

e - c o m p a n i o n

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—"Managing Know-How"
by Deishin Lee and Eric Van den Steen,
Management Science, DOI 10.1287/mnsc.1090.1101.

Online Appendix to 'Managing Know-How'

Deishin Lee and Eric Van den Steen*

September 2009

A Proofs

Proof of Proposition 1: 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 $A \setminus B$ for $t = 2$ and $A \setminus \tilde{B}_t$ for $t = 3$. Given this employee behavior, it follows that (i) the firm never records an action with $v_{a_j} \leq \mu_G$: doing so (at most) 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 proposition. ■

Proof of Proposition 2: We prove this by backwards induction. From Proposition 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 $B_2 = \emptyset$ or $v_{\hat{a}_2} \leq \mu_G$ then period-2 employees will select actions at random from $A \setminus B_2$, 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 $B_2 = \emptyset$ (or $v_{\hat{a}_2} \leq \mu_G$) equals (given that $\tilde{v}_2 = \mu_G$)

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

*Harvard Business School (dlee@hbs.edu; contact author) and Harvard Business School (evandensteen@hbs.edu).

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

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

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

Proof of Proposition 3b: 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 2 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 $x > \mu_G$ (and $\mu_F = \mu_G$). 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. ■

Proof of Proposition 3a: 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 2, $V_2 = (1 + \delta)\mu_G + \delta \int_{\mu_G}^{\infty} (v - \mu_G) dG_{(I)}(v)$. The value of a knowledge system with $n = 1$ equals

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

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 N for the first-period G and M for the second-period G , we can write the above expression as

$$\frac{\delta}{1 + \delta} P_N[v_{(I)} \leq \tilde{v}_1] \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) + \int_{\tilde{v}_1}^{\infty} (u - \mu_N) dN_{(I)}(u) \quad (3)$$

$$= \frac{\delta}{1 + \delta} \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) + \int_{\tilde{v}_1}^{\infty} \left(u - \mu_N - \frac{\delta}{1 + \delta} \int_{\mu_M}^{\infty} (v - \mu_M) dM_{(I)}(v) \right) dN_{(I)}(u) \quad (4)$$

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 3b, it is clear from (3) that a mean-preserving stretch of M increases the whole expression, and again analogously for N and (4), using then the fact that $\bar{v}_1 > \mu_G$. This proves the proposition. \blacksquare

Proof of Proposition 4: 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,

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

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 $j < I$,

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

which (using Leibniz' rule) leads to

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

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

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

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 $x = \bar{v}$, $\Delta V(\bar{v}, j) = \mu_G - \bar{v} < 0$, so that it is optimal to have no one experiment. Again, since the difference is strict and continuous in x , there exists $\hat{v}_1 < \bar{v}$ such that this holds for all $x > \hat{v}_1$.

We proved the proposition for $\mu_G < \hat{v}_I \leq \dots \leq \hat{v}_1 < \bar{v}$ 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 $\Delta V(x, j)$

strictly decreases in j . To see this, note that

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

where the one-before-last step is an integration by parts and we use the fact that $dG = g du$ and $g_{(j)}(u) = \frac{dG_{(j)}(u)}{du} = \frac{dG^j(u)}{du}$. (Note that if $\mu_G < v_a \leq \hat{v}_I$ then \hat{a} is recorded but not communicated to anyone in period 2. The information will be communicated in period 3 if no better action is discovered.) This concludes the proof. \blacksquare

Proof of Proposition 5: 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 (among B_t) that were shown to be successful last period, period $t - 1$. Let $\bar{\bar{B}}_t \subseteq B_t$ be the set of actions (among B_t) that were 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 almost surely different actions (since they are indifferent and there are an infinite number of potential actions) 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

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

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 and will choose at random (and, with an infinite number of potential actions, almost surely select I different actions). 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 period 1, employees undertake (almost surely) I different actions.
3. In period 2, employees try all stored successes if there are any, and undertake (almost surely) I different new actions otherwise.
4. In period 3, employees 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 write the expected payoff per employee with n units of capacity as follows, with x denoting the number of successes in period 1.

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

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:

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

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

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

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

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

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:

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

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 $n = 1$

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

while for $n \geq 2$,

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

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

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

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 $p \in (\underline{p}, \bar{p})$ 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. ■

B More Discussion of Robustness, Limitations, and Extensions

This appendix presents more formal results and further discussion, first of the issues discussed in Subsection 3.2 and then of the issues discussed in Section 6.

