20 June 2011

INTERDEPENDENCE & ORGANIZATION DESIGN

Phanish Puranam & Marlo Goetting[1]

5826 words, main text only

INTRODUCTION

If organizations are multi-agent systems with goals, there must exist a mapping from organizational level goals to agent level tasks. Such a mapping, at least when explicitly recognized (even if not intentionally crafted) is what we think of typically as *division of labor*. Because the results of the efforts so divided must be integrated back, the division of labor results in interdependence between the agents performing the tasks contributing to the overall goal of the organization.

Many perspectives in organization theory build on the premise that organizations "solve" the problems of cooperation and coordination that arise when integrating the efforts of interdependent actors – albeit with varying degrees of success. This premise is common to the analysis of organizational structures using a diverse set of constructs such as information processing (e.g. Simon, 1945; Thompson, 1967), contingency and fit (e.g. Lawrence & Lorsch, 1967), complementarities (e.g. Milgrom & Roberts 1990), epistatic interactions (e.g. Levinthal 1997), power (e.g. Pfeffer & Salancik, 1978), reward interdependence (e.g. Kelley & Thibaut, 1978) and asset specificity (e.g. Williamson, 1975).

We have two objectives for our essay. First, we consider the conditions under which interdependence between agents gives rise to cooperation problems vs. coordination problems.

Second, we consider some ways in which the cooperation and coordination problems arising from interdependence can be analyzed jointly, rather than separately as traditionally has been the case; indeed it has been noted, tongue in cheek, that alien social scientists visiting our planet might be stunned by the vastly different conceptualization of organizations that result from an exclusive focus on one or the other (Dosi, Levinthal & Marengo, 2003). We offer some thoughts on how to bridge the gap between how interdependence is treated in these two conceptualizations of organizations- as systems for obtaining cooperation as opposed to systems of coordinated action.

While much of our discussion, for expositional simplicity, will focus on an interdependent *dyad*, we emphasize the point that the same basic ideas can be applied at larger scales and multiple levels of aggregation – that interdependence between (multiple) individuals, groups, departments and firms can be analyzed using the same

[1] This chapter draws on joint work with several co-authors. Carliss Baldwin, Rich Burton and Dan Levinthal provided extensive feedback on the related work this chapter draws on. Donal Crilly and Niro Sivanathan offered useful suggestions on this draft. Puranam acknowledges funding from the European Research Council under Grant # 241132 for "*The Foundations of Organization Design*" project.

concepts. Indeed, this "fractal" nature of interdependence principles is a basis for optimism regarding the construction of a science of organization design (see also Grandori & Furnari, 2008).

INTERDEPENDENCE AS A CONSEQUENCE OF THE DIVISION OF LABOR

The division of labor in an organization has two logically distinct components: the decomposition of the overall goal into (clusters of) tasks, and the allocation of these task clusters to individual agents (Smith, 1776). We refer to these two processes as task decomposition and task allocation.

Task Decomposition: There are typically a variety of feasible ways to decompose a given organizational goal into a set of contributing tasks. We define a *task structure* as any means-end decomposition (Newell & Simon, 1972) of a goal into its constituent tasks and the interdependency relationships between these. A task is the fundamental, indivisible unit of a task structure and may be thought of as a production technology - it is a transformation of inputs into outputs in a finite time period. Since tasks have inputs and outputs, they have an associated value (the difference between the benefits of the outputs and the cost of the inputs) as seen by the designer of the system.

Let T^R be the most fine-grained task structure visible to an omniscient organization designer. There are multiple ways one could represent T^R - Workflow diagrams, Design Structure Matrices, and Linear Programming models being some of them. These representations capture the possible relationships between tasks - whether one produces an input to another (Thompson, 1967), whether their joint outputs are complements or substitutes (Milgrom & Roberts, 1990), or whether they draw on the same, possibly limited, inputs (Burton & Obel, 1984). These different kinds of interdependencies between tasks can be represented analytically in terms of the different ways that each task's inputs and outputs enter a combined value function.

Following Puranam, Goetting and Knudsen (2011) we say that two tasks are *interdependent when the value generated from performing each is different when the other task is performed vs. when it is not*. The tasks are *independent* if the value of performing each task is the same whether the other task is performed or not. As a consequence the combined value created when independent tasks are performed is the same as the sum of the values created by performing each task alone (e.g. "pooled interdependence" in Thompson, 1967, in which each task makes a discrete contribution to the whole). Specialization represents task decomposition into heterogeneous components and by the definition above introduces interdependence between tasks. In contrast, non-specialized task decomposition can leave tasks independent (Leijonhufvud, 1986).

For the n-task case, Task Structure Matrices (Baldwin & Clark, 2000; Eppinger, 1991; Steward, 1981) provide a compact matrix form representation of patterns of interdependence between tasks. In a Task Structure Matrix (TSM), different tasks are mapped on both the rows and the columns of the matrix – the cells are then populated by 'x' if 'column task' is dependent on 'row task'. Each task is dependent on itself. Dependencies can be symmetric or asymmetric (see Figure 1). While information of the functional form of the interdependence is not given, in a TSM the overall patterns of interdependence between tasks are easy to see.

