07-039

Managing Know-How

Deishin Lee
Eric Van den Steen (MIT Sloan)

Managing Know-How

Deishin Lee and Eric Van den Steen*

December, 2006

Abstract

We use an economic model to study the optimal management of know-how, defined here as employee-generated information about the performance of specific solutions to problems that may or will recur in the future.

We derive three main results. First, information about successes is typically more useful than information about failures, since successful methods can be replicated while failures can only be avoided. This supports firms' focus on 'best practice'. Second, recording mediocre know-how can actually be counter-productive, since such mediocre know-how may inefficiently reduce employees' incentives to experiment. This is a strong-form competency trap. Third, the firms that gain most from a formal knowledge system are also the ones that should be most selective when encoding information (i.e., the ones that are most at risk from the competency trap); namely, large firms that repeatedly face problems about which there is little general knowledge and that have high turnover among their employees.

Beyond these main principles, we also show that it may be optimal to disseminate know-how on a plant-level but not on a firm-level, and that storing back-up solutions is most valuable at medium levels of environmental change.

Keywords: Knowledge Management, Know-how, Competency Trap, Best Practice

1 Introduction

Know-how is a key resource for business, and know-how management a potential lever for competitive advantage. Lew Platt, a former CEO of HP, famously said, "If HP knew what HP knows, we would be three times as profitable." Some companies take this to heart and developed formal systems to manage the know-how generated by their employees. McKinsey & Company and Booz-Allen & Hamilton, for example, both implemented online knowledge management systems to guide analyses and recommendations for future clients (Bartlett 1996, Christiansen and Baird 1998). Siemens implemented ShareNet to record technical and functional solutions to problems (MacCormack 2002). Ford Motor company's "Best Practices Replication Process", which requires improvements to be quantifiable before they can be codified in the system, is estimated to have resulted in savings of $850M over four years (Stewart 2000). Finally, many companies codify best practices in the form of standard operating procedures and ISO 9000 documentation.[1]

* Harvard Business School (dlee@hbs.edu) and MIT-Sloan School of Management (evds@mit.edu). We thank Amy Edmondson, Francesca Gino, Gary Pisano, Scott Stern, Tunay Tunca, and seminar participants at the 2006 Annual Conference of the Production and Operations Management Society and Harvard Business School for helpful comments and suggestions.

[1] While the focus of this paper is on knowledge systems within firms, the ideas may also be applicable to other knowledge repositories, such as Biological Resource Centers (Stern 2004).

This raises a number of interesting and important questions. Should all companies engage in formal knowledge management? If not, which companies derive most value from a formal knowledge system? Conditional on implementing such a system, should the company focus more on learning from successes or on learning from failures? Should such knowledge systems simply capture all experience, or should it be more selective?

The purpose of this paper is to develop and apply an economic framework to address some of these questions. Following the above discussion, we will focus in particular on two first-order questions (from an economic perspective):

- Which firms derive the most value from a knowledge system?
- What type of information (e.g., success versus failure) should such a system collect?

To analyze these questions, we study a model in which a firm's employees are repeatedly faced with a particular problem. Through trial and error, the employees learn about the performance of alternative approaches, i.e., they develop know-how as to what works and what doesn't. The focal issue is whether and what types of knowledge the firm should record for future use. Formally, this setting is captured by looking at multiple generations of employees facing a multi-armed bandit problem[2] and studying what knowledge generated by one generation should be made available to future generations. The multiple generations of employees can be interpreted literally as representing turnover or – more broadly – as successive selections of employees (from the firm's total workforce) who are faced with a given problem. This setting corresponds, for example, to the situation of a management consultancy, a technical support service, or a product development group.

We first show that information about successful practices is typically more useful than information about failures. The reason is that information about a success tells you exactly what to do, while information about a failure merely excludes one of many possible courses of action. This may explain why firms tend to be more focused on 'best practice' than on learning from failure. At the same time, the analysis also generates some insights as to what factors may make information about failures more valuable.

Second, we show that codifying moderately successful practices may actually be harmful to an organization. The reason is a strong-form competency trap: once a successful practice is codified, employees have no incentive to experiment further with actions that could lead to even higher performance than the current best practice. The firm thus faces a trade-off between exploiting a known best practice and continuing experimentation in order to get even better performance in the future.

Third, we show that the factors that make a firm value a knowledge system also make it optimal for that firm to be more selective in encoding information (to avoid the competency trap); namely, firms that are larger, have more employees facing the same issue, face the issue more frequently, and face issues with more uncertainty about the performance of alternative solutions. The effect of firm size derives from two scale effects: on the knowledge *generation* side, large firms have more employees experimenting, leading to better solutions; on the knowledge *application* side, large firms can apply the same knowledge more broadly. Increased uncertainty about the performance of alternative solutions increases the option value of experimenting, and thus the opportunity cost of forgoing experimentation.

[2] A multi-armed bandit problem is a setting where a player has to choose among a number of alternatives that differ in their payoffs, where there is uncertainty about each alternative's payoff, and where the player learns about these payoffs by sequentially trying out different alternatives. For general treatments of multi-armed bandit problems, see Berry and Fristedt (1985) or Gittins (1989). The term derives from the problem a gambler faces when playing a multi-armed bandit in a casino.

We also show that – as a consequence of this tension between experimentation and exploitation – it may be optimal to disseminate (mediocre) know-how at the plant level but not disseminate it at the firm level, so that part of the organization exploits the current best practice, while the rest continues to experiment. We further consider how the results may be affected by the fact that the effectiveness of a particular practice may change over time and show that backup solutions (i.e., alternative solutions to the same problem) are most useful when the rate of change is intermediate: in a stable environment, it suffices to remember one best practice to continue to be successful in the future; with very rapid change, information from the past is unlikely to help in the future. Finally, a knowledge system is – in expectation – also most useful for problems of intermediate-level difficulty: for very simple problems a formal knowledge system is overkill, while extremely difficult problems only rarely generate know-how, thus making the knowledge system (in expectation) superfluous.[3]

Literature The research on knowledge management in the management literature focuses primarily on empirical studies of how firms can create, retain, and transfer knowledge (see surveys in Argote 1999, Argote et al. 2003, Holsapple 2003, and Levitt and March 1988) and how that is affected by different factors. For example, Cohen and Levinthal (1990) suggest that an organization's ability to leverage new information depends on its "absorptive capacity," which is a function of its prior knowledge in a related area. Brockman and Morgan (2003) study how existing knowledge affects new product performance and innovativeness, while Sorenson (2003) studies how environmental volatility affects knowledge transfer in vertically integrated and non-vertically integrated firms. Approaching the resource-based view of the firm from an operations management perspective, Carrillo and Gaimon (2004) study – with a constrained optimization model – the effect of uncertainty on the optimal level of investment in knowledge-based resource capabilities, as embedded in process change and employee knowledge. On the issue of learning from success vs. failure, the popular business press tends to emphasize lessons from best practice (Bartlett et al. 2003, O'Dell and Grayson 1998), although some side with a body of academics who advocate learning from failure (Sitkin 1992, Leonard-Barton 1995, Miner et al. 1999, Canon and Edmondson 2001). Our findings suggest that the *informational value* of a success is greater than that of a failure, although failures can also provide information and lead to adaptive skills that benefit the organization. Relative to this literature, our contribution is to use a formal economic (game-theoretic) model and thus take an explicit economic perspective on the question of know-how management. Our analysis is also more focused on the informational value of specific knowledge and how it interacts with the incentives for experimentation versus exploitation of existing knowledge.

While this is – to our knowledge – the first paper to take a formal economic perspective on knowledge systems, there are a number of related areas in the economic literature. Conceptually closest is the literature on the value of information (Blackwell 1951, Marschak and Radner 1972, Hilton 1981, Athey and Levin 2001), which studies the value of information to an individual decision maker, and how that value is influenced by the characteristics of the decision maker, the characteristics of the information, and the structure of the problem. It focuses, however, on the static problem of one decision maker. Athey and Levin (2001), for example, study the demand for information by decision makers faced with "monotone decision problems," i.e., decision problems in which actions and signals can be ranked such that higher actions are chosen in response to higher signals.

Another body of literature studies how a firm's organization depends on the need to communi-

[3] Note that this statement is "in expectation." Of course, from an ex-post perspective, the value of information (once you have it) always increases in the difficulty of the problem.

cate and use information or knowledge (e.g., Marschak and Radner 1972, Radner 1993, Van Zandt 1999, Bolton and Dewatripont 1994, Garicano 2000, Alonso et al. 2006). The focus of this literature is not on the knowledge itself, but on how it affects organization design. Van Zandt (1999), for example, studies decentralized computation as a model for information processing by organizations.

Finally, a paper that nicely complements ours is Manso (2006), who studies how to give agents incentives to innovate. Among other things, he shows that a tolerance, or even reward, for failure and timely feedback are key ingredients of incentive schemes that are conducive to innovation.