B.1 More Detailed Analysis and Discussion of Robustness

No Overlap among Generations We start with the formal model with overlap among the different groups or generations of employees. Consider, for that purpose, the model of Section 2 with one employee in each period ($I = 1$) and no discounting ($\delta = 1$). To capture a potential overlap in generations, we will assume that each employee will stay on for the next period with some positive probability α (and that an employee does not forget information about payoffs). To fix things, assume that the employee learns (privately) whether he will stay for the next period at the beginning of the current period. Therefore, in each period: i) the employee learns privately whether he stays for the next period or leaves at the end of the current period, ii) the employee selects an action, iii) the firm decides whether to record the action in the knowledge system (not knowing whether the employee stays or not). The assumptions on timing and information simplify the analysis (considerably) but do not drive the essential results. (An alternative specification that would also be analytically tractable is an infinite-period model where it is revealed at the start of a new period whether we keep last period's employee or draw a new employee.) The following proposition indeed shows that the firm will still record only successes, will record at most one such success, and is still subject to the competency trap.

Proposition 1 *An optimal knowledge system has at most one unit of capacity ($n^* \leq 1$). In each period t , at most the action with the highest payoff, $\tilde{a}_t \in B_t \cup C_t$, is recorded. In particular, there exists thresholds $\tilde{v}_1 > \tilde{v}_2 = \mu_G$, 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.*

Proof : Note first that in the first period, the employee always (necessarily) chooses a random action (with expected payoff μ_G). For the rest of the proof, we will follow a backwards induction argument. Consider first the third period. In the third period, the employee simply chooses the action with the highest expected payoff. In particular, if the highest payoff among the recorded or remembered actions

exceeds μ_G , then the employee chooses that action, else the employee chooses randomly among all actions with expected payoff μ_G (i.e. among all unknown actions plus all known action that happen to have a payoff μ_G). Consider now the decision whether and what information to record at the end of period 2. Since the firm does not know whether the employee stays, it cannot condition its policy on that event. Moreover, recording an action with payoff less than μ_G can at most cause the employee to avoid that action and, with an infinite number of potential actions, doing so does not affect the expected payoff. It is then straightforward that the firm will record (from among all actions of which it knows the payoffs) the action with the highest payoff if that payoff strictly exceeds μ_G and will not record any action otherwise.

Consider now the employee's action choice in period 2. If the employee knows that he will not stay on for the next period, the employee chooses the recorded or remembered action with the highest payoff iff that payoff is strictly larger than μ_G , and otherwise chooses a randomly selected action of which the payoff is not known. If the employee knows that he will stay on for the third period, then he can either choose the recorded or remembered action with the highest payoff, denoted \hat{v} , or choose an unknown action. Obviously, if $\hat{v} \leq \mu_G$ then it is strictly optimal to choose an unknown action (given the option value of doing so). Assume therefore that $\hat{v} > \mu_G$. Given the optimal continuation equilibrium (and the fact that the employee will remember \hat{v} even if he doesn't choose it in this period), the payoff from choosing the best recorded or remembered action equals $2\hat{v}$ while the payoff from choosing an unknown action equals $\mu_G + \hat{v} + \int_{\hat{v}}^{\infty} (u - \hat{v}) dG(u)$. There does exist a \check{v} , defined by $\int_{\check{v}}^{\infty} (u - \check{v}) dG(u) = \check{v} - \mu_G$ (which implies that $\check{v} > \mu_G$), such that the employee will choose an unknown action iff $\hat{v} \leq \check{v}$ and will choose \hat{v} otherwise.

Consider now the firm's decision what information to record at the end of period 1. (Remember that in the first period, the employee simply took a random action.) Let \check{a}_1 be the action undertaken in period 1. Since recording an action with $v \leq \mu_G$ at most makes employees avoid that action in the future (which does not affect expected payoffs since there are an infinite number of potential actions), the firm will not record any information if $v_{\check{a}_1} \leq \mu_G$. Consider next the case that $v_{\check{a}_1} > \check{v}$. If \check{a}_1 is recorded, then any period 2 employee will choose \check{a}_1 . The best known action at the end of period 2 is then again \check{a}_1 , which (since $v_{\check{a}_1} > \mu_G$) the firm will again record and the period 3 employee will again choose. The expected payoff for periods 2 and 3 from recording $v_{\check{a}_1} > \check{v}$ is thus $2v_{\check{a}_1}$. If the firm does not record the payoff, the expected payoff for periods 2 and 3 equals $\alpha(2v_{\check{a}_1}) + (1 - \alpha)[\mu_G + \mu_G + \int_{\mu_G}^{\infty} (u - \mu_G) dG(u)]$ where the first term corresponds to the situation that the first period employee stays on for the second period (and thus remembers and undertakes \check{a}_1). Some algebra then implies that the firm will always record \check{a}_1 when $v_{\check{a}_1} > \check{v}$.