It is possible to transform other representations of task interdependence into a Task Structure Matrix with some loss of information. For instance, Thompson's (1967) classic discussion of interdependence is based on a workflow-based representation of tasks. Three increasingly complex levels of interdependence are discussed: pooled, sequential, and reciprocal interdependence, as shown in Figure 2. Pooled interdependence corresponds to no interdependence between the tasks; they contribute independently to overall task performance. Sequential and reciprocal interdependencies can be portrayed by asymmetric and symmetric "x"s about the diagonal, respectively. Burton and Obel (1984) represent task interdependence using a linear programming model: the task interdependence is given in the form of a series of decision variables and constraints posed on them. Interdependence exists to the extent that multiple decision variables X_i are linked to the same constraint C (see Figure 3). If we consider each decision variable to be equivalent to a task, then to the extent that different tasks are linked to the same constraint, they are interdependent. It is possible to convert the linear program into a TSM - albeit with loss of information about the functional form of the interdependence, and on the assumption that the interdependence is always symmetric. The matrix of coefficients in the linear programming model and the corresponding TSM are shown in the lower part of Figure 3.

Task Allocation: Lacking omniscience, bounded rational individuals must work with imperfect representations of T^R. As a consequence, there will typically be many different ways in which organization designers with bounded cognitive capacities can cluster the tasks in T^R, and these clusters of task may be allocated in different ways among the agents in the organization.

Let the matrix T^A capture the interdependencies between the *clusters* of tasks allocated to the agents. The noteworthy difference is that whereas T^R is an n X n matrix, T^A will be an m X m matrix, where m is the number of agents in the organization. Since T^A embodies both a decomposition of the overall goal into clusters of tasks as well as an allocation of these task clusters among the agents, ***it is a concise abstract representation of a division of labor*** (Goetting, Puranam & Warglien, 2011). Thus to consider the original example provided by Adam Smith, pin making could be divided into "*eighteen distinct operations, which, in some manufactories are all performed by distinct hands, though in others, the same man will sometimes perform two or three of them*" (1776: 5). These would correspond to two different T^A of dimensions m=18 and m <18 respectively, for the same underlying T^R.

Given an allocation of tasks to a dyad of agents A and B, interdependence between the agents exists when the returns to A from A's actions depend on B's actions and vice versa. This conceptualization of interdependence between agents appears explicitly in the analysis of reward interdependence (Kelley & Thibaut, 1978), power (Emerson, 1962; Pfeffer & Salancik, 1978) and in game theory in general (e.g. von Neuman & Morgenstern, 1944).

In Puranam et al (2011) we show that interdependence between tasks is neither necessary nor sufficient for interdependence between the agents (that these tasks are allocated to) to arise. Unlike task interdependence, interdependence between agents depends entirely on a key feature of their reward structure- *incentive breadth.* This refers to the level of aggregation at which an agent's actions (or their results) are measured and rewarded. In the case of two agents, narrow incentives correspond to the measurement and reward of individual actions or their results, whereas broad incentives correspond to the measurement and reward of dyad-level joint actions or results. For instance if A provides a critical input to B and B is measured on the final output, then the reward structure is *de facto* broad for B unless B's actions can be measured and rewarded independently of whether A has provided the critical input. Agents are interdependent when they face broad incentives, but are independent when they face narrow incentives.

Put differently, interdependence between tasks is assessed by examining the value function representing the combined system of tasks but interdependence between agents depends on the reward function of the agents. Since in general these will not be identical, there will be a corresponding divergence between task and agent interdependence. For instance, even if the tasks assigned to each agent are interdependent, the agents may be measured and rewarded narrowly for their own tasks.

Consider again the pin factory example popularized by Smith (1776): As an example of independent tasks but interdependent agents, agents A and B are to produce 100 pins each (i.e. the value of A's task output does not change with B's output, and vice versa) but they are paid only if a total of 200 pins are produced. On the other hand, if this was a specialized production process, A and B could be tasked to produce a total of 300 pins, of which A produces the tails and B produces the heads. In this case, the tasks are clearly interdependent; however, if both A and B are rewarded on their individual output, respectively (i.e. A is rewarded if she produces 300 tails regardless of B's performance and vice versa), the agents are effectively independent.

Thus, the division of labor creates interdependence between tasks as well as between the agents to whom the tasks are assigned - even though these two patterns of interdependence may not be identical. Note that the same principles apply if we replace this two-task two-agent case, with the n-task m–agent case. Every dyad of agents can be rewarded with narrow or broad incentives, and their allocated tasks or clusters of tasks can be interdependent or independent with those allocated to other agents.

IMPEDIMENTS TO THE INTEGRATION OF EFFORT

The problem of the integration of effort may be stated *as getting the agents to take actions that maximize the value of the system as a whole, for a given division of labor*. For instance, integration in a multi-department organization has been defined as the "quality of the state of collaboration that exists among departments that are required to achieve unity of effort by the demands of the environment" (Lawrence & Lorsch, 1967a: 11).

