

# Balancing Act: How to Capture Knowledge Without Killing It

BY JOHN SEELY BROWN  
AND PAUL DUGUID

Top-down processes designed to institutionalize new ideas can have a chilling effect on creativity. But they don't have to. Managers can learn to walk the fine line between rigidity—which smothers creativity—and chaos—where creativity runs amok and nothing ever gets to market.



ARTWORK BY KATHY OSBORN

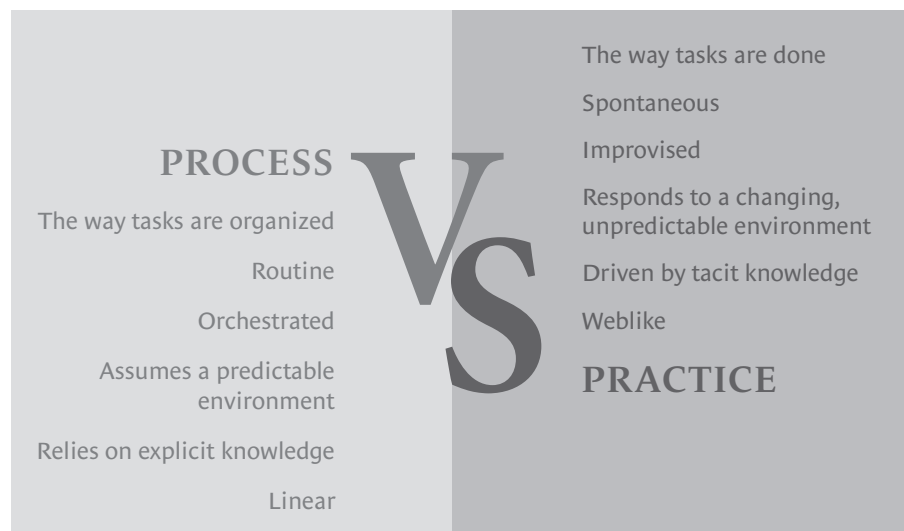
**H**ISTORY WILL PITY THE MANAGERS of the 1990s. The Internet touched down in their midst like a tornado, tearing up the old game book, disrupting every aspect of business, and compelling them to manage for a new economy. When managers sought help, they found the experts were offering two radically different theories about what such management should look like. The first approach – reengineering – focused on process. Organizations that reengineered their business processes would gain sustainable competitive advantage, according to an army of highly paid consultants. Major corporations spent millions of dollars and man-hours trying to do exactly that. But just as scores of reengineering VPs took their seats at *Fortune* 500 companies, word came down that process was stale. The new new thing was knowledge management – businesses that could capture the knowledge embedded in their organizations would own the future.

Reengineering and knowledge management are profoundly different approaches – as all those businesspeople who got whiplash from the turnaround soon realized. Reengineering is about the structured coordination of people and information. It's top-down. It assumes that it's easy to codify value creation. And it assumes that organizations compete in a predictable environment. Knowledge management focuses on effectiveness more than efficiency. It's bottom up. It assumes that managers can best foster knowledge by responding to the inventive, improvisational ways people actually get

**John Seely Brown** is chief scientist of Xerox and director of Xerox's Palo Alto Research Center (PARC).

**Paul Duguid** is a historian and social theorist affiliated with the University of California at Berkeley and Xerox PARC.

*This article is adapted from their book The Social Life of Information (Harvard Business School Press, March 2000).*



things done. It assumes that value-creating activities are not easy to pin down. And it assumes that organizations compete in an unpredictable environment.

Of course, management fads shift all the time. (How else could consultants stay in business?) But we think this shift from process engineering to knowledge management represents something more substantial than a change of fashion. It suggests a dilemma that all managers grapple with: the organizational tension between process, the way matters are formally organized, and practice, the way things actually get done.

Managers find this tension difficult to handle. They're paid to resolve or overcome tensions, but this is one they have to live with. Successful companies are not those that work around the problem; they are those that turn it to their advantage. For in the delicate art of balancing practice and process lies the means both to *foster* invention – by allowing new ideas to spark – and to *further* it – by implementing those same ideas. (See the exhibit, "Process Versus Practice.")

