Industrial and Corporate Change, Volume 12, Number 5, pp. 1099–1124

The flow of ideas and timing of evaluation as determinants of knowledge creation

Sridhar Seshadri and Zur Shapira

There are different perspectives on the study of knowledge in organizations, developed in economics, sociology, anthropology and organization theory. Several authors followed Schumpeter's idea that innovations are new combinations of existing knowledge and incremental learning. Kogut and Zander further developed this idea and defined knowledge as a portfolio of options, and emphasized the importance of combinative capacities in knowledge creation. In a similar vein, Garud and Nayyar developed a notion of the transformative capacitiy of a firm in analyzing technological innovations. This paper follows in this tradition by pointing at the effects of different organizational structures on the flow of ideas and on the possibility of combining proposals in organizations. Communication structures, span of control and timing of evaluation are shown, using computer simulations, to have a large impact on the degree to which commonalities and complementarities among ideas and proposals can be detected and on the eventual combination of ideas for knowledge creation. Implications for organizational design are discussed.

1. Introduction

The knowledge revolution has had a large impact on research in organization theory and strategic management. Research on knowledge creation and knowledge transfer proliferated in the last decade. Much of the discussion on organizational design moved away from static ideas to dynamic notions. More dynamic models have replaced old models that portrayed organizations as passive reactive systems, with models where organizations are considered as acting on their environments (Cyert and March, 1992) and enacting them (Weick, 1979). Researchers dealing with knowledge creation and knowledge transfer moved from the individual level of analysis to focusing on the group and the organizational levels of analysis. New models of organizational learning acknowledge the origin of individual cognition and learning and examine whether organizations can cognize, learn and create knowledge (March, 1991, 1999; Lant and Shapira, 2000). The salience of knowledge creation is evident in studies in the areas of organizational learning (March, 1991; Schulz, 2001), resource based view in strategy (Garud and Nayyar, 1994) and social capital (Cohen and Prusak, 2001).

There are different perspectives on the study of knowledge in organizations. In economics, for example, Schumpeter's (1934) ideas have been rejuvenated and recent writings by organizational economists acknowledge the idea of specific knowledge (Jensen and Meckling, 1992). The anthropological/sociological perspective focuses on the instrumentality of groups and communities in fostering innovation and creativity in organizational setting, culminating in ideas about organizational structure and design (Brown and Duguid, 1991). These perspectives emphasize different principles around which organizations should be designed to foster knowledge creation and knowledge transfer, which are important determinants of the strategic behavior of firms. Cohen and Levinthal's (1990) notion of absorptive capacities has been used to describe the firm's readiness to accept new knowledge, and these capacities relate to the firm's ability to add new knowledge to existing knowledge (Grant, 1996).

The ideas of technology transfer and imitation have been discussed, among others, by Nelson and Winter (1982) and Kogut and Zander (1992). In a thorough analysis, Kogut and Zander focus on the tacit dimension of knowledge. Following Polanyi's (1966) dictum, they distinguish between information and 'know how', and claim that the latter is the dimension of importance for organizational knowledge creation. They argue that 'the knowledge of a firm can be considered as owning a portfolio of options, or platforms, on future developments' (Kogut and Zander, 1992: 385).

The focus on knowledge aggregation and on portfolios of ideas is of major importance in knowledge creation processes in organizations. Indeed, the very thing that distinguishes organizations from single individuals is the ability to draw ideas from different individuals and put up and create combined proposals that may give the organization an advantage. Attempting to create projects that combine elements from different parts of an organization is not easy. Even at the individual level, combining ideas from different sources in a creative manner requires integrated efforts by the individual and other persons. Work by de Bono (1986) emphasized the need to consider relationships among ideas, what he called lateral thinking, to be able to solve problems. However, the aggregation over knowledge and the creation of portfolios of projects and ideas is not easily done. Research on organizational processes suggests that there are several cognitive and structural aspects that operate as barriers against attempts to make use of ideas in different parts of the organization.

We follow the argument that 'knowledge is a portfolio of options' (Kogut and Zander, 1992; Garud and Nayyar, 1994) and examine what can be done to deal with organizational structural aspects and procedures that create barriers against the creation of portfolios of options, ideas and projects. Two issues stand out in our analysis of knowledge creation: the flow of ideas or proposals and the timing of evaluation and decisions on these proposals. We examine how different organizational structures facilitate or inhibit the flow of ideas and the eventual possibility that single ideas get combined together to form potentially more successful projects.

2. The flow and combination of ideas and knowledge creation

Organic structures allow ideas to flow in organizations. However, to create knowledge and generate projects, some mechanisms that facilitate the matching and combination of ideas are needed. In attempting to facilitate this matching, the issue of timing becomes critical, as ideas need to be considered in proximity of time to be candidates for matching and eventual combination. An example that illustrates such a process is the invention of the Walkman at the Sony Corporation.

In an account of the development of the Walkman, Nayak and Ketterimgham (1986) describe the processes that went in parallel at two different departments at Sony. The tape recorder department was working on improving the Pressman, a small tape recorder that Sony produced in 1978. The tape recorder department's engineers were busy trying to minimize the size of the tape recorder while assuring it had stereophonic qualities and would be able to play as well as record. It was the latter function they had difficulties with when they attempted to minimize the size of the instrument. Not succeeding in their endeavor, they left the tape recorder around and used it to play their favorite music cassettes while working. The team turned their non-recording tape recorder into a background music player and almost ignored it. When Ibuka, Sony co-founder and honorary chairman, came to visit this group he noticed the little device that played high quality music. He also remembered that an engineer in another, unrelated department was working on developing lightweight portable headphones. 'What if we combine them?' asked Ibuka, and was met with skepticism.

Ibuka's idea would have at least reduced the needed battery power to operate the larger speakers. It actually led to developing the portable headphones with higher sound quality and to doing away with the recording function. As Nayak and Ketteringham (1986: 134) comment: 'In the world of tape recorders, Ibuka's thought was heresy. He was mixing up functions.' In Sony there was no reason for tape recorder people to ever communicate with headphone people. As one of them is quoted as having said 'We're not very interested in what they do in the Headphone Division' (Nayak and Ketteringham, 1986: 135). Indeed, if it were not for Ibuka, the Walkman idea, which led to the most popular entertainment instrument (for an individual person) in history, might have never seen fruition.

Such eventually successful but random events capture the attention of managers in creative firms, who attempt to create some structures or procedures that will facilitate interaction and cross pollination among people and departments. The notion of communities of practice (Brown and Duguid, 1991) is in this spirit, as well as the dynamic and loose control culture of IDEO, the innovative product design-consulting firm (Sutton and Hargadon, 1996).