Why is it harder to achieve coordinated action around certain patterns of interdependence but not others? There are a number of reasons why the integration of effort may be problematic, but it is helpful to think of these as falling broadly into two categories, related to "knowledge" and "motives" of the agents, respectively (Hoopes & Postrel, 1999). Within each, we can also distinguish between situations where the agents are interdependent vs. when they are independent.

Skill and Coordination problems

Bounded rationality (Simon, 1945) implies that there are limits to the agents' knowledge of how to perform their assigned task (for the moment we will assume that the right actions generate rewards large enough to cover their costs of efforts). Inadequate knowledge of the agents may act as an impediment to the integration of effort in two ways. Consider first the case where the agents are independent of each other. There may be a failure of integration if an agent does not know enough to perform the task cluster assigned to him perfectly. Even though the agents are independent of each other, the impact of this *skill* problem on the integration of effort in the system depends crucially on the interdependence of tasks in the system. In a system with highly interdependent tasks, skill failures at individual tasks will have different implications for system performance compared to that in a system with independent tasks.

Returning to the pin factory example, absent specialization in the case where agents produce 100 pins each, a daily output of 100 fully functional pins is possible even if B (alone) lacked the necessary skill; however, in the case where agents are specialized, the 300 faulty heads produced by B due to a lack of skill will reduce the total value of the tasks to zero, even if A produces 300 perfect pin tails.

Now consider the case where the agents are interdependent. A knowledge-related impediment to integration of effort arises if the agents' interdependence takes the form that the optimal action of each agent depends on a prediction of what the other agents will do. Puranam et al. (2011) describe this as a situation of *epistemic interdependence*- each agent needs to know enough about the other to be able to act *as if* they can predict the other's actions- they need "predictive knowledge". In this view, a coordination failure is a failure to predict the actions of another in situations where such a prediction is essential for optimal action by oneself. In other words, a coordination failure occurs when there is epistemic interdependence but the agent(s) do not possess the necessary predictive knowledge. (To scale this up to the n-task m-agent case, we can examine the predictive knowledge requirements for all agents jointly).

Informally, we often equate coordination failures with communication challenges, and the notion of epistemic interdependence helps to see why this intuition is robust, for two reasons. First, a failure of communication prevents the formation of predictive knowledge, and second, indeed communication itself can be seen as a coordination problem, as the modern view of linguistics does: when communicating, I need to predict which among several possible meanings you chose to attach to the words you used. Talk, if it is to be understood, is seldom cheap- because communication itself is a coordination problem in the domain of meaning (Clark, 1996).

Note that the presence of agent interdependence is a necessary but not sufficient condition for the occurrence of coordination problems. A second necessary (and jointly sufficient with the first) condition regards the timing of the actions - only if at least one agent (who faces broad incentives) needs to act before knowing the other agent's actions will epistemic interdependence exist. For instance, if A provides an input to B and both are rewarded on final output, both agents are interdependent with each other; but only A needs predictive knowledge about B (this is the familiar backward induction problem).

To make this concrete, consider the pin factory with independent tasks but interdependent agents (i.e. the agents are jointly rewarded only if 200 pins are produced at the end of the day and each agent is to produce 100 pins). If A works the morning shift and B in the afternoon, and assume B can see A's total output before he starts his shift, he can decide to take the rest of the day off if A has not reached her target. However, A will have to make inferences (based, for instance on prior experience of working with B) regarding B's productivity since her final reward will be contingent on B producing his 100 pins once A has completed her part.

Skill and coordination problems have been central to classical organization theory (March & Simon, 1958). Both individual competence as well as predictive knowledge can be formed through what have been broadly termed "information processing activities" - communication, mutual observation, learning and (joint) decision making by the agents (e.g. Galbraith, 1973; March & Simon, 1958; Tushman & Nadler, 1978). While communication is the most obvious means of creating predictive knowledge (so much so that we sometimes loosely equate communication with coordination), it is not the only one. For instance, Srikanth and Puranam (2010) in their study of coordination processes in business process offshoring discuss "tacit coordination mechanisms" that allow the formation of predictive knowledge across locations through enhancing observability of context, actions, and outcomes rather than through direct communication.

More generally, predictive knowledge need not involve shared knowledge of any order. For instance, precedents -actions used in the past that are psychologically prominent– and conventions -established principles of action that are not questioned- are forms of predictive knowledge that may arise from shared knowledge (Camerer, 2003; Lewis, 1969; Schelling, 1960)[i]. On the other hand, the mutual adaptation of agents may result in coordinated action with very little overlaps in knowledge, as each agent forms reinforced habits of action in response to the other. Interpersonal routines are quite frugal in terms of shared knowledge requirements (Cohen & Bacdayan, 1994; Nelson & Winter, 1982).