Consider, finally, the case that $\mu_G < v_{\check{a}_1} \leq \check{v}$. In that case, the period 2 employee will choose \check{a}_1 iff the employee knows that he will quit at the end of the period. The payoff (to the firm for periods 2 and 3) from recording the action is then

$$\alpha[\mu_G + v_{\check{a}_1} + \int_{v_{\check{a}_1}}^{\infty} (u - v_{\check{a}_1}) dG(u)] + (1 - \alpha)2v_{\check{a}_1}$$

while the payoff from not recording the action is

$$\alpha[\alpha[\mu_G + v_{\check{a}_1} + \int_{v_{\check{a}_1}}^{\infty} (u - v_{\check{a}_1}) dG(u)] + (1 - \alpha)2v_{\check{a}_1}] + (1 - \alpha)[\mu_G + \mu_G + \int_{\mu_G}^{\infty} (u - \mu_G) dG(u)]$$

Straightforward analysis then shows that there exists some critical value \tilde{v}_t such that the firm records the action iff $v_{\check{a}_1} > \tilde{v}_t$. Note that this also implies that at most one action will be recorded. This concludes the proof. ■

No Prior Information (or Identical Distribution) Our next results are two formal models that relax the assumption that all payoffs are drawn from the same distribution and that – before any information is revealed through trial and error – neither the potential employees nor the firm have any information about these payoffs. The first model considers a variation on the basic model in which each potential employee has full knowledge about a finite number of actions. The second model considers a variation on the basic model where each potential employee has an imperfect signal about each potential action. We will study each variation in turn.

For the first variation, assume that each potential employee knows the payoffs of K randomly selected actions, i.e., as if the employee has already worked on the same problem and tried K different (randomly selected) solutions. For simplicity, assume that the firm hires one employee in each period (i.e., $I = 1$) and that there is no discounting (i.e., $\delta = 1$). Let w_t denote the highest payoff among the actions that the period- t employee knows. The following proposition shows that the firm will still record only successes, will record at most one such success, and is still subject to the competency trap.

Proposition 2 *An optimal knowledge system has at most one unit of capacity ($n^* \leq 1$). In each period t , at most the action with the highest payoff, $\check{a}_t \in B_t \cup C_t$, is recorded. In particular, there exists thresholds $\tilde{v}_1 > \tilde{v}_2 = \mu_G$, 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.*

Proof : We will follow a backwards induction argument. In the third period, the employee chooses the known or recorded action with the highest payoff if that payoff is strictly larger than μ_G , else the employee randomly chooses an action from all actions with expected payoff μ_G (i.e. from all unknown and unrecorded actions plus all known or recorded actions that happen to have a payoff equal to μ_G). Consider now the end of period 2. Given the continuation equilibrium, if $v_{\check{a}_2} > \mu_G$ then the firm records \check{a}_2 , else the firm does not record any action (since the firm prefers not to record an action when otherwise indifferent and when $v_{\check{a}_2} \leq \mu_G$ the effect of recording an action is at most that the period 3 employee will avoid that action, which has no effect on payoffs since there is an infinite number of potential actions). Consider now the employee's action choice in the second period. The employee chooses the known or recorded action with the highest payoff iff that payoff is strictly larger than μ_G , and else randomly chooses an action among all unrecorded and unknown actions. Consider finally the end of period 1. If $v_{\check{a}_1} \leq \mu_G$ then recording an action at most makes the employee avoid that action which has no effect on payoffs (since there is an infinite number of potential actions) so that the firm will record no action in that case. Assume therefore that $v_{\check{a}_1} > \mu_G$. Recording \check{a}_1 will make the employee choose that action if $v_{\check{a}_1} \geq w_2$. To determine the expected payoffs, note that in each period the distribution of w_t is G^K , which we will denote as $H = G^K$. The expected payoff for periods 2 and 3 from recording \check{a}_1 is then

$$H(v_{\check{a}_1}) \left[v_{\check{a}_1} + H(v_{\check{a}_1})v_{\check{a}_1} + \int_{v_{\check{a}_1}}^{\infty} udH(u) \right] + \int_{v_{\check{a}_1}}^{\infty} \left[v + H(v)v + \int_v^{\infty} udH(u) \right] dH(v)$$

which increases in $v_{\check{a}_1}$ and goes to ∞ as $v_{\check{a}_1} \rightarrow \infty$, while the expected payoff from not recording the action is

$$H(\mu_G) \left[\mu_G + G(\mu_G) \left(H(\mu_G)\mu_G + \int_{\mu_G}^{\infty} udH(u) \right) + \int_{\mu_G}^{\infty} \left(H(v)v + \int_v^{\infty} udH(u) \right) dG(v) \right] \\ + \int_{\mu_G}^{\infty} \left(v + H(v)v + \int_v^{\infty} udH(u) \right) dH(v)$$

Some algebra shows that the expected payoff from recording is strictly smaller as $v_{\tilde{a}_1} \downarrow \mu_G$.