2.1 The flow of ideas, timing and knowledge creation

We assume that in dynamic organizations there are streams of ideas generated at the individual as well as at the group level. The task of the organization is how to create procedures, routines and structures where related ideas have a chance to get combined into projects, as described in the Walkman example above. Firms are looking for ways that facilitate their combinative capabilities that are embedded in the organizing principles by which people cooperate in organizations. To increase the ability of managers to observe areas where features of different proposals or ideas overlap, timing becomes an important issue. Going back to the Sony example, the timing for Walkman was ripe because the two parts of the Walkman were 'in process'. It is not clear what would have happened had the tape recorder department's engineers been successful in plugging the recording function into the instrument. The Walkman example shows how recombination of elements can also include the deletion of elements. In this paper we refer to recombination of projects as modification in a general sense, that is, projects can be recombined so that the number of elements increase and they can be recombined while some elements get removed. Thus the term recombination actually includes modification.

The issue of timing is rather complex. A supervisor who receives project proposals from her subordinates searches for some complementarities among them. If she discovers such features, she may be able to combine proposals. This very attempt may lead to delays in moving project proposals up the hierarchy and may lead to missed opportunities if other firms are faster to the market with similar products. At other times the ability of a supervisor to hold on to a project proposal until another proposal arrives that has complementarities with the first project proposal, may result in a more profitable project (see Garud and Nayyar, 1994). We initially consider the delay to be a detriment to the efficient processing proposals but in a later section relax this constraint and consider the benefits of deliberate delays.

We extend the notion of combinative capabilities to a microanalysis of decision processes in organizations. We believe that combining ideas fits nicely with the notions raised by Kogut and Zander and by Schumpeter, but we also acknowledge the reluctance of managers to combine proposals that emanate from cognitive, timing and other organizational aspects. We further assume that there is a flow of ideas in an organization in the form of ideas or proposals that are presented by managers to their supervisors. It is at that level that the main obstacle to the combination of ideas, exist. Two issues are important in enhancing the combination of ideas (proposals): the commonalities and complementarities between proposals that describe the chances that they can be combined; and the timing of review and evaluation of these proposals, assuming that synchronization of ideas and proposals is facilitated if they are evaluated simultaneously by a supervisor.

3. A model of flow and combining of ideas

What is the role of a hierarchical structure in evaluating project proposals? The supervisors at every level obtain a number of proposals. They possibly go over some basic checks such as practicality of objectives, time frame and resource required. They also ascertain whether the proposals fit within a broader set of goals. Another possible role is to seek clarification and give advice on special features that they are familiar with. They could also look at parts of the proposal and see if they fit. These roles are of an auditory nature. We suggest still another role that supervisors sometime play, that is to combine project proposals or to see complementarities in developing certain sets of expertise within the organization. Once such a process begins, it could probably result in a still bigger idea, like a snowball rolling down a hill. We formally model such a process using three elements: namely, the potential for proposals (ideas) to be combined, the processing efficiency of combining proposals (ideas) and the processing delay in combining proposals (ideas).

Using these three elements, we show that it is possible to construct step-by-step the process by which ideas are generated, examined, combined and communicated within a given organization structure. In this section, the detailed modeling framework is presented. In the next section, the framework is used to explore the effect of several factors, such as delays, timing, communications and the number of levels in an organizational hierarchy, on the propensity of an organization to produce rich projects that combine ideas from different parts of the organization. We assume that:

1. The organization has a fixed structure, namely, it does not change during the time in which proposals are being evaluated or combined. All project proposals (ideas) are generated at the lowest level of the organization. We shall use the label ideas (or proposals) for the proposals and project for the final output. All proposals are sent to the next level for processing after scrutiny.
2. Proposals are submitted by subordinates to their immediate and direct supervisor (manager) for evaluation. A manager scrutinizes all incoming proposals submitted by his/her subordinates and has the choice of combining any two of these when there are complementarities between them.
3. Ideas submitted by the same subordinate cannot be combined at a higher level.
4. There is a fixed probability that two ideas submitted by different subordinates have complementarities. If a proposal has no commonalities/complementarities with other proposals, then it is immediately passed to that manager's immediate supervisor for evaluation.
5. A manager can combine a proposal with only one other proposal (submitted by a different subordinate). This is elaborated upon below.
6. Combining ideas either takes no time with a given probability, or takes one unit of time. This delay can be avoided by sacrificing efficiency, that is, by accepting a lower probability of detecting complementarities between ideas as discussed below.
7. Proposals arrive at different discrete times (0, 1, 2, . . .) to a supervisor. Only those proposals that arrive at the same time to a supervisor can be combined.

We now illustrate how this framework can be applied to model the flow and combining of ideas in a two-level hierarchy. Consider two organizations. In organization 1, two managers, say A and B, report to a supervisor C. Each of A and B considers two proposals. The probability that the two ideas of A (likewise of B) have complementarities and thus can be combined is p, independent of everything else. Each manager, A and B, examines a set of two proposals and determines whether they can be combined. If they can be combined, then to do so it either takes no time (zero time) with probability q or one unit of time (say a week) with probability $(1 - q)$. Therefore, manager A sends to C either one (combined) project proposal or two separate proposals immediately, with probabilities pq and $(1 - p)$ respectively; or one combined proposal after one week with probability $p(1 - q)$. Manager B does the same.

Now consider the decision problem of supervisor C. She obtains 1 or 2 proposals from managers A and B without delay, or one (combined) proposal from A and/or one (combined) proposal from B after a week. Ideas of A that were not combined earlier cannot be combined by C. Similarly, ideas of B that were not combined earlier cannot be combined by C. Ideas that do not arrive at the same time cannot be combined. For example, if A sends up all his proposals with no delay and B sends all his proposals with delay, then C cannot combine any of them. However, proposals across managers A and B that arrive at the same time can be combined if they have some commonalities/complementarities. A specific manager cannot combine more than two ideas together. The probability that two proposals across managers can be combined is also p, independent of everything else.

For example: if there are four ideas perused by C, she can combine any two of them if they arrive from two subordinates and have complementarities, but not three or four. This assumption is not unrealistic, in fact it is consistent with the notions that managers possess limited ability for synthesizing ideas due to interruptions, lack of specialized knowledge and data processing limitations (Seshadri and Shapira, 2001). In this example, the outcome is one of the following: (i) four original proposals; (ii) two proposals: where each is comprised of two original proposals; or (iii) three proposals with one comprised of two original ideas and the other two as they were before perusal.

