations in Action*, New York, McGraw-Hill.
- Thomsen, J. 1998. The Virtual Team Alliance (VTA): Modeling the Effects of Goal Incongruency in Semi-Routine, Fast-Paced Project Organizations. Stanford University.
- Utterback, J.M. 1996. *Mastering The Dynamics of Innovation*, Harvard University Business School Press, Boston, Massachusetts.
- Van den Bosch, F.A.J., H.W. Volberda, M. de Boer. 1999. Coevolution of Firm Absorptive Capacity and Knowledge Environment. *Organization Science* **10:5**(5) 551-568.
- Vaughan, D. 1990. Autonomy, Interdependence and Social Control: NASA and the Space Shuttle Challenger. *Administrative Science Quarterly* **35** 225-257.
- Von Hippel, E. 1994. 'Sticky Information' and the Locus of Problem Solving: Implications for Innovation. *Management Science* **40**(4) 429-439.
- Zollo, M., S.G. Winter. 2002. Deliberate Learning and the Evolution of Dynamic Capabilities. *Organization Science* **13**(3) 339-351.