Contribution The contribution of this paper is to study the optimal knowledge management strategy with a formal economic model. We show that information about successes is more valuable than information about failures, and that recording information about mediocre practices may actually be counter-productive. We also analyze which firms derive most value from a knowledge system and which firms should be most selective in encoding information. It is, to our knowledge, the first paper to study knowledge systems with an economic (game-theoretic) framework.

The rest of the paper is structured as follows. Section 2 presents our model. We examine best practices and the strong-form competency trap in Section 3. Sections 4 and 5 consider respectively the issue of local versus global dissemination and how the results may be affected by the fact that the effectiveness of a particular practice may change over time. We discuss extensions to our model in Section 6 and conclude in Section 7.

2 The Basic Model

The model in this paper captures the situation of a firm whose employees repeatedly face similar issues. This may be a white-goods manufacturer that periodically designs new refrigerators, or a consulting firm that often works on post merger management. The firm learns about the performance of alternative approaches to each issue through its employees' trials and errors. The firm has to decide whether to set up a knowledge system to capture such knowledge and, if so, how to manage that system. For simplicity and transparency, we focus here on a finite horizon game in which a firm employs successive generations of employees who each face the same problem. This can also be interpreted as a setting in which the firm keeps the same employees over time, but in each period a different selection of the firm's employees are faced with that particular problem.

Formally, consider a firm that exists for three periods and employs in each period a new generation of I employees who all face the same problem. Each of these employees must choose a particular approach to the problem – an action – from an infinite set $A = \{a_1, \ldots, a_j, \ldots\}$. This action set is common to all employees of all generations. The actions differ in their performance, i.e., in how well they solve the problem or in what payoff they generate: each action a_j results in a fixed but unknown performance v_{a_j}. To capture this variation, the v_{a_j} will be independently and identically distributed draws from a non-degenerate distribution G with mean μ_G. While the realizations of v_{a_j} are (originally) unknown, the distribution G is common knowledge. When an employee undertakes an action, the firm observes the performance of that action so that it becomes (temporarily) part of its know-how. The firm can store this know-how (a_j, v_{a_j}) in its knowledge system to make it available to the next generation of employees.

The sequence of events in the game is described in Figure 1. At the very start of the game, the firm decides whether to build a formal knowledge system and, if so, what its capacity $n \geq 0$ should be. This capacity n is the maximum number of action-performance sets (a_j, v_{a_j}) it can store. The

0	1	2	3
Setup	Period 1	Period 2	Period 3
a Firm chooses $n \geq 0$.	a New generation of employees starts.	a New generation of employees starts.	a New generation of employees starts.
b Performance $v_{a_j} \overset{\text{iid}}{\sim} G$ are drawn.	b Employees simultaneously choose actions a_j.	b Employees simultaneously choose actions a_j.	b Employees simultaneously choose actions a_j.
	c Firm decides which (a_j, v_{a_j}) to store.	c Firm decides which (a_j, v_{a_j}) to store.	

Figure 1: Time line of model

cost of a knowledge system of size n is cn, with $c > 0$. Subsequently, the payoffs $v_{a_j} \overset{\text{iid}}{\sim} G$ are drawn. As mentioned before, the realizations for the v_{a_j} are not publicly revealed at this point.

After this 'setup' period, there are three regular periods which are very similar to each other. In each of these periods, first a new generation of I employees starts. These employees observe any information that is currently in the knowledge system. Second, these I employees simultaneously choose their respective actions from the set A. The firm observes these actions and their resulting payoffs. Finally, the firm decides what know-how (a_j, v_{a_j}), if any, to record in its knowledge system. The only information that can be recorded at the end of period t is whatever know-how was transferred from period $t-1$ plus whatever was revealed through new actions undertaken in period t. The firm has discount factor δ.

The firm maximizes overall performance, i.e., the net present value of all payoffs for all employees. Each employee maximizes his own performance, i.e., the expected payoff from his own action.[4] When indifferent, employees act in the firm's interest. When both are indifferent, employees choose an action at random. Finally, when indifferent, the firm prefers to not store any information. These assumptions on what employees or firms do when indifferent are made for convenience and do not drive the results.

Two variations on the basic model The model as described above is analyzed in the next section. Sections 4 and 5 will consider slight variations:

1. While in the basic model, all employees have automatic access to the knowledge system, Section 4 allows the firm to decide (per period) which employees get access to what information.

2. While in the basic model, v_{a_j} (i.e., the performance of action a_j) remains unchanged throughout the game, Section 5 considers the effect of allowing v_{a_j} to change over time.

In both cases, the respective section will discuss the changes to the model in more detail.

[4]This assumption captures the typical situation at most firms. The results would probably go through if employees cared about a combination of their own and the firm's performance. The reason for using this simpler assumption is to keep the analysis transparent and tractable. The assumption can be endogenized, though at a cost of a considerable increase in complexity. One approach is as follows. Each player has some randomly drawn, player-specific productivity θ_i, which is known by the firm but not by outsiders. The payoff from player i's action is $v_{a_j} + \theta_i$. That performance is observable by outsiders but not verifiable. After the game, each employee will get paid (in the labor market) according to his expected productivity. Actions and payoffs are not contractible, and the firm has no way to dictate what an employee should to do. The only possible contract is then a fixed wage. In these circumstances, the employee will want to choose the action that maximizes his own performance to maximize his post-game wage. An alternative approach would be to introduce moral hazard for each employee, combined with risk aversion. This will typically lead to a contract that puts different weights on current and future payoffs, so that the qualitative results would be preserved.

3 Best Practice and Strong Competency Trap

In this section, we derive three key results:

1. The firm records (at most) the 'best practice' among its potential know-how. Implicit in this result is the effect that information about successes is generally more useful than information about failures.

2. Firms may strictly prefer *not* to record information about mediocre practices even if recording information were free. This is caused by a strong-form competency trap.

3. The drivers that make a firm value a knowledge system also make it optimal for the firm to be more selective in what information to record, i.e., it makes the firm more sensitive to the competency trap.

We will treat these results in sequence, each time discussing the underlying intuition, potential caveats, and the formal proof. Consider first the result that the firm only records the 'best practice' in its knowledge system. It is useful to unpack this in two parts: the firm records only successes and – conditional on recording only successes – it will record only one, the best practice. These two results are driven by two effects. The first and most important effect is that information about successes (i.e., above-average performance) is more useful than information about failures (i.e., below-average performance): if an employee learns about an extremely good action, he can take that action and receive the payoff, but if he learns about an extremely bad action, he still has to choose among the remaining actions, giving him something close to the average payoff. For the stylized model of this paper, this takes the stark form that failures will never be recorded. Section 6 explores under what conditions this result gets relaxed and failures may be recorded.

The second effect that drives this result is that among the actions in the knowledge system, it is in the employees' best interest to undertake the action that leads to the highest performance (conditional on that performance being greater than μ_G). Once employees stop experimenting with new actions, it is also in the firm's best interest that the employees undertake the best known action. Since failures never get recorded, it follows that only *one* action is worth remembering from the firm's perspective, and that is the best practice.

To derive this result formally, let $B_t \subset A$ denote the set of actions about which information is available in the knowledge system at the start of period t, let $C_t \subset A$ denote the set of actions selected by employees in period t, and let $a_t^* = \text{argmax}_{a_j \in C_t \cup B_t} v_{a_j}$ be the highest-payoff action in either B_t or C_t. The following lemma describes the optimal capacity for a knowledge system and what know-how the firm should collect. We use superscript $*$ to denote optimized values.

Lemma 1 *The optimal capacity of a knowledge system is $n^* \leq 1$. In each period t, at most the action with the highest payoff, $a_t^* \in B_t \cup C_t$, is recorded.*

Proof : Consider first the employees' action choices in periods $t = 2, 3$. Since $\#A = \infty$, the employees' prior on any action on which no information is stored is always G. Assume now that information from previous period(s) on a subset $B_t \subset A$ has been recorded and communicated to employees. Let $\tilde{B}_t = \{a_j \in B_t : v_{a_j} \neq \mu_G\}$. Let $\hat{a}_t$ be the highest-payoff action in B_t, i.e., $v_{\hat{a}_t} \geq v_{a_j}$ for all $a_j \in B_t$. The optimal employee strategy is as follows: if $v_{\hat{a}_t} > \mu_G$ choose $\hat{a}_t$, else choose some action from $A \setminus \tilde{B}_t$. Given this employee behavior, it follows that (i) the firm never records an action with $v_{a_j} \leq \mu_G$: doing so makes employees avoid the action, but the benefit of doing so is zero since $\#A = \infty$, and (ii) the firm never records a_k if it also knows about a_j, where $v_{a_j} > v_{a_k}$. It follows that the firm records at most one action, so that

$n^* \leq 1$. Furthermore, if it records an action, it is the action with the highest payoff in either B_t or C_t. This proves the lemma. ■

Section 5 will show that more than one action will sometimes be recorded when the actions' performance may change over time. In particular, at intermediate levels of change, it will be worthwhile to record backup solutions in case the best practice becomes less effective. Section 6 explores how factors such as the structure of the action set or the ex-ante distribution of payoffs may make it optimal to record failures or to record more than one action's performance. It will turn out that no single modification to the model suffices to make either of these results disappear, so that they are relatively robust. This result, and its relative robustness, may explain why firms and management tools such as TQM or ISO9000 certification, focus so much on 'best practice'.