Even if task interdependence cannot be modified, epistemic interdependence in principle can be. The designer may understand the architecture of the system sufficiently to measure and sequence actions to reduce the need for predictive knowledge between agents. In such a system, there may be effective integration of effort even though none of the agents consciously needs to coordinate with any other. Note that a sophisticated level of architectural knowledge can help the designer choose a division of labor that makes the integration of effort easier (e.g. Baldwin & Clark, 2000). In particular, by dividing the task into clusters that are mutually independent, the designer can effectively reduce interdependence between agents, while allocating the tasks to those individuals with some existing predictive knowledge of each other will mitigate coordination problems. In general, many aspects of organization design can be seen as the interplay between the architectural knowledge of the designer and predictive knowledge among the agents (Puranam et al, 2011) - though the "designer", in some cases may of course be evolution through selection.

Agency and Cooperation Problems

We now turn to a class of impediments to integration of effort that are traceable to the motivation of the agents. For the moment, we assume that there are no knowledge-related impediments to the integration of efforts- there are no skill or coordination problems. We will assume that agents are purposive (if not maximizing) and only take actions whose returns exceed their private costs. In this sense we can say that they are always "self-interested" as long as we acknowledge that agents may derive utility from things other than personal cash rewards.

Thus just as bounded rationality is the key behavioral assumption when considering impediments to integration arising from knowledge related issues (i.e. skill and coordination problems), the central behavioral assumption for understanding impediments to integration arising from the motives of the agents is that agents have costly actions. In general, agents do not take these costly actions unless the rewards at least cover the costs. A stronger version would assume optimization of profits (rewards minus costs). A typical assumption also is that the agent cannot be compelled to take these costly actions (else the problem disappears) –writing enforceable contracts is problematic. The problem of opportunism ("self-interest with guile"- Williamson, 1975) makes contracting even more complicated because organizational contracts become harder to agree on (between agents and designers because of opportunistic bargaining with misrepresentations) and implement (because of the possibility of reneging) - but the challenges of creating enforceable contracts could remain in principle even without opportunism, due to bounded rationality alone (Hodgson, 2004).

As before, we will first consider the case of independent agents. In the simplest case, if the reward structure does not cover the agent's costs of efforts (and possible disutility from risk), then the agent's assigned task will not be performed adequately. For instance if the actions are not easy to contract on, and only noisy outcomes can be observed and rewarded, we get the standard principal-agent problem (for instance, see Levinthal, 1988). Even though the agents are independent of each other, the impact of this *agency* problem on the integration of effort in the system depends crucially on the interdependence of tasks in the system. In a system with highly interdependent tasks, agency failures at individual tasks will have different implications for system performance compared to that in a system with independent tasks - regardless of the number of agents and/or tasks that the system is comprised of.

Returning to the pin factory example, the impact of a poorly designed reward structure for A will have more serious consequences with specialization than without. Consider the case where A produces 300 faulty heads compared to A producing 100 faulty pins - in the former case the day's production will fall to zero, while in the latter A's faulty output leaves B's output (of 100 pins) unaffected.

In the case of interdependence between agents, given that an agent will not take costly actions unless the rewards compensate these efforts sufficiently, the key question is how interdependence between agents affects their rewards and costs of efforts. In other words, interdependence between agents may cause an agent to be less likely to take an action either because their own reward is lowered because of the other's actions or their cost of effort is raised by the other's action. Conversely, interdependence may cause an agent to be more likely to take an action if the other agent's action either raises the rewards or lowers the costs of efforts. In game theoretic terms, interdependence between the agents may thus make their efforts strategic complements or strategic substitutes (Gibbons, 1992).

Thus a key dimension of interdependence between agents that matters for the integration of efforts is whether the returns to one agent's actions increase or decrease based on the other agent's actions. We may think of this as the valence of interdependence- whether their actions are strategic complements (positive valence of interdependence) or strategic substitutes (negative valence). The magnitude of this valence depends both on the complementarity/substitution between the agents' efforts in the production function of each agent and on the incentive breadth. [ii] To see this, consider two agents A and B each assigned a task. Assume first that A is rewarded on the performance of his own task. Let B's actions improve the returns to A's actions because their efforts are complements in A's production function. Then all else being equal, as the magnitude of this complementarity effect increases, A is more likely to invest effort in performing his task. Now consider the impact of letting B partly share in the rewards for A's task performance. Then the returns to A's actions effectively decline, and A is less likely to invest effort in his task.[iii]

When the valence of interdependence between agents is such that it lowers their willingness to invest efforts and thus impedes the integration of effort, we have a *cooperation* problem. Consider the well-known phenomenon of free riding: the issue is that adding agents decreases the probability of each agent taking the desirable action from the systemic perspective (see Prendergast, 1999: 39-44, for a review of empirical tests of free-riding). This is because adding agents dilutes the reward function for each agent without creating complementarities in their efforts- the valence of interdependence is weakened. An extensive literature in agency theory explores various solutions to this problem, such as the use of monitoring and sanctions (Alchian & Demsetz 1972), target rate incentives (Holmstrom, 1982; Petersen, 1992), and expected future interactions (Baker, Gibbons & Murphy, 2002). Scholars investigating free riding in the social psychology tradition argue that communication enables individuals to better understand the impact of their actions on individual and group outcomes through a process of discussion and learning (Dawes 1980), and that group identity induces individuals to take the group interest into account when making their own decisions (Bouas & Komorita 1996; Kollock 1998).