There are many reasons why combining ideas is of importance to organizations: the first is the scale and scope of the project. Organizations incur certain fixed overheads to process and support a project during its lifetime. The project potential should be large enough to cover such overheads. Thus, there is a hurdle, only beyond which projects with greater potential for making impact are more easily supported. For example, a large corporation with each division having sales of several billion dollars would hardly be able to focus and support a project with a potential impact of a few hundred thousand dollars! At the same time, recombination of two projects can lead to an aggregate project that is actually smaller than the original two. In this case, the combination of two projects can lead to a more ‘lean and mean’ project that saves the organizations cost. A second reason is that a combined proposal is usually multi-disciplinary and thus has a greater chance of success due to the balance of views (which could sometimes also prove detrimental). Finally, a combination of ideas is not simply an amalgamation (mechanical addition) of independent projects but a true combination so that both expected returns and variance of returns get scaled. Therefore, it offers greater chances of windfall profits. For example, we often observe ideas becoming enriched due to dialog, especially, dialog amongst participants with diverse backgrounds, motivations, and skill levels. Thus, true combination of ideas is similar to creating a richer tapestry. Another metaphor we offer is that of bubbles (ideas) filtering up through the hierarchy; smaller bubbles combine with others at times to produce the spectacular large ones at times. The bubbles have to come together at the same time and should also be able to stick together. We model such a phenomenon. Additionally, we predicate that the variance gets scaled when proposals are combined, whereas returns on uncombined proposals are independent. Further, it is obvious that if only projects with certain minimal returns will prove profitable then the potential impact of the project matters. The foregoing is not to imply that only projects that are combination of ideas will be successful, but that when there is potential for combining ideas to increase the final impact, that is, to enhance the value of the separate ideas, then the firm benefits.

From the viewpoint of assessing performance, we propose to study three metrics: the average number of ideas combined per final project, the variance of the number of ideas combined, and the average processing delay in the system. The processing delay is computed by weighting the delay of each project by the number of ideas combined and dividing the weighted sum by the number of final projects produced.

3.1 *Delay and flow of ideas*

We first examine the performance criteria for the four-level hierarchy shown in Figure 1. In this hierarchy, there are 16 managers at the bottom-most or first level, and each mamager generates two ideas. These managers are grouped in fours, each group of four reporting to a manager at the second level. Thus, there are four managers at the second level. The second level managers are grouped in twos, each group reporting to a manager at the third level; and the third level managers report to the single top manager at the fourth level. The flow of ideas is simulated 5000 times. In each simulation, the 32 proposals are combined by applying the rules for processing, combining and delay explained earlier. The simulation is carried out by initializing the bottom level managers with two proposals each. The processing is done level by level. The proposals

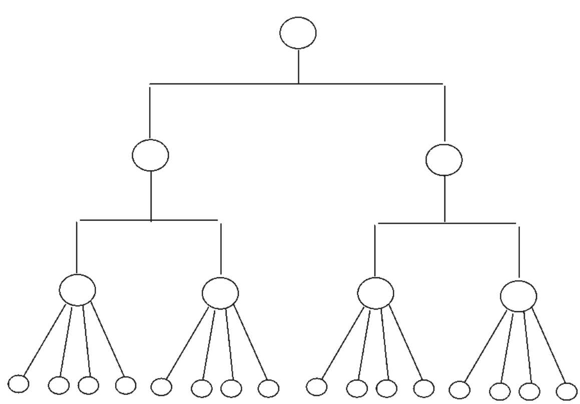

Figure 1 Four level hierarchy.

either single or combined ones move to the next level. The statistics are collected when the proposals are processed by the level four managers. Records are kept regarding the number of ideas combined in the (final) projects and the processing delays if any. For example, when $p = 0.3$, $q = 0.5$, the distribution of the number of ideas combined is as shown in Table 1. We see that on average 3.87 projects consisting of a single proposal are processed with zero delay, 2.85 projects wherein two ideas are combined with delay of one period, etc. Projects that consist of a single proposal are never delayed because they are never combined. Only combining project proposals takes time. It is clear that the projects with more ideas combined suffer the most delays. At the bottom of the table we show the expected number of projects regardless of the number of ideas combined that are delayed by 0, 1, 2 or 3 periods. More than half the projects are delayed by one period or more. The average number of ideas combined in a project is 2.59 and the variance of this quantity is 9.25. The average delay is 0.60 periods.

We now study the effect of delay on the performance metrics by fixing the value of p and varying q. The results are tabulated in Tables 2 and 3. In Table 2, we show the distribution of delay. We see that if combining ideas always takes one period (probability of delay = 1), then the maximum number of proposals combined cannot exceed eight due to the delay. This highlights the effect of lock-step flow of combined proposals in a system that always introduces delays when combining project proposals. We see from Table 3 that as the probability of delay decreases the variance in the number of proposals combined grows significantly and the average number of proposals combined increases too. These effects are primarily due to processing of proposals having little or no effect at higher levels when there is small or no delay in combining proposals.

We should also note that it is possible that even if there is sometimes a long delay,

Table 1 Four level hierarchy, $p = 0.3$, $q = 0.5$: distribution of proposals combined and delays

Number of proposals combined	0	1	2	3
1.00000	3.86960			
2.00000	1.20300	2.85160		
3.00000	0.60280	0.87180		
4.00000	0.32440	0.83420	0.58660	
5.00000	0.17080	0.38480	0.24360	
6.00000	0.06780	0.21140	0.15480	
7.00000	0.02520	0.08600	0.06100	
8.00000	0.00780	0.02860	0.02720	0.00280
9.00000	0.00300	0.00960	0.00800	0.00160
10.00000	0.00040	0.00280	0.00160	0.00080
11.00000	0.00020	0.00080	0.00080	0.00040
Total	6.27500	5.28160	1.08360	0.00560

Table 2 Expected number of projects in a four level hierarchy, $p = 0.3$

Number of proposals combined	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	4.81640	4.33080	3.86960	2.68380
2	5.72220	4.86980	4.05460	1.58320
3		0.99940	1.47460	1.18360
4	3.48560	2.30760	1.74520	1.03880
5		0.55220	0.79920	0.84760
6		0.27880	0.43400	0.68460
7		0.06360	0.17220	0.52120
8	0.22460	0.07080	0.06640	0.36520
9		0.02200	0.02220	0.20240
10		0.00420	0.00560	0.09780
11		0.00120	0.00220	0.04760
12		0.00020		0.01360
13				0.00300
14				0.00020

Table 3 Four level hierarchy: delay and performance metrics, $p = 0.3$

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
1.00	4.82	5.72	3.49	0.22	2.29	7.15	0.96
0.75	5.40	5.76	2.29	0.05	2.42	7.99	0.80
0.50	6.28	5.28	1.08	0.01	2.59	9.25	0.60
0.00	9.27				3.55	18.75	0.00

a terrific project can emerge. Based on this, we may consider the idea of optimal delay. For example, ensuring that proposals get through the different levels quickly requires expenditure of effort by managers and consumption of other organizational resources. The increase in effort required could be more than proportional to the reduction in delay. This costly effort is offset by the production of a richer set of project proposals. Thus, there is a clear trade-off between costly effort and generation of viable projects.