It thus follows that there exists some critical value $\tilde{v}_1 > \mu_G$ such that if $v_{\tilde{a}_1} \leq \tilde{v}_1$ it is optimal not to record any action while if $v_{\tilde{a}_1} > \tilde{v}_1$, it is optimal to record \tilde{a}_1 . This concludes the proof. ■

In the second variation (to relax the assumption of no prior information), we will assume that each potential employee has some imperfect information about each action, i.e., each potential employee can make educated guesses about the actions' payoffs.

Consider, in particular, the following variation on the original model of Section 2. Each action's payoff is drawn from a finite set of discrete values, $x_1 < \dots < x_K$ with each value equally likely. (The number of potential actions is still infinite.) Moreover, each employee gets a private signal about each action's value that is fully informative with probability p and that is completely uninformative otherwise. Let $\mu = \sum_{k=1}^K x_k / K$ denote the ex ante expected value absent any private signal. When an employee has a signal \tilde{x} for some action, then the employee's expected value for that action equals $p\tilde{x} + (1-p)\mu$. For simplicity, we will assume that $\delta = 1$ (but we do not necessarily assume that $I = 1$ as we did in some of the earlier variations). The following proposition again shows that the firm will still record only successes, will record at most one such success, and is still subject to the competency trap.

Proposition 3 *An optimal knowledge system has at most one unit of capacity ($n^* \leq 1$). In each period t , at most the action with the highest payoff, $\tilde{a}_t \in B_t \cup C_t$, is recorded. In particular, there exists thresholds $\tilde{v}_1 > \tilde{v}_2 = px_K + (1-p)\mu$, 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.*

Proof : We will follow a backwards induction argument. Consider first the third period, and let \tilde{x}_3 denote the highest signal that the employee received. (Note that $\tilde{x}_3 \stackrel{a.s.}{=} x_K$.) The player's expected payoff from the corresponding action is $p\tilde{x}_3 + (1-p)\mu$. (Note that each player gets a signal about each action, so there are no actions with expected payoff μ , unless obviously the action's signal also happens to be $x = \mu$.) It follows that in period 3 the employee will choose the recorded action with the highest payoff if that payoff is larger than $p\tilde{x}_3 + (1-p)\mu$, else he chooses the action with the \tilde{x}_3 signal. Consider now the end of period 2. Almost surely (in a statistical sense), the highest signal that the period 3 employee will receive is x_K . It follows that if $v_{\tilde{a}_2} > px_K + (1-p)\mu$ then the firm records \tilde{a}_2 , else the firm does not record any action (since the firm prefers not to record an action when otherwise indifferent and, with $v_{\tilde{a}_2} \leq px_K + (1-p)\mu$, the effect of recording an action is almost surely at most that the period 3 employee will avoid that action, which has no effect on payoffs since there is an infinite number of potential actions). Consider now the employee's action choice in the second period. Let \tilde{x}_2 denote the highest signal that the employee received. The player's expected payoff from the corresponding action is $p\tilde{x}_2 + (1-p)\mu$. It follows that in period 2 the employee will choose the recorded action with the highest payoff if that payoff is larger than $p\tilde{x}_2 + (1-p)\mu$, else he chooses the action with the \tilde{x}_2 signal. Consider now the end of period 1. Almost surely (in a statistical sense), the highest signal that the period 2 employee will receive is x_K . Recording (from period 1) any action with expected payoff smaller than $px_K + (1-p)\mu$ almost surely has at most the effect that the action will be avoided. But with an infinite number of potential actions, that does not affect the payoffs. It follows that the firm will never record an action with a payoff smaller than $px_K + (1-p)\mu$. Assume thus that $v_{\tilde{a}_1} > px_K + (1-p)\mu$. Recording this action implies that every employee will undertake that action in period 2. By the above argument, the firm will record this action again at the end of

period 2, and all employees will choose it again in period 3. The payoff (for periods 2 and 3) from doing so is thus $2Iv_{\check{a}_1}$. The expected payoff from not recording this action, on the other hand, is independent of $v_{\check{a}_1}$ and strictly larger than $2I(px_K + (1-p)\mu)$. It follows that there is a critical value $\tilde{v}_1 > px_K + (1-p)\mu$ such that the firm records $v_{\check{a}_1}$ if $v_{\check{a}_1} > \tilde{v}_1$ and does not record anything otherwise. This proves the proposition. ■