These diverse approaches to resolving free riding problems broadly fall into two categories: either they modify the rewards or costs of effort to compensate for the negative valence of interdependence in a multi-agent system with shared rewards (group identity, monitoring and sanctions work in this way), or they create/highlight positive valence (communication, repeated interaction or target rate incentives work in this way).

AN "INTEGRATED" APPROACH TO ACHIEVING THE INTEGRATION OF EFFORT

It is well known that knowledge problems (skill and coordination) and motivation problems (agency and cooperation) are independently sufficient to create integration failures (Camerer & Knez, 1996, 1997; Heath & Staudenmayer, 2000). For instance, consider the canonical cooperation problem - the famous prisoner's dilemma. In a standard two agent prisoner's dilemma, the existence of both a temptation payoff and a sucker penalty imply that the rewards from acting cooperatively for either agent are never large enough relative to the rewards from defecting. However, there is no epistemic interdependence between rational agents as each agent has a dominant strategy- their optimal actions do not depend on a prediction of the other agent's actions.

As others have noted, if there were strong synergies from cooperating such that the temptation payoff would disappear, the game would cease to be a prisoner's dilemma (Camerer & Knez, 1996, 1997). If the sucker's penalty still existed it would become a pure coordination problem - such as stag hunt - in which the valence of interdependence encourages cooperation, but a lack of predictive knowledge may still deter its achievement.

Note that with perfect predictive knowledge, it is still possible for the agents to coordinate successfully on a bad outcome from the perspective of the system designer-as in the case of two employees who collude to shirk on the job. To analytically separate coordination problems from cooperation problems, we believe it is important to define successful coordination independently of whether integration of effort is achieved - rather success is defined in terms of mutual predictability of action. If this mutual predictability leads the agents to take actions that maximize the value of the system as a whole, then it leads to successful integration of efforts. Thus, sufficiently positive valence of interdependence and the existence of predictive knowledge are both necessary to achieve the integration of efforts. [iv]

These analytical distinctions notwithstanding, it is probably true that in the real world, problems of organization design do not come neatly and separately packaged into knowledge and motivation problems. Interdependence between agents may simultaneously have a valence that discourages the agents from investing effort into their own tasks, as well as creates doubts about how to invest those efforts because of epistemic interdependence. One obvious consequence of this is that organization designs that ignore one or the other aspect can be disastrous.

Consider the practice in many multi-unit organizations of creating "broad" firm level incentives to encourage inter-unit collaboration, realization of synergies etc. However, to the extent that each sub-unit cannot influence the other unit's output as effectively as its own, total output may decline with such incentives, relative to the case where each unit is rewarded only for its own output. This is typically the case in organizations where sub-units are specialized to different tasks, and the coordination challenges between specialists implies that inter-unit collaboration effort is in general less productive than production effort. Kretschmer and Puranam (2008) show formally that ignoring the coordination challenges that arise when stimulating collaboration between specialized units through broad incentives not only impedes the achievement of synergies, but could lower organizational performance below the levels achieved when such synergies were simply ignored.

Integration With an Omniscient Designer

Even though coordination and cooperation problems are analytically distinct, their solutions may be interlinked in complex ways. Take the simplest case, in which the designer knows what the ideal actions of the agents should be, and agents who pursue, if not maximize (subjective) expected utility. Such agents, in a situation of epistemic interdependence, may be only able to estimate a probability distribution over the other agent's actions. If the probability of an agent taking the action desired by the designer increases in its expected utility, then it is trivial to show that the designer's investments in increasing the agents' predictive knowledge or in increasing the valence of interdependence between them are complements in terms of improving the overall integration of effort. (Also see Hardins & Higgins (1996) for a discussion of how improvements in predictive knowledge may also effectively increase the valence of interdependence). Yet in practice, organizational arrangements to improve predictive knowledge may simultaneously suppress motivation.

Consider the problem of post-merger integration in technology acquisitions- the acquisitions of small entrepreneurial firms by larger firms for their technological capabilities (Puranam et al, 2006). In such acquisitions, unlike mega-mergers, there are no gains from eliminating redundancies and consolidating administrative overheads. Yet acquirers often structurally integrate such acquisitions –fold them into existing organizational units – despite the organizational disruptions and weakening of incentives this creates. Structural integration results in common procedures, common goals, and common authority between acquired and acquiring firm's technical employees, as they are located within common organizational units. This enhances predictive knowledge as all interacting parties adhere to the same procedures, are aware of a common goal, and are directed by the same source of authority. This strong "coordination effect" however comes at the cost of disruption and de-motivation within the formerly autonomous acquired organization. These dual effects of structural integration are well documented. There is empirical evidence showing that following structural integration in technology acquisitions, the employees in the target organization may file fewer patents, though their work may be cited more often by the acquirer's employees (Puranam & Srikanth, 2007); initial products based on the target's technology may have lower hazards of being launched, but conditional on the first one being launched future versions may appear more rapidly (Puranam et al, 2006). When acquirer and target have a significant pre-acquisition overlap in technical knowledge (a basis for predictive knowledge between them), then they are less likely to engage in structural integration in the first place (Puranam et al, 2009).