In Table 4, we show the effect of diminished probability of combining project proposals on the expected number of ideas combined. When $p = 0.1$, unless the probability of delay is quite low, projects with a large number of ideas combined are quite unlikely. Thus, when it is harder to discover commonalities and complementarities, it becomes harder to create such project proposals unless extraordinary efforts are made to move the proposals along speedily. In Table 5, we contrast the cases when $p = 0.1$ and $p = 0.5$. Apart from the obvious effect of more ideas combined and greater variance of ideas combined when p is high, notice that delays too escalate. This is due to proposals being combined more frequently. This leads to lower efficiency due to the hierarchy processing smaller number of proposals per unit time. In some sense, this stickiness amongst project proposals vitiates the specialization characteristic of a hierarchy.

Hierarchies appear to serve another purpose than envisioned previously, namely, they create richer projects. To illustrate this effect, we compare the performance metrics for two other hierarchies, a three and a two level structure. In the three level structure there are 16 managers as before in level 1, in groups of four managers. Each group reports to a manager at the second level, thus there are four level 2 managers. These in turn report to a single top level (level 3) manager. In the two level structure all 16 managers at level 1 report to a single level 2 manager. The number of ideas initiated is 32 (2 per level 1 manager). We illustrate for $p = 0.3$. In Table 6 we show the distribution of the number of ideas combined for the three and two-level hierarchies. Clearly, the probability of delay still affects the number of ideas combined. However, delay does not affect the distribution as much in a flat organization. Flat organizations also do not

Table 4 Expected number of projects in a four level hierarchy, $p = 0.1$

Number of proposals combined	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	9.96920	9.63020	9.12980	8.26660
2	9.16660	8.17000	7.06660	4.41700
3		1.00580	1.80720	2.41460
4	0.92280	0.63700	0.62620	1.13900
5		0.07400	0.12680	0.41780
6		0.01380	0.02540	0.12460
7		0.00120	0.00300	0.03020
8	0.00080	0.00040	0.00040	0.00600
9				0.00040
10				0.00000
11				0.00000

Table 5 Performance metrics for four level hierarchy, $p = 0.1$ and 0.5

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
			$p = 0.1$				
1.00	9.97	9.17	0.92	0.00	1.61	3.11	0.56
0.75	11.28	7.78	0.47	0.00	1.65	3.33	0.45
0.50	12.69	5.93	0.17	0.00	1.72	3.69	0.34
0.00	16.82				1.93	5.05	0.00
			$p = 0.5$				
1.00	3.12	3.73	3.34	1.01	2.94	12.86	1.24
0.75	3.44	3.99	2.81	0.31	3.12	14.02	1.03
0.50	4.02	4.04	1.62	0.04	3.38	16.38	0.78
0.00	6.21				5.33	42.03	0.00

Table 6 Expected number of projects in a three and two level hierarchy, $p = 0.3$

Number of proposals	Probability of delay in combining proposals							
	Three level				Two level			
	1	0.75	0.5	0	1	0.75	0.5	0
1	3.05600	2.90920	2.80420	2.54140	2.46280	2.48720	2.48520	2.39540
2	7.57360	6.47740	5.57000	2.78740	11.83580	11.11920	10.31520	8.11840
3		1.38820	2.24840	2.92500		1.10400	2.08360	3.83940
4	3.41160	2.45320	1.98920	1.88800	1.46640	0.99060	0.65840	0.46240
5		0.34060	0.51840	0.98020				
6		0.06600	0.11340	0.35320				
7		0.00600	0.00980	0.07140				
8	0.01880	0.00220	0.00160	0.00460				

produce projects that are made of great many individual proposals. In Table 7 we show the performance metrics for the two hierarchical structures. We notice a remarkable fact: the average number of ideas combined seems to be the same (except when delay

Table 7 Performance metrics for three and two level hierarchy, $p = 0.3$

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
			Three level				
1.00	3.06	7.57	3.41	0.02	2.30	6.46	1.04
0.75	4.46	7.22	1.96	0.00	2.37	6.88	0.83
0.50	6.30	6.16	0.80	0.00	2.44	7.31	0.59
0.00	11.55				2.81	9.88	0.00
			Two level				
1.00	2.46	11.84	1.47	0.00	2.04	4.71	0.94
0.75	5.15	9.84	0.71	0.00	2.05	4.69	0.72
0.50	8.07	7.22	0.26	0.00	2.07	4.75	0.50
0.00	14.82				2.17	5.24	0.00

equals zero) for all three hierarchies! The average delay also seems to be the same. The flatter structures seem to excel at producing projects with medium number of ideas combined whereas the deeper hierarchy combines the most ideas. This accounts for the substantial difference in variance.

A third aspect we now attempt to model is the role of informal communication in organizations. The viewpoint here is that such communication helps discover complementarities and commonalities in proposals, albeit, at a lower efficiency than through the hierarchy. There is another role for such interactions that we do not model, namely, through such informal discussion ideas become enriched through dialogue. The metaphor we offer is again that of weaving a tapestry. Parts of the organization weave a different part. But when they interact the whole tapestry is somehow richer and more meaningful.

We model communication by allowing managers at the same level to share their proposals and thus discover scope for combining ideas. This is done through specifying a parameter, p_1, which is the probability that two project proposals at a given level but with different managers have commonalities/complementarities independent of everything else. Thus, the model is unchanged except that when a manager does not find proposals that can be combined amongst the ones submitted to him or her, s/he also scans the proposals with other managers at his or her level. We study the effect of communicating in this fashion either at all levels or only at level 1, the bottommost level.

In Tables 8 and 9 we show the distribution of ideas combined under these two

Table 8 Expected number of projects in a four level hierarchy, $p = 0.3$, $p_1 = 0.01$, all levels communicate

Number of proposals	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	4.29320	3.80740	3.34160	2.42740
2	5.45340	4.66640	3.94080	1.49860
3		0.99840	1.47260	1.16580
4	3.60040	2.39540	1.78460	0.95520
5		0.57920	0.79980	0.82800
6		0.30280	0.45740	0.67080
7		0.08040	0.19120	0.52360
8	0.29980	0.08680	0.09040	0.38160
9		0.02360	0.03200	0.23980
10		0.00780	0.00920	0.13260
11		0.00180	0.00280	0.05240
12			0.00040	0.01840
13				0.00520
14		0.00020		0.00160
15				0.00020
16				

assumptions. The probability of being able to find complementarities across proposals initiated by different managers is 0.01. Despite such a low probability there is significant increase in the expected number of the ideas combined. When communication is allowed at all levels, the expected number of projects with eight or more proposals combined increase by 0.29900, 0.11980, 0.13440 and 0.82540 respectively for the four cases of delay (delay with probability 1, 0.75, 0.25 and 0, respectively). These increases look even more significant considering that the expected numbers of projects with eight or more ideas combined without communication were 0.2246, 0.0984, 0.0964 and 0.7298 respectively. Thus, increases are between 13 and 40%. What is possibly even more remarkable that when communication is allowed only at the lowest level, the increases are larger in every case—0.36860, 0.15500, 0.12660 and 0.36140. These increases are in the range of 50–164%.