It's undoubtedly a hard balancing act. Lean too much toward practice, and you may get new ideas bubbling up all over the place, but you'll lack the structure to harness them. (And in the modern business world, worthwhile ideas that you don't harness end up in your competitors' hands.) Lean too much toward process, and you get lots of structure but too little

freedom of movement to strike that initial spark. Finding the right balance is a central task for managers everywhere. It's embodied in a million business fads, and it transcends them all.

It is possible to strike the right balance. In this article, we'll look closely at an example drawn from a company we know well because we work there. This is the story of how Xerox Corporation learned to foster best practice among a particular group of employees and then to circulate their expertise using the organizational support that process can provide.

### The Limits of Process-Based Thinking

One way managers attempt to resolve the tension between process and practice is by compartmentalizing. They do everything possible to foster invention and creativity among highly paid, elite workers (designers and scientists, for example). At the same time, they try to make everyone else's work completely predictable and to hold the majority of workers tight within the clamps of process. As a result, searches for underutilized knowledge round up the usual suspects – the output of the obviously inventive – and ignore everyone else, whose work practices are thought of as purely routine.

But this compartmentalization doesn't reflect the way most businesses currently operate. Today even the people involved in seemingly

routine work practices have to be inventive because the world they're working in changes so quickly. Their routines are always a little out of kilter. They must improvise to make up the difference between the conditions their routines were designed for and the actual conditions thrown up by a mutable world.

Consider an ordinary business form. Even the most recently printed (or Web-posted) form usually has boxes that are no longer used, categories that no longer apply. These redundant boxes are signposts of change. Employees quickly devise ways to fix the slightly out-of-date process. "Oh, leave that box," they'll tell customers, "but make sure to check under 'c.'" That will ring a bell in the marketing department, and they'll take care of you."

This particular example is insignificant. But such conversations—which happen all the time—are evidence of practical inventiveness used to get around the limits of process. These

course, and they also show us a lot about the tension between process and practice.

### Knowing What You Know

Identifying a company's best practices is not easy, for a couple of reasons. First, there's a large gap between what a task looks like in a process manual and what it looks like in reality. Second, there's a gap between what people think they do and what they really do. Actual work practices are full of tacit improvisations that the employees who carry them out would have trouble articulating. The manager who wishes to understand the company's best practices must bridge both of those gaps.

To illustrate the difficulty of identifying best practices, we'll look at the customer service representatives who fix Xerox machines. From the process perspective, a rep's work can be described quickly. Customers having difficulty call the Customer Service Center. The center, in turn, no-

a doubt. But they succeed primarily by departing from formal processes; those processes followed to the letter would soon bring their work (and their clients' work) to a halt.

For example, the company's documented repair processes assume that machines work predictably. Yet large machines, made up of multiple subsystems, are not so predictable. Each reflects the age and condition of its parts, the particular way it's used, and the environment in which it sits, which may be hot, cold, damp, dry, clean, dusty, secluded, in traffic, or otherwise. Any single machine may have profound idiosyncrasies. Reps know the machines they work with, Orr suggests, as shepherds know their sheep. While everyone else assumes one machine is like the next, a rep knows each by its peculiarities and sorts out general failings from particular ones.

Consequently, although the documentation gives the reps a map, the critical question for them is what to do when they fall off the map—which they do all the time. Orr found a simple answer to that question. When the path leads off the map, the reps go...to breakfast.

**When the Going Gets Tough.** Orr began his account of the reps' day not where the process view begins—at nine o'clock, when the first call comes in—but at breakfast beforehand, where the reps share and even generate new insights into these difficult machines. Orr found that a quick breakfast can be worth hours of training. While eating, playing cribbage, and gossiping, the reps talked work, and talked it continually. They posed questions, raised problems, offered solutions, constructed answers, laughed at mistakes, and discussed changes in their work, the machines, and customer relations. Both directly and indirectly, they kept one another up to date about what they knew, what they'd learned, and what they were doing.

The reps' group breakfast shows that work goes on that formal processes don't capture. But it shows more. It demonstrates that a job that seems highly independent on paper is in reality remarkably so-

## What happens when the reps fall off the map? There's a simple answer to that question. When the path leads off the map, the reps go ... to breakfast.

small fixes may be part of a company's best practice—where local inventiveness has enabled practices on the ground to outstrip processes on paper. All the small, individually insignificant best practices scattered around a company add up to an enormous amount of knowledge.