The second key result – that it may be optimal for a firm to explicitly not record information about successes (i.e., actions with above-average performance) even if doing so were free – is more surprising. It is caused by an agency problem: since the benefits from experimentation are realized in the next period (or elsewhere in the firm), employees don't benefit from these gains while the firm does. The optimal level of experimentation is thus lower from the perspective of an employee than from the perspective of the firm. This takes a stark form in the stylized model of this paper: once any above-average action is known, employees will choose that action and stop experimenting. To force employees to continue to experiment, the firm will therefore avoid recording information about actions with mediocre above-average performance. In fact, the firm would even be willing to pay to ensure that information on actions with mediocre payoffs (i.e., with $\mu_G < v < \tilde{v}_1$ in terms of the proposition below) disappears.

This issue is essentially a strong-form "competency trap." The notion of "competency trap," introduced by Levitt and March (1988), captures the idea that an (at first sight) valuable competency may have unexpected consequences that end up hurting the firm. An essential component of the notion of a competency trap as suggested by Levitt and March (1988) is that the firm must become proficient using the inferior procedure. In particular, they describe it as a situation "when favorable performance with an inferior procedure leads an organization to accumulate more experience with it, thus keeping experience with a superior procedure inadequate to make it rewarding to use." The competency trap we identify here is actually stronger: the simple fact of knowing about the moderately successful practice inefficiently reduces the incentives for employees to experiment.

Sitkin (1992) also suggests that success can be a liability to an organization because it can breed complacency and homogeneity, and also because organizations tend to punish failure disproportionately more than inaction (or status quo). We find that the liabilities of success can be even more extreme: risk neutral, performance-maximizing employees who are rewarded symmetrically for failures and successes will stop experimenting once information about a moderately successful practice is revealed to them. The alternative approach of giving employees explicit incentives to experiment is discussed below.

To show this result formally, we will now derive the firm's optimal strategy when it has a knowledge system with one unit of capacity ($n = 1$). The following proposition indeed says that the firm will sometimes not record information that would be valuable to employees.

Proposition 1a *There exists $\tilde{v}_1 > \tilde{v}_2 = \mu_G$, such that for $t = 1, 2$,*

- *actions in period t with payoff $v \leq \tilde{v}_t$ never get recorded,*
- *the first time any payoff in period t strictly exceeds $\tilde{v}_t$, the action corresponding to the highest payoff in that period gets recorded, and*

• *from that time on, all employees undertake that one recorded action.*

Proof : We prove this by backwards induction. From Lemma 1, $\#B_t \leq n^* \leq 1$. Let $\hat{a}_t$ denote the recorded action, if any, available to period t employees. Let $\tilde{B}_t = \{a_j \in B_t : v_{a_j} \neq \mu_G\}$.

No information is recorded in period 3 since it is the last period. Employees in period 3 undertake $\hat{a}_3$ if $v_{\hat{a}_3} > \mu_G$ and otherwise randomly select from $A \setminus \tilde{B}_3$. It follows that the continuation value equals μ_G if $\tilde{B}_3 = \emptyset$ or $v_{\hat{a}_3} \leq \mu_G$, and equals $v_{\hat{a}_3}$ otherwise. It also follows that at the end of period 2, the firm will store the highest-payoff known action, as long as that action has payoff $v > \mu_G$. In other words, if we set $\tilde{v}_2$ such that $\tilde{v}_2 = \mu_G$, then it will indeed be the case that the action with the highest known payoff with $v > \tilde{v}_2$ is recorded, and that that action gets executed in period 3. If no such action is known, then $B_3 = \emptyset$ and the players select randomly from A.

In period 2, if $\tilde{B}_2 = \emptyset$ or $v_{\hat{a}_2} \leq \mu_G$ then period-2 employees will select actions from $A \backslash \tilde{B}_2$ (and coordinate on choosing different actions since, them being indifferent, they do as the firm prefers), and the continuation of the game is as derived above. However, if $v_{\hat{a}_2} > \mu_G$ then all period-2 employees will choose $\hat{a}_2$. At the end of the period, the firm only knows the payoff $v_{\hat{a}_2} > \mu_G$, and thus continues to store that information, so that $\hat{a}_3 = \hat{a}_2 > \mu_G$ and all employees take that action again in period 3. It follows that once an action with payoff $v > \mu_G$ gets recorded, all employees will undertake that action, no new information gets generated, that same action will be kept in the knowledge system henceforth, and that same action will be undertaken henceforth.

To complete the proof, we now determine what actions will be recorded at the end of period 1. From the above, it follows that the continuation payoff at the start of period 2 when $\tilde{B}_2 = \emptyset$ (or $v_{\hat{a}_2} \leq \mu_G$) equals (given that $\tilde{v}_2 = \mu_G$)

$$V_2 = \mu_G + \delta \int_{\mu_G}^{\infty} u \, dG_{(I)}(u) + \delta \mu_G P[v_{(I)} \leq \mu_G] = (1+\delta)\mu_G + \delta \int_{\mu_G}^{\infty} (u - \mu_G) \, dG_{(I)}(u),$$

where $G_{(i)}$ and $g_{(i)}$ are respectively the distribution and density functions of the i^{th} order statistic, $v_{(i)}$, of the random variable drawn from distribution G. The continuation payoff when B_2 contains an action with payoff $v > \mu_G$ equals $(1+\delta)v$. It follows that the firm will store a value $v > \mu_G$ if and only if

$$v > \tilde{v}_1 = \mu_G + \frac{\delta}{1+\delta} \int_{\mu_G}^{\infty} (u - \mu_G) \, dG_{(I)}(u). \tag{1}$$

Since the second term on the righthandside is positive, $\tilde{v}_1 > \mu_G$. This completes the proof. ■

This result captures a very important issue for firms: employees typically bear the costs of local experimentation while the firm as a whole gets the benefits. There seems to be an obvious solution to this problem: give employees a share of the overall benefits from experimentation.[5] Such incentive schemes, however, run into all kinds of problems. First, the benefits of new methods are notoriously difficult to quantify (which is necessary to base incentives on it): not only is the counterfactual typically missing, but nearly any change interacts with many others. Second, due to a difference in timing between the (immediate) costs of experimenting and its (future) benefits, the firm may be tempted to renege on such promises. Third, the extra cost induced by such pay-for-performance reduces the incentives for the firm to implement the idea, which reduces the expected benefits for an employee and thus the incentives to experiment. And then there are a host of smaller issues such as concerns for fairness, the fact that ideas are difficult to protect, or the need for team effort in trying new methods. To capture such issues generically, we simply assumed that employees only care about their own performance, but the qualitative results seem to extend to the case where such incentive schemes are partially feasible. This is an issue for further research.

We now turn to the question of which firms should be most selective when recording information (to avoid the competency trap). That – and how it relates to a firm's value from a knowledge system

[5] For a more sophisticated view on incentives for innovation, see Manso (2006).

– was the subject of the third key result. Proposition 1b below will show that a firm should be more selective when it is larger, has more employees facing the same issue, faces the issue more frequently, and faces issues with more underlying uncertainty.

To see the intuition behind these different results, it is useful to remember that when codifying a best practice, the firm trades off the short-term benefits of using the current know-how with the opportunity cost of forgoing experimentation. The comparative statics all follow from how the opportunity cost of forgoing experimentation changes with each factor. First, the opportunity cost of forgoing experimentation increases with the size of the firm I (or with the number of employees facing that issue) since with more employees experimenting, the expected payoff of the best practice – which is a first order statistic – increases. Second, more uncertainty about the performance of the actions increases the option value of experimenting, and thus the opportunity cost of forgoing experimentation. Finally, the opportunity cost of forgoing experimentation also increases in the discount factor, δ, since a higher δ makes the future – and thus the benefits from further experimentation – more valuable. This latter comparative static on δ has an alternative useful interpretation. In particular, if we look at periods as time intervals of length T in a continuous-time model with discount rate r, then $\delta = e^{-rT}$ is also a measure for the time between intervals. A high discount factor thus corresponds to facing that problem more frequently (though every time by different employees). It follows that firms that are faced more often with the problem or that have high turnover should be more selective in encoding best practice.

These comparative statics are derived in the following proposition. To capture the notion that one probability distribution differs from another in the ex-ante uncertainty about the payoffs, we will use a *mean-preserving stretch* to represent that one distribution is a stretched-out version of another. In particular, F is a *mean-preserving stretch* of G if $\mu_F = \mu_G$ and there exists a continuous function $\phi(x)$ such that $F(\mu_F + \phi(x)) = G(\mu_G + x)$ with $\phi(0) = 0$ and $1 < \phi'(x) < A$ for some $A < \infty$ and $x \in \mathbb{R}$.[6] A special case is that F is a *linear* mean-preserving stretch of G if $\mu_F = \mu_G$ and $F(\mu_G + \beta x) = G(\mu_G + x)$ for some $\beta > 1$ and $x \in \mathbb{R}$.