Independent Draws For the formal model that relaxes the assumption that the realized payoffs of the different actions are independent draws from a non-degenerate distribution G , consider the following relatively simple variation on the main model of Section 2. Let there be one employee in each period (so that $I = 1$), no discounting ($\delta = 1$), and let the set of available actions be ordered in an infinite series $\{a_j\}_{j=-\infty}^{\infty}$. Assume now that the actions' payoffs have a multivariate normal distribution with means normalized to zero and with a covariance matrix such that $\sigma_{j,k} = V$ if $j = k$, $\sigma_{j,k} = C \geq 0$ if $k \in \{j-1, j+1\}$ and $\sigma_{j,k} = 0$ otherwise, where $C < V$ by the Cauchy-Schwartz inequality and the fact that a random variable cannot be perfectly correlated simultaneously with two independent random variables. This means that two adjacent values have a bivariate normal distribution while non-adjacent values are independent.^a To fix things and simplify the analysis, we will also assume that the firm can only record observed actions (i.e., it cannot record inferred payoffs unless they were also observed).

The following proposition shows that for this case, the firm will still record only successes, will record at most one such success, and is still subject to the competency trap.

Proposition 4 *An optimal knowledge system has at most one unit of capacity ($n^* \leq 1$). In each period t , at most the action with the highest payoff, $\check{a}_t \in B_t \cup C_t$, is recorded. In particular, there exists thresholds $\tilde{v}_1 > \tilde{v}_2 = 0$, 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.

Proof : Note first that, when $m \in \{k-1, k+1\}$ and with $\rho = C/V < 1$:

$$\begin{aligned} E[v_{a_m} | v_{a_k}] &= \rho v_{a_k} & \in & \begin{cases} (0, v_{a_k}) & \text{if } v_{a_k} > 0 \\ \{0\} & \text{if } v_{a_k} = 0 \\ (v_{a_k}, 0) & \text{if } v_{a_k} < 0 \end{cases} \\ \text{Var}[v_{a_m} | v_{a_k}] &= (1 - \rho^2)V & \in & (0, V) \end{aligned}$$

Let f_N denote the normal density with mean 0 and variance V . Let $f_{N|v}$ denote the normal density with mean and variance conditional on having observed an adjacent value v .

For the rest of the proof, we will follow a backwards induction argument. Consider first the third period. Note that if the firm recorded an action, then that also gives information about the payoffs of the two adjacent actions, but the expected payoff of these adjacent actions is either strictly smaller than that of the recorded action (if the latter is strictly positive) or weakly smaller than 0 (if the latter is weakly smaller than 0). It follows that the employee will choose the recorded action with the highest payoff if that payoff is strictly larger than 0, else he randomly chooses an action among

^aOne setting where that would be the case is when there exists some series of independently and normally distributed variables $\{x_k\}_{k=-\infty}^{\infty}$ and $v_k = x_k + x_{k+1}$. Clearly, any two adjacent v_k will be bivariate normal, but two v_k that are more than one spot apart will indeed be independent.

those with expected payoff 0 (i.e., from the set of unrecorded actions plus any actions with a recorded or inferred payoff that happens to be exactly 0). Consider now the end of period 2. If $v_{\tilde{a}_2} > 0$ then the firm records \tilde{a}_2 , else the firm does not record any action (since the firm prefers not to record an action when otherwise indifferent and, with $v_{\tilde{a}_2} \leq 0$, the effect of recording an action is at most that the period 3 employee will avoid that action, which has no effect on payoffs since there is an infinite number of potential actions with expected payoff equal to 0). Consider now the employee's action choice in the second period. Following an analogous argument as for period 3, the period 2 employee will choose the recorded action with the highest payoff if that payoff is strictly larger than 0, else he chooses at random from those actions about which no payoff information is known (directly or inferred from others).