For a company to make the most of that knowledge—to "know what it knows," in the famous phrase of former Hewlett-Packard CEO Lew Platt—it needs to take practice, practitioners, and the communities that practitioners form seriously. That requires two steps. First, managers need to learn what local knowledge exists. Then if the knowledge looks valuable, they need to put it into wider circulation. Let's take those tasks one at a time. They lie at the heart of knowledge management, of

ties a rep. He or she then goes to the customer's site. With the help of error codes, which report the machine's state, and documentation, which says what those codes mean, the rep diagnoses the problem and follows instructions for fixing it. Practice here would seem to involve little more than following the map you are given and doing whatever it tells you to do.

It would seem that way, if someone hadn't bothered to look more closely. Julian Orr, formerly an anthropologist at Xerox's Palo Alto Research Center (PARC), studied what reps actually did, not what they were assumed to do. And what they actually did turned out to be quite different from the process we've just described. The reps' work is organized by business processes, without

cial. Reps get together not only at the parts drop and the customer service center but also on their own time for breakfast, at lunch, for coffee, or at the end of the day – and sometimes at all of those times. This sociability is not just a retreat from the loneliness of an isolating job. The constant chatting is similar to the background updating that goes on all the time in any ordinary work site.

There, too, chatting usually passes unnoticed unless someone objects to it as a waste of time. But it's not. Orr showed that the reps use one another as their most critical resources. In the course of socializing, the reps develop a collective pool of practical knowledge that any one of them can

whom “what” and “why” are critical but often hard matters to discern.

We all tell stories this way. Economists tell stories in their models. Scientists tell stories in their experiments. Executives tell stories in their business plans (see “Strategic Stories: How 3M Is Rewriting Business Planning,” HBR May–June 1998). Storytelling helps us discover something new about the world. It allows us to pass that discovery on to others. And finally, it helps the people who share the story develop a common outlook. Orr found that war stories give the reps a shared framework for interpretation that allows them to collaborate even though the formal processes assume they are working independently.

## Often what one person thinks useful others find flaky, idiosyncratic, incoherent, redundant, or just plain stupid.

draw upon. That pool transcends any individual member's knowledge, and it certainly transcends the corporation's documentation. Each rep contributes to the pool, drawing from his or her own particular strengths, which the others recognize and rely on. Collectively, the local groups constitute a community of practice. (For a detailed description, see “Communities of Practice: The Organizational Frontier,” HBR January–February 2000.)

**Storytelling.** Much of the knowledge that exists within working groups like the one formed by our Xerox reps comes from their war stories. The constant storytelling about problems and solutions, about disasters and triumphs over breakfast, lunch, and coffee serves a number of overlapping purposes. Stories are good at presenting things sequentially (this happened, then that). Stories also present things causally (this happened because of that). Thus stories are a powerful way to understand what happened (the sequence of events) and why (the causes and effects of those events). Storytelling is particularly useful for the reps, for

**Improvisation.** Not all of the reps' problems can be solved over breakfast or by storytelling alone. Experimentation and improvisation are essential, too. One day, Orr observed a rep working with a particularly difficult machine. It had been installed recently, but it had never worked satisfactorily. Each time it failed, it produced a different error message. Following the established process for each particular message – replacing or adjusting parts – didn't fix the overall problem. And collectively the messages made no sense.

Having reached his limits, the rep summoned a specialist. The specialist could not understand what was going on, either. So the two spent the afternoon cycling the machine again and again, waiting for its intermittent crashes and recording its state when it did. At the same time, they cycled stories about similar-looking problems round and round until they, too, crashed up against this particular machine. The afternoon resembled a series of alternating improvisational jazz solos, as each man took the lead, ran with it for a little while, then handed it off to the other,

this all against the bass-line continuum of the rumbling machine.

In the course of this practice, the two gradually brought their separate ideas closer together toward a shared understanding of the machine. Eventually, late in the day, everything clicked. The machine's erratic behavior, the experience of the two technicians, and the stories they told finally formed a single, coherent account. They made sense of the machine and worked out how to fix it. And the solution quickly became part of the community lore, passed around for others in their group to use if they encountered the same problem.