Proposition 1b *The optimal $\tilde{v}_t$ increases in I, δ, and a mean-preserving stretch of G (strictly for $t = 1$ and weakly for $t = 2$).*

Proof : Since $\int_{\mu_G}^{\infty} (u - \mu_G)\, dG_{(I)}(u)$ is strictly positive and $\frac{d}{d\delta}\left(\frac{\delta}{1+\delta}\right) > 0$, equation (1) in the proof of Proposition 1a implies that $\tilde{v}_1$ increases in δ. Second, an increase in I causes a strict first-order stochastic dominance shift of $G_{(I)}$. Since the integrand is strictly increasing, it follows that $\int_{\mu_G}^{\infty} (u - \mu_G)\, dG_{(I)}(u)$, and hence $\tilde{v}_1$, strictly increases in I. Finally, let F be a (strict) mean-preserving stretch of G. Note that this implies that $F(x) < G(x)$ for $x > \mu_G$. The function

$$H(u) = \begin{cases} G(u) & \text{when } u \leq \mu_G = \mu_F \\ F(u) & \text{when } u > \mu_G = \mu_F \end{cases}$$

is also a distribution function, and one which first-order stochastically dominates G. Moreover,

$$\int_{\mu_F}^{\infty} (u - \mu_F)\, dF_{(I)}(u) = \int_{\mu_F}^{\infty} (u - \mu_F)\, dH_{(I)}(u) = \int_{\mu_G}^{\infty} (u - \mu_G)\, dH_{(I)}(u) > \int_{\mu_G}^{\infty} (u - \mu_G)\, dG_{(I)}(u),$$

since $H_{(I)}$ is a first-order stochastic dominance shift of $G_{(I)}$ (and strictly so on the relevant support) and the integrand is strictly increasing. This completes the proof. ■

[6]Note that mean-preserving stretch is a stricter notion than mean-preserving spread. We want to capture the fact that the two distributions differ in their variance and in their variance only. Mean-preserving spread allows the distributions to differ completely, as long as they can be ordered by their riskiness.

The third key result is that the factors that make a firm more selective also make it value a knowledge system more; namely, firms that are larger, have more employees facing the same issue, face the issue more frequently, and face issues with more uncertainty about the performance of alternative solutions. That is the content of the following proposition.

Proposition 1c *If* $\delta P[v_{(I)} \leq \tilde{v}_1] \int_{\mu_G}^{\infty} (v-\mu_G) dG_{(I)}(v) + (1+\delta) \int_{\tilde{v}_1}^{\infty} (u-\mu_G)\, dG_{(I)}(u) \geq \frac{c}{I}$ *then* $n^* = 1$ *and the firm remembers the highest-payoff action, else* $n^* = 0$*. The value of a knowledge system increases in* I*,* δ*, and a mean-preserving stretch of* G*, and decreases in* c*.*

Proof : Let V_t be the firm's expected payoff per employee when it starts period t without any information (but with a knowledge system with capacity $n = 1$). From the proof of Proposition 1a, $V_2 = (1+\delta)\mu_G + \delta \int_{\mu_G}^{\infty} (v - \mu_G) dG_{(I)}(v)$. The value of a knowledge system with $n = 1$ equals

$$\begin{aligned}
&\left(V_2 P[v_{(I)} \leq \tilde{v}_1] + (1+\delta) \int_{\tilde{v}_1}^{\infty} u\, dG_{(I)}(u) - (1+\delta)\mu_G \right) I - c \\
&= \left(\delta P[v_{(I)} \leq \tilde{v}_1] \int_{\mu_G}^{\infty} (v - \mu_G) dG_{(I)}(v) + (1+\delta) \int_{\tilde{v}_1}^{\infty} (u - \mu_G)\, dG_{(I)}(u) \right) I - c \qquad (2)
\end{aligned}$$

By the envelope theorem, we can disregard any changes in $\tilde{v}_1$ when considering comparative statics. The effect of c and δ follow from visual inspection of equation (2), taking into account that $\tilde{v}_1 > \mu_G$. Consider next the effect of I. Since the term in brackets is positive, it suffices to show that

$$\frac{\delta}{1+\delta} P[v_{(I)} \leq \tilde{v}_1] \int_{\mu_G}^{\infty} (v - \mu_G) dG_{(I)}(v) + \int_{\tilde{v}_1}^{\infty} (u - \mu_G)\, dG_{(I)}(u)$$

increases in I. The effect of I goes completely through $G_{(I)}$. To prove now the result, we will consider the effect of I on $G_{(I)}$ separately for the first period and for the second period. In particular, using M for the first-period G and N for the second-period G, we can write the above expression as

$$\begin{aligned}
&\frac{\delta}{1+\delta} P_N[v_{(I)} \leq \tilde{v}_1] \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) + \int_{\tilde{v}_1}^{\infty} (u - \mu_N)\, dN_{(I)}(u) \qquad (3) \\
&= \frac{\delta}{1+\delta} \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) + \int_{\tilde{v}_1}^{\infty} \left(u - \mu_N - \frac{\delta}{1+\delta} \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) \right) dN_{(I)}(u) \qquad (4)
\end{aligned}$$

Equation (3) shows that the whole expression increases upon a FOSD shift in $M_{(I)}(v)$, induced by an increase in I. Equation (4) shows the same for $N_{(I)}(u)$. Finally, using an argument completely analogous to that in Proposition 1b, it is clear from (3) that a mean-preserving stretch of M increases the whole expression, and again analogously for $N_{(I)}(u)$ and (4), using then the fact that $\tilde{v}_1 > \mu_G$. This proves the proposition. ■

The comparative static on c is new but rather trivial: as the cost of implementing a knowledge system decreases, its net value to the firm obviously increases. The intuition for the other factors is for the most part similar to the intuition in the case of selectivity, although there are some important differences. The most important difference is that, with respect to the value of a knowledge system, the firm size I now matters for *two* reasons:

1. Experimentation: In larger organizations, more employees experiment, so that the quality of the know-how increases, making it more valuable to remember.

2. Exploitation: Larger organizations benefit from returns to scale since they have more employees who can apply the best practice, at the same cost.

In other words, knowledge systems generate scale effects in both the experimentation (knowledge generation) and exploitation (knowledge application) phases. Casual observation indeed suggests that it is the large Fortune 500 companies, such as Ford and Siemens, and the large consulting companies, such as McKinsey and Booz-Allen & Hamilton, that have implemented knowledge management systems.

The intuition for the effect of (ex-ante) uncertainty about the performance of alternative solutions and for the discount factor are the same as before. First, uncertainty increases the option value of being able to remember the best practice, and thus the value of the knowledge system itself. Second, a higher discount rate makes the future – and thus the benefits of experimentation – more important. The interpretation of δ in terms of frequency is again very useful. In particular, this comparative static gives one explanation why, for example, McDonald's with its high employee turnover codifies processes as standard operating procedures (SOP's) and then measures quality as adherence to these SOP's (Margolis and Upton 1996), while a family-run restaurant with long-term employees is less likely to codify processes so rigidly.

It should actually be no surprise that the value of a knowledge system and the optimal selectivity depend on the same factors: the value of a knowledge system is driven by the benefits of (remembering the results of) experimentation while optimal selectivity is driven by the opportunity cost of forgone experimentation. These are, to a large extent, two sides of the same coin.[7] The implication of this higher-level insight is quite important: it means that the firms that are most likely to implement a knowledge system should also try to be most selective in recording best practice.

4 Local versus Global Dissemination

The previous section assumed that the firm can control what information is recorded, but not who can access it. In other words, any recorded information is automatically available to all employees. This is the case for many formal knowledge systems, such as those of McKinsey or of Booz-Allen & Hamilton (Christiansen and Baird 1998). In some settings, however, the firm may be able to control which employees get what information. A typical case is one where the firm decides whether certain know-how should be disseminated on a plant level versus on a firm level. This section studies this alternative by assuming that the firm can choose in each period which employees observe the stored knowledge.

Since the firm captures benefits from experimentation that employees do not, the firm faces a trade-off between experimentation and exploitation. All-or-nothing access presented the firm with a stark choice between experimentation and exploitation since the only way to force an employee to experiment was to not record information. When the firm can explicitly control who can access the know-how, however, the firm can have both experimentation and exploitation. In particular, the following proposition shows that it is often optimal for the firm to communicate mediocre information to a subset of employees, while forcing the other employees to experiment further. Let $v_{\hat{a}}$ denote the best known payoff at the end of period 1, corresponding to action $\hat{a}$. Let $\underline{v} = \inf\{\operatorname{supp} G\}$ and $\overline{v} = \sup\{\operatorname{supp} G\}$, with $\underline{v} < \overline{v}$, since G is non-degenerate.