Integration With a Bounded Rational Designer

Obtaining integration of effort when at least the designer knows what the agents ought to do is already challenging. This may be seen as an *execution* problem. The more general, and unfortunately more complicated case arises when even the designer does not know what the ideal actions for the agents are, so that we confront a *search* problem. Organizational life is replete with instances of specialists from different domains searching for optimal interdependent actions. For instance, managers within a multi-business company, despite incentives to pursue synergies, may have limited knowledge of the complementarities in production or the cross-elasticities of demand across divisions. Teams of engineers developing sub-systems may know that certain design choices in each sub-system could lead to dramatically enhanced performance of the system as a whole, but do not know which ones. Because the boundaries of specialization often constitute barriers to interpretation, these *joint search* problems are characterized by communication constraints arising from differences in perspectives, jargon, languages and technical backgrounds in addition to ignorance about how the key actors are actually interdependent (Lawrence and Lorsch 1967; Dougherty 1992; Heath and Staudenmayer, 2000).

One may argue that merely sensitizing the agents to the fact of their interdependence may be a major contribution that an organization designer can make towards the integration of effort. Heath and Staudenmayer (2000) present a fascinating series of examples of "partition focus" - a tendency by individuals to place more emphasis on the task division process and less to the process of achieving the integration of efforts, as well as "component focus"- a tendency to focus exclusively on the components tasks so created. Can a designer help in these situations? The existence of interdependence can only be discovered as the agents perform their actions and notice an exogenously generated change in the value of their action. Perhaps one way in which a designer may be useful is to help highlight such surprises for the agents, and rule out incorrect alternative explanations (such as bad luck).

Another possibility is that designers may provide (possibly faulty, but still useful) initial beliefs to agents searching for optimal interdependent actions. Puranam and Swamy (2010) develop a formal model of agents engaged in joint search through a trial and error (reinforcement learning) process. They show that the rates of individual learning as well as the initial beliefs of the agents influence the speed with which the discovery process is successful, in surprising ways. For instance, moderate rates of learning of the agents may lead to speedier discovery than rapid rates of learning; both agents possessing faulty initial beliefs is better than possessing no beliefs at all, or one agent possessing accurate beliefs but not the other. The common mechanism they identify underlying all these results is the *confounding effect*- the actions of one agent creating false negatives for the other agent. While maintaining the assumption that the designer faces the same ignorance about the nature of interdependence as the agents, their results suggest that a designer who can influence the learning rates and/or initial representations of the agents can be useful. For instance, the use of appropriate incentives by the designer that makes agents more responsive to immediate feedback is one way in which to control learning rates.

Even if learning rates cannot be controlled, the analysis suggests that a designer may still be useful when the learning rates are exogenously high. By imposing her own faulty initial representation on all the agents, the designer can offer an advantage relative to the cases when the agents have no representation at all, or when some have the correct representation but others have incorrect representations. By effectively coordinating the agents' state of error (rather than letting the agents be differentially knowledgeable), the designer can minimize the chances of coordination failure through the confounding effect.

CONCLUSION

We have argued that task interdependence influences the magnitude of the consequences of integration failures arising from skill and agency problems at the individual level, but cannot by itself cause integration failures. On the other hand interdependence between agents can lead to failures of integration, either because of the absence of predictive knowledge or negative valence of interdependence. While we have noted a few initial steps towards developing a deeper understanding of how knowledge and motivation problems may be jointly analyzed, it seems to us that this remains a rich area of further inquiry into the links between interdependence and organization design.

[i] When coordinated action arises from shared knowledge, it may also be of interest to understand exactly to how many orders knowledge must be shared for coordinated action. This is a question central to the research in epistemic game theory (e.g. Aumann & Brandenburger, 1995), but one which we are agnostic about here.

[ii]The magnitude of the complementarity/substitution effects in the production function can also be modeled in a reduced form as the "gains from integration" - Kretschmer & Puranam (2008). The idea is that regardless of whether there is substitution or complementarity between actions, there is value in "managing" this interdependence collaboratively-discovering it, modifying it, redistributing the costs and benefits etc. The notion of "synergy" in corporate strategy is typically used in this reduced form sense.

[iii] When the incentive structure is symmetric across A and B, then an additional tradeoff arises in that broadening incentives may make B's contributions to A's task increase but at the same time may force A to share some of the rewards. See Kretschmer and Puranam for a formal analysis.

[iv] However, they may not be jointly sufficient. Consider a stag hunt game in which each player correctly predicts that the other is sufficiently doubtful about the other player's intentions- leading to both agents selecting the lower equilibrium. What appears to be missing in this game is common knowledge of payoff structure- if this can be presumed, then predictive knowledge would lead to the selection of the high equilibrium.