In Table 10 we show the performance metrics under communication. Under full communication, the average number of proposals combined has increased by 4%, the variance by 7–9% and the delay (when there probability of delay is greater than zero) by 4–5%. Under communication at level 1 only, the average number of proposals combined has increased between 18 and 22%, the variance between 34 and 51% but the

Table 9 Expected number of projects in a four level hierarchy, $p = 0.3$, $p_1 = 0.01$, level 1 communicates

Number of proposals	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	2.72640	2.39580	2.09040	1.64280
2	4.36760	3.91700	3.35600	1.21900
3		0.92060	1.34980	1.00720
4	3.94820	2.59780	1.84200	0.87180
5		0.66900	0.90620	0.75780
6		0.38220	0.58060	0.63600
7		0.12160	0.26780	0.55560
8	0.59320	0.18020	0.14260	0.43840
9		0.05100	0.06000	0.31880
10		0.01760	0.01540	0.19320
11		0.00420	0.00400	0.09460
12		0.00020	0.00080	0.03540
13		0.00020	0.00020	0.00920
14				0.00120
15				0.00040
16				

delays also increased (when delay probability is non-zero) by 21–27%. It seems rather extraordinary that permitting communication at the lower level seems to be more effective than communication at all levels. We traced a few simulations to understand this effect. We found that the particular effect was caused not due to communication alone but also due to the rule for searching for combining proposals. When communication is allowed at all levels, proposals get combined at higher levels across managers that are ‘far apart’ in the sense between managers 1 and 4. Thus, big project proposals sometimes get pushed apart and thus are less likely to combine again. When communication is allowed only at the lowest level, if big proposals are ‘close’ originally then they continue to stay close together. This seems to point back at the analogy of snowball effect or bubbles filtering through a hierarchy. They not only have to move together to stick to one another, but also have to be close enough to get stuck.

3.2 *The effects of heterogeneity in the ability to detect and combine proposals*

In this section we explore another aspect of organizations, namely, the probability of combining proposals could vary due to different reasons such as the capabilities of the person who peruses proposals. Thus, first we allow variations in p. We expect that such

Table 10 Four level hierarchy: delay and performance metrics, $p = 0.3$, $p_1 = 0.01$, communication at all levels and level 1

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
			Communication at all levels				
1.00	4.29	5.45	3.60	0.30	2.38	7.81	1.01
0.75	4.86	5.56	2.46	0.07	2.51	8.60	0.84
0.50	5.78	5.19	1.14	0.01	2.69	9.97	0.63
0.00	8.90				3.69	20.24	0.00
			Communication at level 1				
1.00	2.73	4.37	3.95	0.59	2.79	10.77	1.23
0.75	3.40	4.86	2.86	0.13	2.89	11.21	0.99
0.50	4.36	4.84	1.40	0.01	3.06	12.63	0.73
0.00	7.78				4.19	25.27	0.00

variations affect the distribution of the number of proposals combined. The variations are modeled by permitting p to be uniformly distributed in the interval [0.2,0.4]. Therefore, the average value of p is 0.3. The simulation is performed as follows. Select a value of p within this range for each manager and then process the 32 project proposals. This is repeated 5000 times. The results are reported in Tables 11–14. In Tables 11 and 12 there is no communication. These results should be compared with those in Tables 2 and 3. We see that the average number of proposals combined has gone down slightly and the variance has also reduced slightly. The effect is a bit more pronounced when there is no delay than when there is delay. Thus, a little heterogeneity seems to have very little effect on the distribution of proposals combined unless the organization is already very efficient. By efficiency we mean that there are no delays in the organization to begin with. In that case, heterogeneity might slightly reduce the average number of proposals combined.

The effect of heterogeneity is even less pronounced when there is communication, see Tables 13 and 14 and compare them with Tables 8 and 10 (top panel). This is to be expected. Communication allows the managers with higher combinative capability to look at other managers' proposals. Thus, the results with and without heterogeneity are very similar. In fact in one case, when the probability of delay equals one, the variance actually goes up unlike the case when there is no communication.

We also examine the expected number of projects with eight or more ideas combined. These are virtually unchanged when heterogeneity is introduced. In the case

Table 11 Expected number of projects in a four level hierarchy, *p* varied in the range [0.2,0.4], no communication

Number of proposals	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	4.88800	4.40860	3.97020	2.77200
2	5.78360	4.94440	4.14980	1.62220
3		1.01640	1.49500	1.22720
4	3.42580	2.28620	1.70840	1.01280
5		0.55020	0.79180	0.86060
6		0.24920	0.41420	0.67900
7		0.07140	0.16380	0.51180
8	0.23020	0.06900	0.06880	0.36140
9		0.01660	0.02300	0.20120
10		0.00500	0.00520	0.09520
11		0.00080	0.00060	0.04340
12		0.00020	0.00040	0.01180
13				0.00140
14				
15				
16				

Table 12 Four level hierarchy: delay and performance metrics, *p* varied in the range [0.2,0.4], no communication

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
1.00	4.89	5.78	3.43	0.23	2.28	7.12	0.96
0.75	5.50	5.77	2.29	0.05	2.40	7.85	0.79
0.50	6.44	5.28	1.07	0.01	2.56	9.07	0.59
0.00	9.40				3.50	18.32	0.00

when no communication is permitted, the expected numbers (for the four delay probabilities) are 0.23020, 0.09160, 0.09800 and 0.71440 with heterogeneity versus 0.2246, 0.0984, 0.0964 and 0.7298 without. Notice that when the probability of delay

Table 13 Expected number of projects in a four level hierarchy, *p* varied in the range [0.2,0.4], all levels communicate

Number of proposals	Probability of delay in combining proposals			
	1	0.75	0.5	0
1	4.32040	3.84340	3.46840	2.50200
2	5.44300	4.75680	3.94780	1.58260
3		0.98820	1.46900	1.15100
4	3.58440	2.34580	1.77160	0.96340
5		0.58320	0.82000	0.82380
6		0.29940	0.44300	0.66200
7		0.08360	0.19300	0.51200
8	0.30620	0.08960	0.08280	0.35660
9		0.02200	0.02960	0.22840
10		0.00740	0.00700	0.14260
11		0.00080	0.00160	0.06200
12			0.00100	0.02020
13			0.00040	0.00560
14				0.00140
15				
16				