Proposition 2 *There exist $\hat{v}_1, \ldots, \hat{v}_I$, where $\mu_G < \hat{v}_I < \cdots < \hat{v}_1 < \overline{v}$ such that*

- *if $v_{\hat{a}} \leq \hat{v}_I$, then no information is recorded and all employees experiment in period 2,*

[7] The one key difference (except obviously for the role of the cost c) is that the firm size I has a double effect for the value but not for selectivity. The reason is that with selectivity, the cost of experimenting is also proportional to the firm size (since it means forgoing the benefit of best practice for all employees).

- *if $\hat{v}_{j+1} < v_{\hat{a}} \leq \hat{v}_j$, $j = 1, \ldots, I-1$, then $v_{\hat{a}}$ is communicated to $I-j$ employees who undertake action $\hat{a}$ in period 2, while the other j employees experiment,*
- *if $v_{\hat{a}} > \hat{v}_1$, then $v_{\hat{a}}$ is communicated to all employees who undertake $\hat{a}$ in periods 2 and 3.*

Proof : Let $V(x, j)$ denote the continuation value if, at the start of period 2, the best known payoff is $x \geq \mu_G$ and it is communicated to $I - j$ employees, so that the j other employees experiment in period 2. In that case,

$$V(x, j) \quad = \quad (I-j)x + j\mu_G + \delta I \left[x + \int_x^{\overline{v}} (u-x)\, dG_{(j)}(u) \right].$$

Let $\Delta V(x, j)$ denote the gain from an increase from j to $j+1$ employees who experiment, i.e., $\Delta V(x, j) = V(x, j+1) - V(x, j)$. We have, for $j < I$,

$$\begin{aligned} \Delta V(x, j) \quad &= \quad \left((I-j-1)x + (j+1)\mu_G + \delta I \left(x + \int_x^{\overline{v}} (u-x)\, dG_{(j+1)}(u) \right) \right) \\ &\quad - \left((I-j)x + j\mu_G + \delta I \left(x + \int_x^{\overline{v}} (u-x)\, dG_{(j)}(u) \right) \right) \\ &= \quad \mu_G - x + \delta I \left(\int_x^{\overline{v}} (u-x)\, dG_{(j+1)}(u) - \int_x^{\overline{v}} (u-x)\, dG_{(j)}(u) \right), \end{aligned}$$

which (using Leibniz' rule) leads to

$$\frac{d\Delta V(x, j)}{dx} \quad = \quad -1 + \delta I \left(-\int_x^{\overline{v}} dG_{(j+1)}(u) + \int_x^{\overline{v}} dG_{(j)}(u) \right) = -1 - \delta I (G_{(j)}(x) - G_{(j+1)}(x)) < 0.$$

By monotone comparative statics Milgrom and Roberts (1990, 1994), it follows that the optimal number of employees that experiment decreases with x. Moreover, at $x = \mu_G$,

$$\Delta V(\mu_G, j) = \delta I \left(\int_{\mu_G}^{\overline{v}} (u-\mu_G)\, dG_{(j+1)}(u)) - \int_{\mu_G}^{\overline{v}} (u-\mu_G)\, dG_{(j)}(u) \right) > 0,$$

since $G_{(j+1)}$ first-order stochastically dominates $G_{(j)}$ and the integrand increases. Since this holds for all values of j, it follows that at $x = \mu_G$, it is optimal for all employees to experiment. Since the difference is strict and $\Delta V(x, j)$ is continuous in x, there exists $\hat{v}_I > \mu_G$ such that this holds for all $x \leq \hat{v}_I$. Furthermore, at $x = \overline{v}$, $\Delta V(\overline{v}, j) = \mu_G - \overline{v} < 0$, so that it is optimal to have no one experiment. Again, since the difference is strict and continuous in x, there exists $\hat{v}_1 < \overline{v}$ such that this holds for all $x > \hat{v}_1$.

We proved the proposition for $\mu_G < \hat{v}_I \leq \cdots \leq \hat{v}_1 < \overline{v}$ and are left to show that the weak inequalities are in fact strict. Note that, by definition of $\hat{v}_j$ and by the fact that $\Delta V(x, j)$ is continuous and strictly decreasing in x, $V(\hat{v}_{j+1}, j+1) = V(\hat{v}_{j+1}, j)$ or $\Delta V(\hat{v}_{j+1}, j) = 0$. It now suffices to show that $\Delta V(x, j)$

0	1	2	3
Setup	Period 1	Period 2	Period 3
a Firm chooses $n \geq 0$.	a New generation of employees starts.	a Each v_{a_j} gets redrawn with independent probability q.	a Each v_{a_j} gets redrawn with independent probability q.
b Performance $v_{a_j} \overset{\text{iid}}{\sim} G$ are drawn.	b Employees simultaneously choose actions a_j.	b New generation of employees starts.	b New generation of employees starts.
	c Firm decides which (a_j, v_{a_j}) to store.	c Employees simultaneously choose actions a_j.	c Employees simultaneously choose actions a_j.
		d Firm decides which (a_j, v_{a_j}) to store.	

Figure 2: Time line of model with change

strictly decreases in j. To see this, note that

$$\begin{aligned}
&\Delta V(x,j) - \Delta V(x,j-1) \\
&= \mu_G - x + \delta I \left(\int_x^{\overline{v}} (u-x)\, dG_{(j+1)}(u) - \int_x^{\overline{v}} (u-x)\, dG_{(j)}(u) \right) \\
&\quad - \left(\mu_G - x + \delta I \left(\int_x^{\overline{v}} (u-x)\, dG_{(j)}(u) - \int_x^{\overline{v}} (u-x)\, dG_{(j-1)}(u) \right) \right) \\
&= \delta I \left(\int_x^{\overline{v}} (u-x)\, g_{(j+1)}(u) du - \int_x^{\overline{v}} (u-x)\, g_{(j)}(u) du - \int_x^{\overline{v}} (u-x)\, g_{(j)}(u) du + \int_x^{\overline{v}} (u-x)\, g_{(j-1)}(u) du \right) \\
&= \delta I \int_x^{\overline{v}} (u-x) \frac{d\big[G^{j+1} - 2G^j + G^{j-1}\big]}{du}\, du \\
&= -\delta I \int_x^{\overline{v}} (u-x) \frac{d\big[(1-G)(G^j - G^{j-1})\big]}{du}\, du \\
&= -\delta I \left[(u-x)(G^j(u) - G^{j-1}(u))(1-G(u)) \right]_x^{\overline{v}} + \delta I \int_x^{\overline{v}} (G^j - G^{j-1})(1-G) du \\
&= 0 + \delta I \int_x^{\overline{v}} (G^j - G^{j-1})(1-G) du < 0.
\end{aligned}$$

where the one-before-last step is an integration by parts and we use the fact that $dG = g\, du$ and $g_{(j)}(u) = \frac{dG_{(j)}(u)}{du} = \frac{dG^j(u)}{du}$. This concludes the proof. ■

Note that the incremental benefit from an extra employee experimenting decreases with the number of employees that are already experimenting. The incremental cost, however, is linear since it always equals $v_{\hat{a}} - \mu_G$. This implies that there is a finite optimal number of employees who should experiment.

5 The Effect of Change

A marked result of Section 3 was that the firm will, at most, record only one best practice. We now show how this may be different when the actions' performance may change over time, e.g., when changes in the environment can make formerly successful methods obsolete.

To formally capture in a straightforward manner the idea that performance may change over time, we will assume that at the beginning of each period, each action's performance is redrawn with independent probability q. That is, with probability q, v_{a_j} is redrawn from distribution G at the beginning of period t, while with complementary probability $1 - q$, it remains the same as in period $t - 1$. The firm and employees do not directly observe whether a payoff has been redrawn (although they will obviously learn this fact when the action is tried again). The timing of the modified game is shown in Figure 2.

To keep the analysis tractable, we will further focus here on a specific case. In particular, assume that the "actions" are potential solutions to a particular problem. Each potential solution either solves the problem, resulting in a success with payoff 1, or does not solve the problem, resulting in a failure with payoff 0. Ex-ante, however, it is not known which are the successful solutions. All employees share a common prior that each potential solution has an independent probability p of being successful, i.e., G is a Bernoulli distribution with parameter p and support $\{0, 1\}$. We can think of $(1 - p)$ as the difficulty of the problem, i.e., high p is an easy problem whereas low p is a difficult problem. We will also assume that $\delta = 1$. As before, the cost of one unit of capacity is c.

We find that if an action's performance may change over time, the firm may want to record more than the best-performing action, i.e., $n^* > 1$, because the additional information serves as backup for when the best practice becomes obsolete. Such backup information will be most useful when the rate of change is intermediate: when change is fast (q is large) yesterday's know-how is less valuable today, limiting the benefit of such extra information; when change is slow (q is small) there is no need to remember more than only the best practice, since the probability that this practice becomes obsolete is small.

Second, from an ex-ante perspective – in particular when designing the knowledge system – the ability to store know-how (and backup know-how) is most useful for moderately difficult problems, i.e., at medium levels of p. For very easy problems (very high p) the solution becomes so obvious that a knowledge system is unnecessary; for very difficult problems (very low p) the probability of discovering a solution becomes so small that there is typically not much to record. Therefore, we will tend to see formal knowledge systems for moderately difficult problems.