Consider now the end of period 1. Again following an analogous reasoning as before, the firm will never record an action if $v_{\tilde{a}_1} \leq 0$. Assume thus now that $v_{\tilde{a}_1} > 0$. The expected continuation payoff from recording \tilde{a}_1 when $v_{\tilde{a}_1} > 0$ equals $2v_{\tilde{a}_1}$. The expected continuation payoff from recording no action is $0 + \int_0^\infty u f_N(u) du$. Note now that at $v_{\tilde{a}_1} = 0$, it is strictly optimal to not record any action. Furthermore, the payoff from not recording any action is independent of $v_{\tilde{a}_1}$ and finite, while the payoff from recording \tilde{a}_1 is strictly increasing in $v_{\tilde{a}_1}$ and goes to infinity as $v_{\tilde{a}_1} \rightarrow \infty$. It follows that there exists a $\tilde{v}_1 > 0$ such that \tilde{a}_1 will be recorded iff $v_{\tilde{a}_1} > \tilde{v}_1$. This implies the proposition. ■

Finite Action Sets and the Role of Recording Failures For the numerical example with a finite action set, consider the setting of Section 2 with 2 employees ($I = 2$), a finite number of actions that is at least 5 ($5 < N < \infty$), and no discounting ($\delta = 1$). Since our interest is in the optimal know-how recording policies, we will further assume that the system capacity is never a constraint, i.e., $n \geq 4$. Let the payoff distribution G be a discrete distribution that puts equal weight on 4 possible payoffs: $\{0, 49, 51, 100\}$. Finally, since the purpose is just to illustrate the possible effects, we will make an assumption that is a tad unrealistic but dramatically simplifies the analysis: assume that the firm has unlimited recall independent of what is recorded in the knowledge system.^b

For the purposes of this analysis, what matters are the recording policies at the end of the first period. In the unique equilibrium, these policies are as follows. If any of the first period actions pays off 100, then that action (or one of these actions) is recorded as the (unique) best practice and will be undertaken by all employees going forward. If no first period action pays off 100, then any action that paid off 0 or 49 (i.e., any failure in the sense of below-average payoffs) will be recorded and will be avoided henceforth by all employees. Actions with payoff 51 (i.e., a mediocre success in the sense of a slightly above-average payoff) are explicitly not recorded.

This gives indeed the following implications. First of all, failures *do* get recorded. In fact, when some players have above-average performance and others have below-average performance, the firm may record only the below-average payoffs. Second, since the benefits from remembering failures are additive, there is essentially no limit on the number of failures that get recorded (though remembering failures and remembering success are mutually exclusive in this particular model). Third, remembering failures does not reduce the incentives to experiment, so that the reasons for not recording failures are less strong than for not recording mediocre successes.

With regard to other settings where remembering failures may be valuable, consider the following context. 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

^bIn other words, the firm remembers (and can record at a later time) the payoffs of the first period actions even if they weren't recorded in the knowledge system at the time. The purpose is to eliminate the need for the firm to make inferences about what it did not record earlier.

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 argued above that the qualitative results of the paper seem to extend to that case.

The most important 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. One example would be a variation on the setting discussed in the above section on ‘No Prior Information’ but now with the signal being common to all players (or public information). This type of setting is an important avenue for future research.

B.2 More Detailed Discussion of Limitations and Extensions

We now turn to a more detailed discussion of the issues mentioned in Section 6.

Competition One important extension is incorporating the role of competition. In particular, the analysis in this paper studied a firm in isolation, which is a reasonable approximation for, for example, know-how about internal operations. In many other settings, however, it matters how the results might be affected by the interaction with competition. Section 5 already argued that when an action’s payoff depends on competitors’ actions, then learning becomes more difficult and the value of a knowledge system is reduced. As a consequence, we are more likely to see knowledge systems for internal operations or settings where such interactions are absent. This analysis assumed that competitive changes are exogenously given. More research is necessary to see whether the results also hold with endogenous competitive interactions. But there are also some other interesting angles.

One important effect is that competition affects the importance of experimentation versus exploration. In particular, when a firm in isolation stops experimenting after recording a mediocre success, it loses the potential gains from discovering an even better success. When that firm is in competition and its competitors do keep experimenting, then its lack of experimentation may be doubly expensive since this may put it in a competitive disadvantage in the future. Closely related is the observation that the competency trap that was derived in this paper may also provide one explanation why incumbents have less incentives to experiment than challengers. In particular, an incumbent often has better current know-how than a challenger. But since its current performance is also the opportunity cost of experimenting, experimentation is more expensive for the incumbent than for the challenger. As a consequence, the challenger is more likely to experiment than the incumbent and thus more likely to make the “next big discovery”. Note that this may still be true even when the incumbent is actually better at experimenting.