Table 14 Four level hierarchy: delay and performance metrics, *p* varied in the range [0.2,0.4], all levels communicate

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
1.00	4.32	5.44	3.58	0.31	2.38	7.85	1.01
0.75	4.91	5.63	2.42	0.07	2.50	8.54	0.84
0.50	5.91	5.20	1.12	0.01	2.67	9.82	0.62
0.00	9.01				3.66	20.04	0.00

equals one, heterogeneity is slightly helpful in producing richer projects. Similarly, when all managers can communicate the expected numbers are 0.30620, 0.11980, 0.12240 and 0.81680 versus 0.2998, 0.1202, 0.1348 and 0.8318. This occurs most likely because when there is delay, the possibility of getting a project with many ideas combined is small. However, due to heterogeneity, managers with high combinative capacity can set the process of generating a rich project into motion in the lower levels of the hierarchy. This combined proposal can then snowball through the hierarchy gathering more proposals to it.

We also explore extreme heterogeneity. We allow p to vary in the range [0,0.6]. For the no communication case, the average number of ideas combined goes down (see Table 15). The effect is more pronounced when there is significant delay. The other effect is in the expected number of project with eight or more ideas. These numbers are now 0.17900, 0.07240, 0.07460 and 0.51440, against 0.2246, 0.0984, 0.0964 and 0.7298 for the four delay probabilities. Thus, extreme heterogeneity affects the expected number of such projects considerably.

We now turn to variability in p_1. This can be equated with variability in communication. We vary p_1 uniformly in the interval [0, 0.02]. The results for the 'all communicate' case are shown in Table 16 (compare with Table 10, top panel). Variability in communication does *not* affect the average number of proposals combined or the variance. On the contrary, we observe a slight increase in three cases in the variance and/or the average number of proposals combined. In fact, the expected numbers of projects with eight or more ideas combined for the four delay probabilities are 0.30780, 0.12580, 0.12720 and 0.83900 versus 0.29900, 0.11980, 0.13440 and 0.82540. This is rather surprising. A possible reason is that we have allowed equal chances for increase or decrease in the value of p_1. Thus, the gains when p_1 increases more than offset the losses when it is decreased. However, it is still surprising because it is not the case that p_1 is increased uniformly for the whole hierarchy but it may increase for one and decrease for another manager in the hierarchy. Thus, even 'random' linkages are helpful so long as the probability of such linkages is positive.

Table 15 Four level hierarchy: delay and performance metrics, p varied in the range [0,0.6], no communication

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
1.00	4.25	5.41	3.62	0.31	2.18	6.51	0.90
0.75	4.92	5.57	2.44	0.07	2.40	7.85	0.79
0.50	5.83	5.22	1.15	0.01	2.56	9.07	0.59
0.00	8.91				3.50	18.32	0.00

Table 16 Four level hierarchy: delay and performance metrics, $p = 0.3$, p_1 varied in the range [0,0.02], all communicate

Prob. of delay	Expected number of projects with delay				Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3			
1.00	4.25	5.41	3.62	0.31	2.39	7.89	1.02
0.75	4.92	5.57	2.44	0.07	2.51	8.57	0.83
0.50	5.83	5.22	1.15	0.01	2.67	9.86	0.63
0.00	8.91				3.69	20.32	0.00

3.3 The effects of deliberate delay

In this section we explore the possibility that managers may deliberately introduce delays. To many observers, bureaucratic delays seem unnecessary. We show that in our model deliberate delay may actually be beneficial and that it should be introduced judicially.

We assume the managers introduce delays so that they wait for one period to see what other proposals are coming and process the combined set of proposals together. Modeling this phenomenon is not that simple. If the managers are clairvoyant they will introduce just the right amount of delays so that similar proposals come together for evaluation. Instead of assuming this we assume that managers proceed according the same rule in the organization, namely, that proposals received during periods 0 and 1 will be processed together in period 1; proposals received during periods 2 and 3 in period 3, etc. This could correspond to a quarterly or semi-annual or annual proposal review cycle.

In Table 17 we show the results of this modification for the all communicate case. The results should be compared with Table 10 (top panel). Except for the case when the delay probability equals zero, the average number of ideas combined has increased by 8–14%, the variance by 21–36% and also the delay by 150–204%. The trade-off is quite evident! Deliberate delay increases combinative capacity. When the original delay probability is zero, there is nothing to be gained by holding up proposals. Thus, is bureaucratic delay a rational organizational response to inefficiency? Our findings suggest that managers could incorporate delays deliberately to cope with the fact that there is randomness in processing proposals. Such deliberate delays serve as coordinating mechanism because they allow similar ideas to come together. There are other reasons why deliberate delays are introduced, for example to safeguard oneself as well as the organization from errors that might be discovered or a change in policy.

Table 17 Four level hierarchy: delay and performance metrics, $p = 0.3$, $p_1 = 0.01$, all communicate, delay introduced

Prob. of delay	Expected number of projects with delay								Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3	4	5	6	7			
1.00	0.00	4.00	2.04	3.08	1.76	1.04	0.10	0.30	2.72	10.63	2.54
0.75	0.00	4.76	1.88	3.31	1.18	0.50	0.02	0.07	2.81	11.09	2.24
0.50	0.00	5.81	1.52	3.16	0.52	0.17	0.00	0.01	2.93	12.11	1.92
0.00	0.00	8.90							3.69	20.31	1.00

In Table 18 we show the results for four situations. These should be compared with the performance measures shown in Table 3. First, as above, all levels use the delay to wait and accumulate proposals. The gains are substantial, with average number of ideas combined increasing by 8–12%, variance by 18–32% and delays increasing by 157–216% (see top panel of Table 18). We explore this phenomenon a bit more. In the second panel we show the performance when only the second level in the hierarchy can deliberately use delays. In the remaining panels we show the effect of deliberate delay only at the third and the fourth (or top) level. Interestingly, the maximum benefits seem to obtain due to deliberate delay at the middle or third level. Whereas the minimum delay through the hierarchy and almost as good a performance is obtained by introducing deliberate delays at the top level.

In summary, deliberate 'bureaucratic' delays seem to be beneficial in producing richer projects. They serve the purpose of waiting on a proposal to see if similar proposals are coming down the pipeline. They may increase the burden of simultaneous processing of more proposals. They also may introduce inordinate delays thus reducing the efficiency to process proposals. It may be necessary to introduce delays at every level of the hierarchy. Most benefit appears to obtain by introducing delay in the middle tiers of the hierarchy. The tradeoff between and delay and proposal complexity seems to be best when delay is introduced at the very top of the hierarchy. Finally, strangely deliberate delay seems to help the most when the probability of delay is high.