The benefit of recording backup solutions also changes in I and c in the same way and for the same reason as was the case for the value of the knowledge system itself. As in the static case, investment in the knowledge system decreases in c: as the cost of adding capacity increases, the expected net benefit of an extra unit of capacity decreases, leading to less recorded knowledge. For the number of employees I, note that we have again the two scale effects that make knowledge systems more valuable for larger organizations: the average quality of know-how is higher since there is more experimentation, and there are also returns to scale in using the information.

The following proposition formalizes these results. It determines the marginal benefit from an extra unit of capacity and derives the comparative statics for the optimal size n^* of the knowledge system. To state this proposition, let $\Pi(n)$ be the expected payoff per employee when information on n payoffs can be recorded and transferred from period $t - 1$ to period t. The marginal benefit of a unit of capacity is then $\Delta\Pi(n) = \Pi(n) - \Pi(n - 1)$.

Proposition 3 *The benefit from adding an additional unit of capacity is*

$$\Delta\Pi(n) = \sum_{j=n}^{I} \frac{I!}{(I-j)!j!} p^j (1-p)^{I-j} (1-(1-p)q)(1-p)^n q^{n-1} (1-q).$$

This marginal benefit function $\Delta\Pi(n)$ is quasiconcave in q, decreases in n, and increases in I. Moreover, $\Delta\Pi(1)$ is also strictly quasi-concave in p.

The optimal capacity of the knowledge system n^* is quasiconcave in q, increases in I, and decreases in c. Moreover, there exists $0 < \underline{p} \leq \overline{p} < 1$ such that $n^* \geq 1$ if and only if $p \in (\underline{p}, \overline{p})$.

Proof : The proof is in appendix. ■

Quasiconcavity of $\Delta\Pi(n)$ in q implies that the value of having the n^{th} piece of backup information is "single-peaked" in q, so that backup information is most valuable in medium change environments. Knowledge systems (with or without backup information) are also most valuable for moderately difficult problems, i.e., $\underline{p} \leq p \leq \overline{p}$.[8]

Figure 3 illustrates the quasiconcavity of the marginal benefit of capacity, $\Delta\Pi(n)$, in q and in p. It shows the (positive) incremental gains – net of the cost of a unit of capacity – from the first three units of capacity (used for recording successes) for both $c = 0.2$ and $c = 0.01$, for a firm with $I = 3$ employees. Note that whenever the graph is positive, it is optimal for the firm to increase the capacity of its knowledge system. (To accentuate the areas where the firm should add the extra capacity, we added the zero surface to each graph.) Since the first unit of capacity determines whether the firm has a knowledge system at all, the first graph shows under which conditions the firm should invest in a system to record know-how. The two other graphs show under which conditions the firm should add the second and third units of capacity. For the case where $c = 0.2$, the firm will build a knowledge system with at most one unit of capacity. It will build such a system when low q and intermediate p are combined. In the $c = 0.01$ case, on the other hand, it nearly always pays to have a knowledge system except when extremely rapid change (unstable environment) is combined with either very easy (high p) or very difficult (low p) problems. The second and third graphs for this case illustrate that the optimal size is indeed largest for intermediate values of q and p.

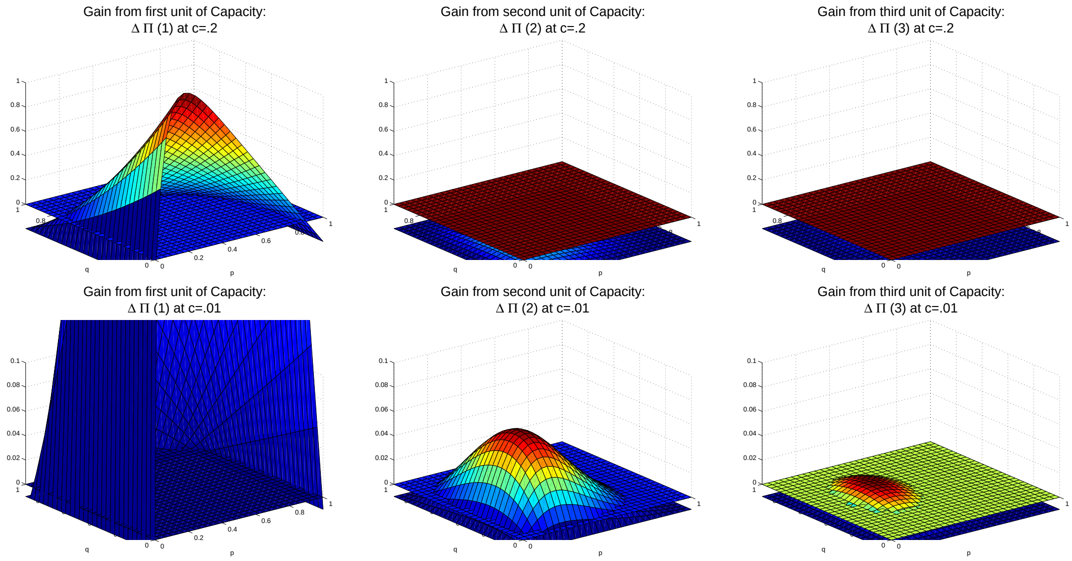

Figure 3: Payoff gains per employee from the first 3 successive units of capacity with $I = 3$, and respectively $c = 0.2$ (upper graphs) and $c = 0.01$ (lower graphs).

[8]Both the intuition and the equations suggest that $\Delta\Pi(n)$ is actually quasiconcave in p for *all* n, but a proof awaits further research.

6 Further Extensions

Thus far, we have considered an environment with an infinite action set. Although business conditions call for creative managerial solutions that usually cannot be exhaustively listed, there are situations where it makes more sense to think in terms of a finite set of options. For example, production workers often have to choose from a finite number of tools to perform each production operation, translating to a finite set of actions in the framework of our model. To see how this would affect the conclusions, suppose that there is a finite set of J actions and I employees. Each action's payoff is still an independent draw from G. Suppose now first that the firm records a low-payoff action a^-, with payoff $v^- < \mu_G$, so that employees can avoid it. Relative to the case when no info is recorded, each employee's expected payoff for the next period increases by $\frac{\mu_G - v^-}{J}$. Essentially, employees will have an equal probability of picking any action except the one they know gives a lower than average payoff. Suppose next that the firm remembers a high-payoff action a^+, with payoff $v^+ > \mu_G$, so that employees can replicate that action. Relative to the case when no info is recorded, each employee's expected payoff for the next period increases by $\frac{J-1}{J}(v^+ - \mu_G)$. Clearly, for the same deviation from the mean, it is still better to record successes than to record failures (unless the firm wants employees to experiment). Moreover, there is still no need to record more than one successful action. So it follows that with a finite set of actions, recording failures may make sense but only temporarily until a sufficiently good best practice is discovered and only if the failures are rather extreme. Note that our result that successes are more valuable than failures does *not* imply that organizations cannot learn from failure. On the contrary, knowing how to avoid a misstep can be a huge advantage. However, knowing a course of action that results in success will also serve this purpose and often more.

Another case where remembering failures may be valuable is the following. Suppose that it is commonly known that among 10 possible actions, there is one with payoff of -5, one with payoff of +1, and the rest with payoffs 0. It is, however, not known which actions correspond to which payoffs. In this case, it would be better to record a -5 payoff than a 0 payoff since employees in future generations will know not to select the "worst failure" option. While in this case, failures do have more informational value than successes, it is difficult to find realistic examples of such a setting where the exact *realized* distribution of payoffs is known but not which action corresponds to which payoff. What often does happen is that payoffs are partially correlated, with the outcome of one action causing an update in the prior beliefs regarding the other actions without revealing them completely. We believe that the qualitative results of the paper extend to that case, although that requires further research.

A natural setting where it may pay to remember failures is the situation where not all actions look the same ex-ante, and some of the actions that look most promising actually lead to failures. Nevertheless, even in this case, successes still seem more useful than failures. Also the other results, such as the competency trap and the comparative statics, seem to apply. This is, however, an important avenue for future research.

Another interesting extension is to consider employees' incentives to share their best practices. If turnover from period to period is less than 100%, and an employee is evaluated relative to his peers, he may not be willing to share information that could benefit others. The firm would have to create an incentive system so that employees find it utility maximizing to reveal information about good and bad outcomes, as appropriate. The informational rents would increase the cost of the knowledge system, thereby affecting the optimal strategy for information collection and optimal knowledge system capacity.

7 Concluding Remarks

For many firms, the ability to create, organize, and disseminate know-how is a key factor in their ability to succeed. We studied some basic questions on the management of know-how. We found that information about successes is generally more useful than information about failures because past successes can be used to guide future actions whereas failures just point out the pitfalls to avoid. The larger the set of possible actions, the less useful information on failures is because the likelihood of repeating the failures decreases.