FIGURE 1
Task Structure Matrix with Four Task Elements

	x1	x2	x3	x4
x1	x	x		
x2	x	x		x
x3		x	X	
x4		x		x

FIGURE 2
Thompson's Workflow Diagrams

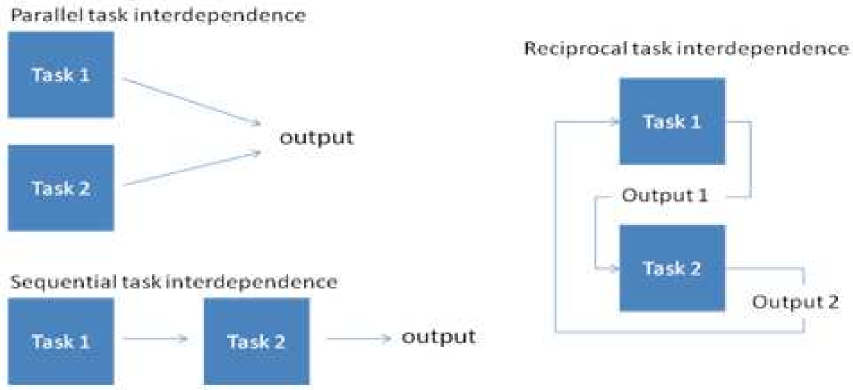

FIGURE 3
Burton & Obel's (1984) Linear Programming Model (pg.24)

$$\underset{X0,\, X1,\, X2,\, X3}{\text{Maximize}} \quad C0X0 + C1X1 + C2X2 + C3X3$$

Subject to:

$$\begin{array}{lllllllll} A0X0 & + & A1X1 & + & A2X2 & + & A3X3 & \leq & b0 \\ & & B1X1 & & & & & \leq & b1 \\ & & & & B2X2 & & & \leq & b2 \\ & & & & & & B3X3 & \leq & b3 \end{array}$$

	X0	X1	X2	X3
b0	A0	A1	A2	A3
b1		B1		
b2			B2	
b3				B3

	X0	X1	X2	X3
X0	1	1	1	1
X1	1	1	1	1
X2	1	1	1	1
X3	1	1	1	1

REFERENCES

Alchian, A.A. & Demsetz, H. 1972. Production, Information Costs, and Economic Orgnaization. ***American Economic Review***, 62(5): 777-795.

Baldwin, C.Y., & Clark, K.B. 2000. ***Design Rules***. Cambridge, MA: The MIT Press.

Baker, G., Gibbons, R., Murphy, K.J. 2002. Relational Contracts and the Theory of the Firm. ***Quarterly Journal of Economics***, 117(1): 39-84.

Bouas, K.S. & Komorita, S.S. 1996. Group Discussion and Cooperation in Social Dilemmas. Personality and Social Psychology Bulletin, 22: 1144-1150.

Burton, R.M., & Obel, B. 1984. ***Designing Efficient Organizations: Modelling and Experimentation***. Amsterdam: North-Holland.

Camerer, C. 2003. ***Behavioural Game Theory: Experiments in Strategic Interaction***. Princeton: Princeton University Press.

Camerer, C. & Knez, M. 1996. Coordination, Organizational Boundaries and Fads in Business Practices. ***Industrial and Corporate Change,*** 5(1): 89-112.

Clark, H. 1996. ***Using Language***. Cambridge, UK: Cambridge University Press.

Cohen, M.D. & Bacdayan, P. 1994. Organizational Routines are Stored as Procedural Memory: Evidence from a Laboratory Study. ***Organization Science***, 5(4): 554-568.

Dosi, G., Levinthal, D. A., & Marengo, L. 2003. Bridging contested terrain: linking incentive-based and learning perspectives on organizational evolution. Industrial and Corporate Change, 12(2): 413-436

Emerson, R.M. 1962. Power-Dependence Relations. ***American Sociological Review***, 27(1): 31-41.

Eppinger, S.D. 1991. Model-based Approaches to Managing Concurrent Engineering. ***Journal of Engineering Design,*** 2: 283-290.

Galbraith, J. R. 1973. ***Designing Complex Organizations***. Reading, MA: Addison-Wesley.

Gibbons, R. 1992. ***Game Theory for Applied Economists***. Princeton, NJ: Princeton University Press.

Goetting, M., Puranam, P. & Warglien, M. 2011. ***The Emergence of the Division of Labor: A Qualitative Lab Study***. Working Paper, London Business School, London, UK.

Grandori, A. & Furnari, S. 2008. A Chemistry of Organization: Combinatory Analysis and Design. ***Organization Studies***, 29: 459-485.

Guilford.Heath, C., & Staudenmayer, N. 2000. Coordination Neglect: How Lay theories of organizing complicate coordination in organizations. ***Research in Organizational Behavior***, 22: 53-191.