A second potential effect of competition is that there may be spillovers from learning. This reduces the benefits from experimentation and search in two ways. First of all, it may be possible to learn from others without incurring the learning expense that goes with trial and error. Second, any competitive advantage that comes from such learning may vanish quickly when competitors imitate. Note that this does not reduce the benefits of the knowledge system since it is still important to maximize performance by applying state-of-the-art knowledge. But it would probably affect the arguments for not recording information about mediocre successes. One potential conjecture is that firms in industries with heavy spillovers would be less critical about what information to record.

Finally, competition may also affect how operational effectiveness (and thus optimal knowledge management) translates to financial performance. The general intuition is that competition generates

more pressure to perform, which would increase the value of a knowledge system. Under cut-throat competition, however, the cost of experimentation may directly threaten a firm's survival, reducing the incentives to experiment. These issues clearly require more formal analysis.

Overall, extending the current analysis to include the potential effects of competition seems to be a very promising avenue.

Absence of Explicit Incentive Schemes The model assumed that employees simply try to maximize their own performance. While this probably approximates the typical situation in most companies quite well, it does obviously raise the question what the effect would be of allowing the firm to offer explicit incentives to experiment. At first sight, it seems that the possibility of explicit experimentation incentives may potentially eliminate the result that it may be strictly optimal not to record mediocre actions. We believe, however, that the result will typically persist, even in the face of such possible incentives, for three complementary reasons.

The first of these reasons is that the results we obtained are actually first-best or close to first-best, even without using incentives to experiment. In particular, when the firm can choose who gets access to which knowledge, the results that are obtained by this strategic use of the knowledge system are completely first best. One way to see this is that essentially nothing is lost by (explicitly) not communicating mediocre payoff actions: the firm does not want its employees to undertake such actions anyway. It thus follows that adding incentives to experiment would have a limited benefit.

This closely interacts with the second reason: experimentation is only one aspect of most employees' performance, and incentives to experiment often run in conflict with other objectives, such as giving incentives for effort. In particular, Manso [2006] shows that incentives for experimentation are often low-powered in the short term or even reward failure, which is very different from incentives for effort or for decisions under certainty. Using the knowledge system to take care of optimal experimentation relaxes the constraints on the incentive system and thus makes it easier to get it right from these other perspectives.

The third part is that incentives to experiment are often difficult or costly to implement, for two reasons. First, such incentives typically require considerable ability to commit since some of the necessary actions to elicit experimentation go ex-post against the firm's best interest. Not punishing failure, for example, is an important factor in eliciting experimentation. But failure often indicates low ability or lack of motivation so that it may be optimal ex-post to replace the employee, and thus difficult to commit ex-ante not to. (Note that even academia does not give tenure for trying, only for succeeding.) A second issue is that the effect of the decision is often difficult to isolate or measure, and thus difficult to contract on. In particular, while we assumed for simplicity that the firm observes all outcomes of the employees' actions, that is often not the case. Instead, firms often have to rely on self-reports by employees who have a clearer picture on the effect of a decision (which would be an interesting avenue for further research). Using such reports as the basis for experimentation incentives would likely lead to reporting distortions that hurt the performance of the knowledge system. And even when the firm observes the outcomes, incentive contracts require more: they require that the outcomes are contractible either explicitly or implicitly. Both issues suggest that incentives for experimentation may be very expensive to implement.

Overall, these arguments suggest two things. First, it may be possible to derive the assumption that employees maximize their own short-term performance endogenously as the outcome of an optimal incentive contract. Second, even when the inclusion of experimentation incentives would be optimal, it seems that they will never completely replace this strategic use of the knowledge system.

Knowledge System Interaction with Other Systems One very promising area for future research is the interaction of the knowledge system with other systems in the firm. We have discussed some possible interactions with the incentive system. Another interesting dimension of incentives are employees' incentives to share their best practices. If turnover from period to period is less than 100%, and an employee is evaluated relative to her peers, she 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 and also without exaggeration. The potential informational rents may increase the cost of the knowledge system, thereby affecting the optimal strategy for information collection and optimal knowledge system capacity.

We also conjecture that there is a nice complementarity between a knowledge system and job-rotation: as a firm has more job-rotation, the value of a knowledge system increases (both to transfer know-how to these employees who are new to the job and because there will be more experimentation); as the firm has a knowledge system, the value of job-rotation increases (because the cost decreases as it is easy to transfer know-how and because the experimentation that follows from such job-rotation can now be captured in the knowledge system).^c

Similar arguments could be made for hiring and turnover. For example, a high rate of turnover or an inexperienced hiring pool both favor the implementation of a knowledge system. In the other direction, the availability of a knowledge system reduces the gains from limiting turnover and from hiring more experienced people.