4. Discussion

Organizational change is a topic that captures the attention of organization theorists. Earlier ideas developed in the realm of the general systems theory suggested that organizations should be structured so as to match the degree of complexity in their environments, as depicted in Ashby's (1956) law of requisite variety. However,

Table 18 Four level hierarchy: delay and performance metrics, $p = 0.3$, no communication, delay introduced

Prob. of delay	Expected number of projects with delay								Average number of proposals combined	Variance of proposals combined	Average delay
	0	1	2	3	4	5	6	7			
1.00	0.00	4.43	2.35	3.14	1.42	0.82	0.08	0.30	2.56	9.37	2.48
0.75	0.00	5.14	2.12	3.34	1.18	0.51	0.01	0.07	2.66	10.03	2.20
0.50	0.00	6.13	1.73	3.21	0.53	0.17	0.00	0.01	2.79	10.99	1.90
0.00	0.00	9.30							3.54	18.69	1.00
1.00	0.00	4.43	5.49	2.78	0.08				2.56	9.37	1.91
0.75	0.00	5.14	5.46	1.69	0.01				2.66	10.03	1.74
0.50	0.00	6.13	4.94	0.70	0.00				2.79	10.99	1.55
1.00	0.00	4.85	5.78	1.67	0.01				2.68	10.15	1.78
0.75	0.00	5.44	5.32	0.83	0.00				2.84	11.33	1.62
0.50	0.00	6.49	4.20	0.27	0.00				3.00	12.70	1.44
1.00	0.00	4.93	6.01	1.53	0.23				2.56	8.62	1.78
0.75	0.00	6.21	4.72	0.92	0.05				2.74	9.96	1.57
0.50	0.00	7.58	3.18	0.39	0.00				2.93	11.77	1.36

changing a hierarchical organization into a flexible innovative organization is not an easy task. Coupling this with the idea that today's innovative organizations focus on knowledge in a much more dynamic environment makes the task highly difficult.

In this paper we tackled a specific part of the problem. Assuming that organizations need to combine ideas (Schumpeter, 1934; Kogut and Zander, 1992), but acknowledging the reluctance of managers to do so, we focused on the structural problems that can facilitate the flow of ideas (proposals) in a hierarchical organization while putting some timing constraints on the evaluation of these ideas/proposals coming up through this flow. Our analysis attempts to depict a realistic organizational setting. Managers who may be reluctant to combine proposals may also be inclined to do so if they see the commonalities/complementarities between ideas and are able to evaluate them at the same point in time. Thus, timing of evaluation becomes important. One can argue that timing costs can be minimized if organizational memory is functioning properly. That is, even if manager A sends his project proposals up with delay after manager B's proposals have already been reviewed by manager C, manager C can remember the features of manager B's proposals and attempt to combine them later (cf. Garud and Nayyar, 1994).

It is unclear, however, if organizational memory is capable of functioning in such a way as it was demonstrated to be fallible (Levitt and March, 1988). Furthermore, our argument goes beyond the memory issue. It is not only delay costs that matter, but delays also mean that the probability of manager C discovering potential complementarities in a scenario such as that given above is reduced. As many managers know, getting the attention of senior managers to seriously focus on project proposals is not easy. And, if those proposals are not available simultaneously for the senior manager's review, the probability is reduced even further. Our reasoning follows managerial attention as the scare resource, and not organizational memory. In this respect our modeling approach is similar to the one described in the garbage can model of Cohen *et al.* (1972), who argued that the simultaneous configuration of different streams eventually determine organizational choices. In the same vein, we argue that presenting proposals simultaneously to superiors is an important element in getting them to focus on the potential commonalities and complementarities between those project proposals. While our analysis shows that deliberate delay may help in producing projects that combine several proposals (that have commonalities or complementarities), these advantages at the exploration stage should be considered along with the potential increased costs of delay at the exploitation stage (March, 1991).

Other authors adopted a network perspective to examine the flow of ideas between organizations. Schulz (2001) examined horizontal and vertical flows of knowledge among subunits of a large multinational corporation. He adopted an organizational learning approach and focused on the different ways ideas are spreading in an organization. Our analysis is more narrowly defined and focuses on the supervisor who examines project proposals rather than ideas in general. The results of Schulz's (2001) study and the present analysis suggest that there are potential benefits in looking at flows of knowledge in organizations as determinants of organizational innovation. Note also that our analysis shows that heterogeneity in the ability to scan project proposals across supervisors and subunits leads to potential improvement in obtaining projects that combine several proposals. Cross-supervisor communication appears to hold a great potential as affecting the ability to combine proposals. This result should be examined in future empirical research.

We simulated the decision process by focusing on the three aspects: The potential for proposals (ideas) to be combined, the processing efficiency of combining proposals (ideas), and the processing delay in combining proposals (ideas). Our results suggests the following implications.

First, the consequences of delay are important in evaluating ideas. However, while delays are often difficult to reduce organizations may incur efficiency losses when they attempt to decrease them. For example, ideas that hold great potential might get overlooked and errors in judgement could increase when the evaluation process is speeded up. The loss could possibly increase with the depth of hierarchy as well as with the delays that already exist in the organization. Therefore, routines and procedures that can speed up the processing or facilitate speedy evaluation of ideas without loss in efficiency can be extremely valuable to organizations that have several layers of the hierarchy.

Second, the possibilities of obtaining rich projects that combine several proposals are rather low unless the organization focuses on creating such projects. The opportunities to create such projects may go unnoticed in an organization without special effort. For example, when p is small (say 0.1), the probability of obtaining one combined project (comprising of all ideas generated) in a organization with L levels is proportional to 10^{-L}!

Third, there might be an optimal delay of processing since there are trade-offs between reduced efficiency and larger or even richer proposals. When there is great scope for detecting commonalities or complementarities within a hierarchy, processing delays increase, even though ideally a hierarchy should house specialists that act as if they were a loosely coupled set of subsystems. Alternative routines might be necessary to counter this inefficiency.

Fourth, flexible communications across levels within an organization have very beneficial effects. Communication linkages help even if the probability of combining ideas from other managers is not that large. Thus relatively ‘weak’ coupling of ideas across managers might suffice, as long as the communication channels are available to all managers at a given level. This outcome is primarily due to allowing each manager to see the proposals that are available at every other manager’s group within a hierarchical level. The effect of such random linkages is probably even more pronounced when the span of control is greater. However, the cost is again increased delays. There is also the cost of managers being able to create and maintain these linkages without regard to the time they spend on pursuing their own specialties.