Hardin, C. D., & Higgins, E. T. (1996). Shared reality: How social verification makes the subjective objective. In E. T. Higgins & R. M. Sorrentino (Eds.), ***Handbook of motivation and cognition: The interpersonal context (Vol. 3).*** New York:

Hodgson, G.M. 2004. Opportunism is not the Only Reason why Firms Exist: Why an Explanatory Emphasis on Opportunism May Mislead Management Strategy. ***Industrial and Corporate Change***, 13(2): 401-418.

Holmstrom, B. 1982. Moral hazard in teams. ***Bell Journal of Economics***, 13(2): 324-340.

Hoopes, D.G., & Postrel, S. 1999 Shared Knowledge, "Glitches", and Product Development Performance. ***Strategic Management Journal***, 20(9): 837-865.

Kelley, H.H., & Thibaut, J.W. 1978. ***Interpersonal Relations: A Theory of Interdependence***. New York: John Wiley & Son.

Kollock, P. 1998. Social Dilemmas: The Anatomy of Cooperation. ***Annual Review of Sociology***, 24(1): 183.

Kretschmer, T., & Puranam, P. 2008. Integration through Incentives in Differentiated Organizations. ***Organization Science.***, 19(6): 860-875.

Lawrence, P.R., & Lorsch, J.W. 1967. ***Organization and Environment: Managing Differentiation and Integration.*** Boston: Harvard Graduate School of Business Administration.

Leijonhufvud, A. 1986. Capitalism and the Factory System. in Langlois, R.N. (ed.) ***Economic as a Process: Essays in the New Institutional Economics***, New York: Cambridge University Press, pp. 203-223.

Levinthal, D. A. 1997. Adaptation on Rugged Landscapes. ***Management Science,*** 43(7): 934.

Levinthal, D. 1988. A Survey of Agency Models of Organizations. ***Journal of Economics Behavior and Organizations***, 9: 153-185.

Lewis, D. 1969 - ***Convention: A Philosophical Study***. Cambridge, MA: Harvard University Press.

March, J.G., & Simon, H.A. 1958. ***Organizations***. Cambridge, MA: Wiley.

Milgrom, P. & Roberts, J. 1990. The economics of modern manufacturing: Technology, strategy, and organization." ***American Economic Review*** 80(3): 511-28.

Nelson & Winter, 1982. ***An Evolutionary Theory of Economic Change***. Cambridge: Harvard University Press.

Newell, A. & Simon, H. 1972. ***Human Problem Solving***, Englewood Cliffs, NJ: Prentice Hill.

Petersen, 1992 Individual, Collective, and Systems Rationality in Work Groups: Dilemmas and Market–type Solutions. ***American Journal of Sociology***, 98(3): 469–510

Pfeffer, J. & Salancik, G.R. 1978. ***The External Control of Organizations: A Resource Dependence Perspective***. New York: Harper & Row.

Prendergast, C. 1999. The Provision of Incentives in Firms. ***Journal of Economic Literature***, 37: 7-63.

Puranam, P., Goetting, M. & Knudsen, T. 2011. ***Organization Design: The Epistemic Interdependence Perspective.*** Working Paper, London Business School, London, UK.

Puranam, P. & Swamy, M. 2010. ***Expedition without Maps: Why faulty initial representations may be useful in joint disovery problems,*** Working Paper, London Business School, London, UK. ssrn

Puranam, P., Singh, H., Chaudhuri, S. 2009. Integrating Acquired Capabilities: When Structural Integration is (Un)necessary. ***Organization Science***, 20(2): 313-328.

Puranam, P. & Srikanth. K. 2007. What They Know vs. What They Do: How acquirers leverage technology acquisitions. ***Strategic Management Journal***, 28(8): 805-825.

Puranam, P., Singh, H., Zollo, M. 2006. Organizing for Innovation: Managing the coordination-autonomy dilemma in technology acquisitions. ***Academy of Management Journal***, 49(2): 263-280.

Schelling, T. 1960. ***The Strategy of Conflict***. Cambridge, MA: Harvard University Press.

Simon, H. A. 1945. ***Administrative Behavior***. New York: Macmillan Press.

Smith, A. 1776. ***An Inquiry into the Nature and Causes of the Wealth of Nations***. Edinburgh.

Srikanth, K. & Puranam, P. 2010. ***Integrating Distributed Work: Comparing Task Design, Communication and Tacit Coordination Mechanisms, Strategic Management Journal***, *Forthcoming.*

Steward, D.V. 1981. ***Systems Analysis and Management: Structure, Strategy and Design***. New York: Petrocelli Books.

Thompson, J.D. 1967. ***Organizations in Action***. New York: McGraw-Hill.

Tushman, M.L., & Nadler, D.A. 1978. Information Processing as an Integrating Concept in Organizational Design. ***Academy of Management Review***, 3(3): 613-624.

Von Neumann, J. & Morgenstern, O. 1944 ***Theory of Games and Economic Behavior***. Princeton University Press.

Williamson, O.E. 1975. ***Markets and hierarchies: Analysis and antitrust implications.*** New York: Free Press.