In codifying successful actions, the firm must trade off its desire to exploit current best practice versus the potential for even better performance through experimentation. Therefore, even if the firm identifies moderately successful practices, it may not want to codify these practices, thereby forcing employees to continue experimenting. Larger firms benefit more from knowledge systems in both the experimentation (knowledge generation) and exploitation (knowledge application) phases. The larger the firm is, the greater the number of employees experimenting with different solutions, therefore, the more likely that someone will discover a high-performance solution. Once a home-run discovery is made, the impact on the organization is also greater because there are more employees to leverage this information. These scale effects also imply that larger firms should also be more selective when codifying information. Other factors that make a firm value a knowledge system, and simultaneously increase its optimal selectivity, are high levels of turnover and facing problems for which there is high uncertainty about the right solutions. Finally, if the firm is able to control who can access the know-how information, it can reap the benefits of experimentation *and* exploitation by communicating "mediocre" practices to a subset of employees (for example, within one plant), while others (throughout the rest of the firm) continue to experiment.

We also studied how potential changes over time in the effectiveness of practices impact the optimal management of know-how. In slow-changing environments, there is no need to keep historical information because simply following the "best practice" will ensure future success so that small knowledge systems suffice. In fast-changing environments, knowledge system capacity should also be small because information from the past will quickly become obsolete. Therefore, we expect to find the largest knowledge systems in medium-stability environments. We also found that knowledge systems were most useful for medium-difficulty problems. Employees can figure out easy problems for themselves while solutions to extremely difficult problems are unlikely to be found so that there is (in expectation) little useful information to pass on.

The framework in this paper can be used to explore other questions on knowledge management. We mentioned already the role of incentives and some variations on the model in this paper. Other potential avenues are, for example, the case where an action's performance is observed only by the employee undertaking it – who can then send cheap talk messages – or the interaction between knowledge management and moral hazard. As knowledge management continues to grow in importance, a systematic economic perspective may shed some important insights on such issues.

A Proofs

Proof of Proposition 3: Note that, since $\#A = \infty$, the employees' prior on any action on which no information is stored is always G, i.e., the expected payoff from an action about which nothing is known is p. If $n = 0$, then employees simply undertake random actions each period and $\Pi(0) = 3p$. We will henceforth assume that $n \geq 1$. Recall that $B_t \subset A$ is the set of actions from previous periods that has been recorded and communicated to employees. Note that $\#B_t$ cannot exceed the capacity of the knowledge system, therefore, $\#B_t \leq n$. As before, recording a failure allows the employees to avoid that action in the future, but with $A = \infty$ the benefit of doing so is zero. Therefore (by the assumption on the firm's preferences when payoff-indifferent) the firm will prefer not to do so. It follows that the firm will only record successes. Let then $\bar{B}_t \subseteq B_t$ be the set of actions that was shown to be successful last period, period $t-1$. Let $\bar{\bar{B}}_t \subseteq B_t$ be the set of actions that was shown to be successful two periods ago, period $t-2$, (and have not been tried since). By the earlier argument regarding successes versus failures and the fact that $t \leq 3$, $B_t = \bar{B}_t \cup \bar{\bar{B}}_t$.

We now first establish the optimal behavior of the firm and employees when there is a knowledge system with capacity n. First consider period 3, the last period. The employee's expected payoff from an action that was shown to be successful in period 2, $\bar{a}_j \in \bar{B}_3$, is $(1-q)+qp$. The employee's expected payoff from an action that was shown to be successful in period 1 (but has not been tried since), $\bar{\bar{a}}_j \in \bar{\bar{B}}_3$, is $(1-q)^2 + (1-(1-q)^2)p$. Finally, any $a_j \in A \setminus B_t$ has expected payoff p. Therefore, the ordering of actions according to decreasing expected payoffs is: $\bar{a}_j$, $\bar{\bar{a}}_j$, and $a_j \in A \setminus B_3$. Thus, the optimal employee strategy is as follows: if $\bar{B}_3 \neq \emptyset$ choose an action from $\bar{B}_3$, else if $\bar{\bar{B}}_3 \neq \emptyset$ choose an action from $\bar{\bar{B}}_3$, else choose an action from $A \setminus B_3$. This is also the optimal strategy for the firm (in period 3). It follows that if there is at least one success in period 2, then that action will be recorded (and will be the one that employees choose in period 3).

Consider the second period. If $B_2 = \bar{B}_2 = \emptyset$, then all employees choose actions $a_j \in A$ and the expected payoff per employee is p. Moreover, all employees will choose different actions (since they are indifferent but the firm prefers maximal experimentation) and the firm will remember one successful action (if one occurred). The expected payoff per employee for periods 2 and 3 is then

$$\Pi = p + (1-p)^I p + (1-(1-p)^I)((1-q)+qp) = 2p + (1-(1-p)^I)(1-q)(1-p).$$

If $\#\bar{B}_2 = k > 0$, then employees are indifferent among any such actions. The firm benefits from verifying the payoffs of as many actions as possible, i.e., learning whether the payoffs have changed, since that will maximize the probability that employees undertake a successful action in period 3 (while there's no cost to the firm from having employees choose different actions from $\bar{B}_2$). It follows that the employees will coordinate to undertake as many different actions among that set of formerly successful actions as possible. With I employees, there can be at most I different actions that were undertaken in period 1 and thus at most I successes that were stored. Since there are (weakly) more employees than recorded successes, the employees will try all recorded successes in period 2. If $\#\bar{B}_2 = k$, then the expected payoff for periods 2 and 3 is

$$\begin{aligned}\Pi &= \underbrace{(1-q)+qp}_{\text{period 2}} + \underbrace{(1-q^k)\left[(1-q)+qp\right] + q^k\left[(1-(1-p)^k)[(1-q)+qp] + (1-p)^k p\right]}_{\text{period 3}} \\ &= 2p + (1-q)(1-p) + (1-q^k)(1-q)(1-p) + q^k(1-(1-p)^k)(1-q)(1-p) \\ &= 2p + 2(1-q)(1-p) - q^k(1-p)^k(1-q)(1-p)\end{aligned}$$

which clearly increases in k. It follows that the firm will store all the successes it can in period 1.

Consider finally the first period. Since no actions are known, all employees are indifferent among all actions. Given this indifference, they will coordinate on choosing different actions (since that is what the firm prefers). The expected first-period payoff is thus always p per employee.

At this point we have established that:

1. The firm will always store information about successes: in the first period all successes up to a maximum of n and in the second period one success, if there is one.
2. In the first period, employees will undertake I different actions.
3. In the second period, employees will try all stored successes if there are any, and undertake I different new actions otherwise.
4. In the third period, employees will undertake any stored success if there is any, and undertake a random action otherwise.

Given the behavior of the firm and employees, we now derive the marginal value of a unit of knowledge system capacity. We can the write the expected payoff per employee with n units of capacity as follows, with x denoting the number of successes in period 1.

$$\begin{aligned}
\Pi(n) &= p + P[x=0]\left[2p + (1-(1-p)^I)(1-q)(1-p)\right] \\
&\quad + \sum_{k=1}^{n-1} P[x=k]\left[2p + 2(1-q)(1-p) - q^k(1-p)^k(1-q)(1-p)\right] \\
&\quad + P[x \geq n]\left[2p + 2(1-q)(1-p) - q^n(1-p)^n(1-q)(1-p)\right] \\
&= 3p + P[x=0]\left[(1-(1-p)^I)(1-q)(1-p)\right] \\
&\quad + \sum_{k=1}^{n-1} P[x=k]\left[2(1-q)(1-p) - q^k(1-p)^k(1-q)(1-p)\right] \\
&\quad + P[x \geq n]\left[2(1-q)(1-p) - q^n(1-p)^n(1-q)(1-p)\right]
\end{aligned}$$

To determine the marginal benefit function, $\Delta\Pi(n) = \Pi(n) - \Pi(n-1)$, note that for $n \geq 2$, $\Pi(n)$ and $\Pi(n-1)$ differ only in how much expected utility $x \geq n$ generates: in $\Pi(n-1)$ they are lumped together with $x = n-1$. Therefore, we get that:

$$\begin{aligned}
\Delta\Pi(n) &= P[x \geq n]\left[[2-(1-p)^n q^n](1-q)(1-p) - \left[2-(1-p)^{n-1}q^{n-1}\right](1-q)(1-p)\right] \\
&= P[x \geq n]\left[1-(1-p)q\right](1-p)^n q^{n-1}(1-q) \\
&= \sum_{j=n}^{I} \frac{I!}{(I-j)!j!} p^j (1-p)^{I-j}(1-(1-p)q)(1-p)^n q^{n-1}(1-q).
\end{aligned}$$

For $n = 1$, it is slightly different:

$$\begin{aligned}
\Delta\Pi(1) &= P[x=0][1-(1-p)^I](1-q)(1-p) + P[x \geq 1]\left[2-(1-p)q\right](1-q)(1-p) \\
&= (1-p)^I[1-(1-p)^I](1-q)(1-p) + (1-(1-p)^I)\left[2-(1-p)q\right](1-q)(1-p) \\
&= [1-(1-p)^I](1-q)(1-p)\left[2-(1-p)q+(1-p)^I\right].
\end{aligned}$$