Fifth, there is limited evidence that it is not only timing but proximity of ideas that results in larger, richer and an additional number of combined project proposals. Thus, there may be some advantage in random linkages amongst managers at the lowest levels. Several proposals may cluster together if they are originally formed in close proximity to one another. More research and experimentation is required to study this effect and understand how routines can help ideas stay together within the same space and also within the same time.

Sixth, deeper hierarchies produce projects that combine several proposals. Looking at the framework we develop here one may be a bit uneasy about the role that we pose for hierarchy in fostering innovation. While it is assumed that flat organizational structures are the right way to foster innovation, hierarchy has a role in creating structures that allows the process of innovation to be longitudinal. We focusesd on the combination of proposals to better describe some of the aspects of the process of fostering innovation in large organizations. If an organization like IBM or 3M is interested in such endeavor, creating flat organizational structures may not work. On the contrary, hierarchical structure that moves proposals up in the organization in an efficient manner can maintain innovation over a long period of time. Hierarchical structures are good for several reasons (Jaques, 1990), and the natural preoccupation with flat structures may be good in small organizations but needs to be coupled with other structural arrangements especially in large organizations. Brown and Duguid (2001: 198) commented that in understanding knowledge in organizations, 'often too much attention is paid to the idea of community and too little to the idea of practice'. Along the same line we suggest that the timing of evaluation is an essential element in attempting to increase the probabilility of combining ideas and proposals in large organizations.

We should note of course that developing more efficient mechanisms that facilitate their combinative capabilities does not guarantee that these are going to lead to more successful projects. Obviously, the demand for the products of such projects carry a significant weight. As Christensen (1997) showed, the selection of projects without anticipating future demand may lead successful firms to demise. In addition, organizations that are aptly aware of this trap should allow considerations about future demand affect the way they deal with knowledge creation.

Finally, research on creativity in organizations identified several elements that stifle creativity and innovation, among them is the enforcement of rigid bureaucratic procedures. Thus, when IBM decided to develop the PC, they sent the developers group away from company headquarters and bureaucracy. Our results suggest that flexible communication structures may facilitate innovation, in a similar vein to Nonaka's (1994) description of the hypertext organization where different layers in the organization are loosely linked together, and Brown and Duguid's (1991) claim that communities of practice should be fluid rather than bounded. However, in focusing on innovations that explore new combinations of existing knowledge and incremental learning (Schumpeter, 1934; Kogut and Zander, 1992), the ultimate goal is that commonalities/complementarities among proposals should be detected early in the game. Innovative companies should not rely on people of the status of Ibuka to point at such commonalities/complementarities but need to develop procedures that will facilitate the flow of ideas and the possibilities for combining proposals for better processes of knowledge creation. Such processes should not only allow 'chaotic' or flexible communication structures, but should also take into account the importance of timing and the cost of delays.

Acknowledgements

We thank Mie Augier, Bill Barnett, Joel Baum, Michael Cohen, Martin Schultz and an anonymous reviewer for comments on an earlier draft of this paper.

Address for correspondence

Stern School of Business, New York University, 40 West 4th St. #7-06, New York, NY 10012–1126, USA. Email: zshapira@stern.nyu.edu.

References

Ashby, R. (1956), *An Introduction to Cybernetics*. Chapman & Hall: London.

Brown, J. S. and P. Duguid (1991), 'Organizational learning and communities of practice: Toward a unified view of working, learning and innovation,' *Organization Science*, **2**, 40–57.

Brown, J. S. and P. Duguid (2001), 'Knowledge and organization: a social-practice perspective,' *Organization Science*, **12**, 198–213.

Cohen, D. and L. Prusak (2001), *In Good Company: How Social Capital Makes Organizations Work*. Harvard Business School Press: Boston, MA.

Cohen, M., J. March and J. Olsen (1972), 'A garbage can model of organizational choice,' *Administrative Science Quarterly*, **17**, 1–25.

Cohen, W. and D. Levinthal (1990), 'Absorptive capacities: a new perspective on learning and innovation,' *Administrative Science Quarterly*, **35**, 128–152.

Christensen, C. M. (1997), *The Innovator's Dilemma*. Harvard Business School Press: Boston, MA.

Cyert, R. and J. G. March (1992), *A Behavioral Theory of the Firm*, 2nd edn. Blackwell: Oxford.

De Bono, E. (1986), *Lateral Thinking*. Viking: New York.

Garud, R. and P. Nayyar (1994), 'Transformative capacity: continual structuring by inter-temporal technology transfer,' *Strategic Management Journal*, **15**, 365–385.

Grant, R. (1996), 'Toward a knowledge-based theory of the firm,' *Strategic Management Journal*, **17**, 109–122.

Jaques, E. (1990), 'In praise of hierarchy,' *Harvard Business Review*, **68**, 127–133.

Jensen, M. and W. Meckling (1992), 'Specific and general knowledge and organizational structure,' in L. Werin and H. Wijkander (eds), *Contract Economics*. Blackwell: Oxford.

Kogut, B. and U. Zander (1992), 'Knowledge of the firm, combinative capabilities, and the replication of technology,' *Organization Science*, **3**, 383–397.

Lant, T. and Z. Shapira (eds) (2000), *Cognition in Organizations: Computation and Interpretation*. Lawrence Erlbaum: Mahwah, NJ.

Levitt, B. and J. March (1988), 'Organizational learning,' *Annual Review of Sociology*, **14**, 319–340.

March, J. G. (1991), 'Exploration and exploitation in organizational learning,' *Organization Science*, **2**, 71–86.

March, J. G. (1999), *The Pursuit of Intelligence in Organizations*. Blackwell: Oxford.

Nayak, P. and J. Ketteringham (1986), *Breakthroughs*. Rawson Associates: New York.

Nelson, R. and S. Winter (1982), *An Evolutionary Theory of Economic Change*. Belknap Press: Cambridge, MA.

Nonaka, I. (1994), 'A dynamic theory of organizational knowledge creation,' *Organization Science*, **5**, 14–37.

Nonaka, I. and H. Takeuchi (1995), *The Knowledge-creating Company*. Oxford University Press: New York.

Polanyi, M. (1966), *The Tacit Dimension*. Anchor Day Books: New York.

Schulz, M. (2001), 'The uncertain relevance of newness: organizational learning and knowledge flows,' *Academy of Management Journal*, **44**, 661–681.

Schumpeter, J. (1934), *The Theory of Economic Development*. Harvard University Press: Cambridge, MA.

Seshadri, S. and Z. Shapira (2001), 'Managerial allocation of time and effort,' *Management Science*, **47**, 647–662.

Sutton, R. and A. Hargadon (1996), 'Brainstorming groups in context: Effectiveness in a product design firm,' *Administrative Science Quarterly*, **41**, 685–718.

Weick, K. (1979), *The Social Psychology of Organizing*, 2nd edn. Addison-Wesley: Reading, MA.