One potentially very important effect is the interaction with the type of projects that the firm selects. In particular, a firm may focus more on problems where the knowledge system generates a competitive advantage. A consulting firm with a strong knowledge system may focus, for example, on non-routine projects of medium complexity.

All these issues, however, require further formal analysis in a model that can capture overlapping generations. But this is definitely a very interesting avenue for future research.

C Numerical Example Illustrating Proposition 5

Figure C.1 illustrates the quasiconcavity of the marginal benefit of capacity, $\Delta\Pi(n)$, in q and in p . It shows the 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. 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 .

^cObviously, job-rotation is often done for different reasons (such as giving people a broad range of experiences, being able to evaluate an employee in different settings, or keeping them interested in their job) and these know-how effects are quite independent of that.

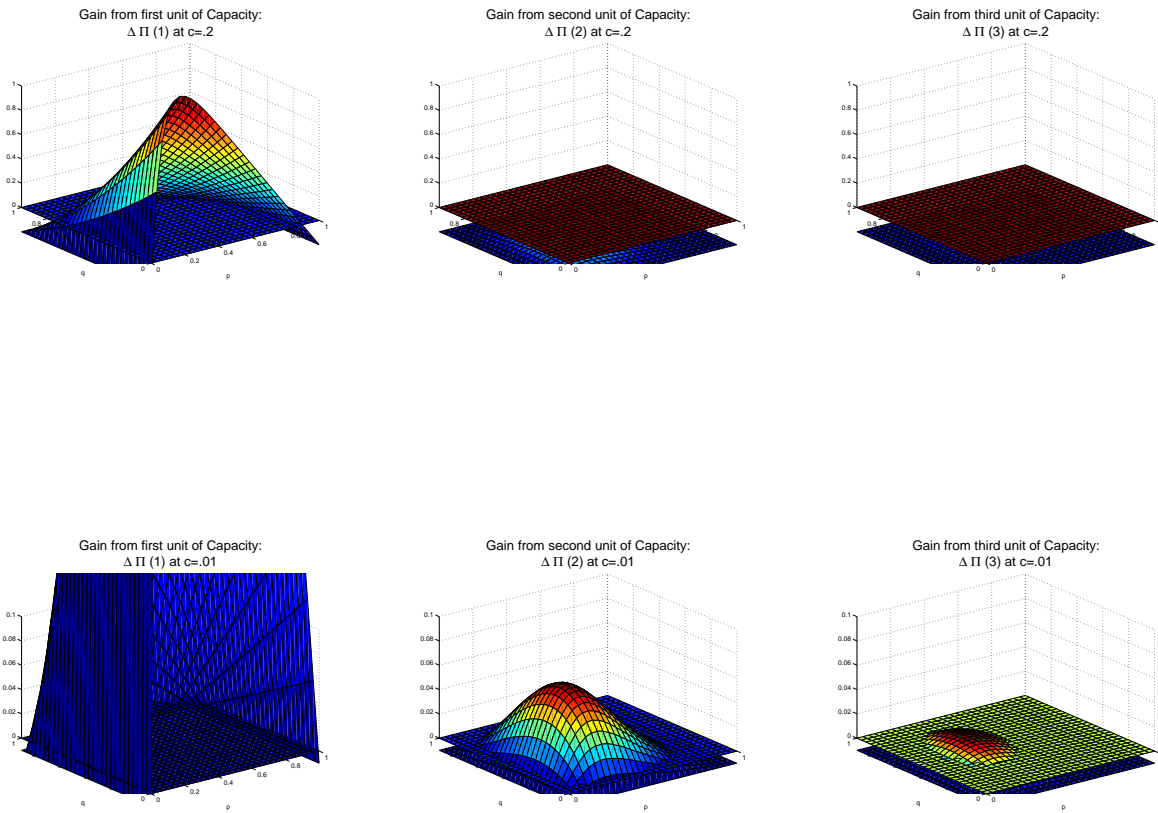


Figure C.1: Payoff gains per employee from the first 3 successive units of capacity with $I = 3$, and respectively $c = 0.2$ (upper graphs) and $c = 0.01$ (lower graphs).

References

- G. Manso. Motivating innovation. Working Paper, 2006.
- P. Milgrom and J. Roberts. The economics of modern manufacturing : Technology, strategy, and organization. *American Economic Review*, 80(3):511– 528, June 1990.
- P. Milgrom and J. Roberts. Comparing equilibria. *American Economic Review*, 84(3):441– 459, June 1994.