For the comparative static of $\Delta\Pi(n)$ in n, note now that:

$$\begin{aligned}
\Delta\Pi(2) &= P[x \geq 2](1-(1-p)q)(1-p)^2 q(1-q) < P[x \geq 1](1-q)(1-p) \\
&< P[x \geq 1](2-(1-p)q)(1-q)(1-p) \qquad < \Delta\Pi(1).
\end{aligned}$$

From this and inspection of the $\Delta\Pi(n)$ formula above, it follows that $\Delta\Pi(n)$ indeed monotonically decreases in n.
For the comparative static in I, consider first the case where $n = 1$. The derivative for I is:

$$\begin{aligned}
\frac{d\Delta\Pi(1)}{dI} &= +[(1-p)^I - (1-p)^{2I}](1-q)(1-p)\log(1-p) \\
&\quad + (-(1-p)^I)\left[2-(1-p)q\right](1-q)(1-p)\log(1-p) \\
&= (1-p)^I[-(1-p)^I + (1-p)q - 1](1-q)(1-p)\log(1-p),
\end{aligned}$$

which is positive since $\log(1-p) < 0$ and $(1-p)q < 1$. For $n \geq 2$, $P[x \geq n]$ increases in I, so that $\Delta\Pi(n)$ also increases in I.
For quasi-concavity of $\Delta\Pi(n)$ in q, note that for $n = 1$

$$\frac{d\Delta\Pi(1)}{dq} = -(1-p)^I[1-(1-p)^I](1-p) + (1-(1-p)^I)\left[-2-(1-p)(1-2q)\right](1-p) < 0,$$

while for $n \geq 2$,

$$\begin{aligned}
\frac{d\Delta\Pi(n)}{dq} &= P[x \geq n](1-p)^n \frac{d\left[1-(1-p)q\right]q^{n-1}(1-q)}{dq} \\
&= P[x \geq n](1-p)^n q^{n-2}\left\{-(1-p)q(1-q) - \left[1-(1-p)q\right]q + (n-1)\left[1-(1-p)q\right](1-q)\right\}
\end{aligned}$$

for which the sign depends on the expression between brackets. That expression can be written

$$\begin{aligned}
&-(1-p)q(1-q) - \left[1-(1-p)q\right]q + (n-1)\left[1-(1-p)q\right](1-q) \\
&\quad = (n-1) - n(2-p)q + (n+1)(1-p)q^2
\end{aligned}$$

is a quadratic expression which is positive at $q = 0$ and negative at $q = 1$ from which it follows that the derivative is indeed first increasing and then decreasing. This establishes quasi-concavity.

Quasiconcavity in p for $n = 1$ follows from standard but tedious algebra and calculus, which is available from the authors.

We now turn to the results on the optimal capacity of the knowledge system n^*. Consider first the effect of q on n^*. Since $\Delta\Pi(n)$ is quasiconcave in q, the area (in q-space) where $\Delta\Pi(n) \geq c$ is convex, i.e., it is a closed interval. Denote this area as D_n. Since $\Delta\Pi(n)$ decreases in n, $D_{n+1} \subseteq D_n$. Moreover, for a given $\tilde{q}$, n^* is optimal if $\tilde{q} \in D_{n^*-1}$ but $\tilde{q} \notin D_{n^*}$. This means that D_{n^*-1} is actually the upper contour set for n^*, so that the upper contour sets are convex. This implies that n^* is quasiconcave in q. The result that $n^* \geq 1$ iff $p \in (\underline{p}, \overline{p})$ follows analogously from the strict quasi-concavity of $\Delta\Pi(1)$ in p (and the fact that $\Delta\Pi(1) = 0$ at both $p = 0$ and $p = 1$).

Since $\Delta\Pi(n)$ decreases in n and adding an extra unit of capacity is profitable if and only if $\Delta\Pi(n) \geq c$, at the optimal n^*, $\Delta\Pi(n^* - 1) \geq c > \Delta\Pi(n^*)$. Consider the effect of increasing I on n^*. Since $\Delta\Pi(n)$ increases in I, $\Delta\Pi(n^* - 1)$ increases so that $\Delta\Pi(n^* - 1) \geq c$ still holds. It follows that the optimal n^* must weakly increase. The argument for c is analogous. This completes the proof. ∎

References

Alonso, R., W. Dessein, and N. Matouschek (2006). When does coordination require centralization? Working Paper.

Argote, L. (1999). *Organizational Learning: Creating, Retaining and Transferring Knowledge.* Norwell, MA: Kluwer Academic.

Argote, L., B. McEvily, and R. Reagans (2003, April). Managing knowledge in organizations: An integrative framework and review of emerging themes. *Management Science 49*(4), 571–582.

Athey, S. and J. Levin (2001). The value of information in monotone decision problems. Working Paper, MIT and Stanford University.

Bartlett, C. A. (1996). Mckinsey & company: Managing knowledge and learning. Harvard Business School, Case 9-396-357.

Bartlett, C. A., R. M. Kanter, R. S. Kaplan, and M. C. Jensen (2003). *Best Practice: Ideas and Insights from the World's Foremost Business Thinkers.* Cambridge, MA: Perseus Publishing.

Berry, D. E. and B. Fristedt (1985). *Bandit Problems: Sequential Allocation of Experiments.* New York: Chapman and Hall.

Blackwell, D. (1951). Comparisons of experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics*, 93– 102.

Bolton, P. and M. Dewatripont (1994). The firm as a communication network. *Quarterly Journal of Economics 109*(4), 809– 839.

Brockman, B. K. and R. M. Morgan (2003, Spring). The role of existing knowledge in new product innovativeness and performance. *Administrative Science Quarterly 34*(2), 385–419.

Canon, M. D. and A. C. Edmondson (2001, March). Confronting failure: antecedents and consequences of shared beliefs about failure in organizational work groups. *Journal of Organizational Behavior 22*(2), 161–177.

Carrillo, J. E. and C. Gaimon (2004, November). Managing knowledge-based resource capabilities under uncertainty. *Management Science 50*(11), 1504–1518.

Christiansen, C. M. and B. Baird (1998, July 27). Cultivating capabilities to innovate: Booz-allen & hamilton. Case 9-698-027, Harvard Business School.

Cohen, W. M. and D. A. Levinthal (1990, March). Absorptive capacity - a new perspective on learning and innovation. *Administrative Science Quarterly 35*(1), 128–152.

Garicano, L. (2000, October). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874– 904.

Gittins, J. C. (1989). *Multi-Armed Bandit Allocation Indices.* New York: John Wiley & Sons.

Hilton, R. W. (1981, January). The determinants of information value: Synthesizing some general results. *Management Science 27*(1), 57–64.

Holsapple, C. W. (Ed.) (2003). *Handbook on Knowledge Management.* Berlin, Germany: Springer-Verlag.

Leonard-Barton, D. (1995). *Wellsprings of knowledge: Building and sustaining the sources of innovation.* Cambridge, MA: Harvard Business School Press.

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

MacCormack (2002, November 5). Siemens sharenet: Building a knowledge network. Case 9-603-036, Harvard Business School.

Manso, G. (2006). Motivating innovation. Working Paper.

Margolis, J. and D. Upton (1996, September 23). Mcdonald's corporation. Case 9-693-028, Harvard Business School.

Marschak, J. and R. Radner (1972). *Economic Theory of Teams.* Yale University Press.

Milgrom, P. and J. Roberts (1990, June). The economics of modern manufacturing : Technology, strategy, and organization. *American Economic Review 80*(3), 511– 528.

Milgrom, P. and J. Roberts (1994, June). Comparing equilibria. *American Economic Review 84*(3), 441–459.

Miner, A. S., J. Y. Kim, I. W. Holzinger, and P. Haunschild (1999). Fruits of failure: Organizational failure and population-level learning. *Advances in Strategic Management 16*, 187–220.

O'Dell, C. and C. J. Grayson (1998, Spring). If only we knew what we know: Identification and transfer of internal best practices. *California Management Review 40*(3), 154–170.

Radner, R. (1993, September). The organization of decentralized information processing. *Econometrica 61*(5), 1109– 46.

Sitkin, S. B. (1992). Learning through failure: the strategy of small loses. *Research in Organizational Behavior 14*, 231–266.

Sorenson, O. (2003, April). Interdependence and adaptability: Organizational learning and the long-term effect of integration. *Management Science 49*(4), 446–463.

Stern, S. (2004). *Biological Resource Centers: Knowledge Hubs for the Life Sciences.* Washington, DC: Brookings Institution Press.

Stewart, T. A. (2000, November 27). Knowledge worth $1.25 billion. *Fortune 142*(13), 302–303.

Van Zandt, T. (1999, July). Real-time decentralized information processing as a model of organizations with boundedly rational agents. *Review of Economic Studies 66*(3), 633– 658.