As Orr's study shows, executives who want to identify and foster best practices must pay very close attention to the practices as they occur in reality rather than as they are represented in documentation or process designs. Otherwise, they will miss the tacit knowledge produced in improvisation, shared through storytelling, and embedded in the communities that form around those activities. Does that mean process has no importance in this context? Of course not. But the processes that support how people work should be deeply informed by how they already work – not imposed from above by process designers who imagine they understand the work better than they actually do. Armed with a sense of what really happens on the ground, it's possible to design processes that prompt improvisation rather than ones that are blindly prescriptive.

### Spreading What You Know

People working in small groups develop very rich knowledge in practice, as we've seen. Assuming a company has correctly identified those practices and the tacit knowledge embedded within them, the question becomes, How can we spread that useful knowledge around? This is the point at which process becomes useful. Process – in the form of organizational coordination – can get that local knowledge into wider circulation.

Let's return to the Xerox reps. The group Orr studied included about a dozen people; the rep force worldwide currently numbers some

25,000. Locally generated fixes and insights circulated pretty efficiently within the small group but rarely made it beyond. So people in different groups spent time grappling with problems that had already been solved elsewhere. The reps as a whole still didn't know what some reps, as a group, knew.

The far-flung communities that made up the entire network of reps needed some organizational support to help them share local knowledge around the world. So Xerox initiated the Eureka project to oversee the knowledge dissemination. The project set out to create a database to preserve resourceful ideas over time and deliver them over space.

Do we hear a yawn? Databases are the most basic of knowledge management tools. They're also among the most ignored. Organizations fill their databases with useful tips and data, and nobody uses them. Why should another be any different? The answer in this case is that it's different because of how the data are judged to be useful.

Most such databases, like most business processes, are top-down creations. Managers fill them with what they think will be useful for the people they manage. And—surprise—surprise—the people usually don't find them so. Yet even when individuals fill databases with their own ideas of what's useful, they aren't much help either. Often what one person thinks useful others find flaky, idiosyncratic, incoherent, redundant, or just plain stupid. The more a database contains everyone's favorite idea, the more unusable it becomes.

The Eureka database was designed to get past that problem by establishing a process to help capture best practices. Reps, not the organization, supply the tips. But reps also vet the tips. A rep submits a suggestion first to a local expert on the topic. Together, they refine the tip. It's then submitted to a centralized review process, organized according to business units. Here reps and engineers again vet the tips, accepting

some, rejecting others, eliminating duplicates, and calling in experts on the particular product line to resolve doubts and disputes. If a tip survives this process, it becomes available to reps around the world, who have access to the tips database over the Web. So reps using the system know that the tips—and the database as a whole—are relevant, reliable, and probably not redundant.


It's interesting to compare this method of circulating knowledge with the established practices and formal processes of the scientific community. The two methods are quite similar. Scientists, too, work in small, local groups. To circulate their ideas more widely, they also put those ideas through a well-established process of peer review. If accepted, the ideas are then published for others to see.

Most scientists don't get paid for scientific articles. Good articles do, however, earn them status among their peers. They become known and respected for careful work, reliable results, and important insights. The reps have followed a similar course. The corporation offered to pay for the tips, but the pilot group of reps who helped design the system thought that would be a mistake, worrying, among other things, that payment for submissions would lead people to focus on quantity rather than quality in making submissions. Instead, the reps chose to have their names attached to tips. Those who submit good tips earn positive recognition. Because even good tips vary in quality, reps, like scientists, build social capital through the quality of their input. At a recent meeting of Xerox reps in Canada, one individual was surprised by a spontaneous standing ovation from coworkers who were expressing their respect for his tips. Of course, as in the scientific community, such recognition may also lead to career advancement. But it is important not to underestimate the value of social and intellectual capital within workplace communities—particularly

those not usually recognized for their knowledge production.

The current Eureka database holds about 30,000 records. And its value is growing as it grows. In one case, an engineer in Brazil was about to replace a problematic high-end color machine (at a cost of about \$40,000) for a disgruntled customer. Experimenting with a prototype of Eureka, he found a tip from a Montreal technician that led him to replace a defective 50 cent fuse instead. In all, Eureka is estimated to have saved the corporation \$100 million.

Process and practice, then, do not represent rival views of the organization. Rather, they reflect the creative tension at the center of innovative organizations. In this, organizations resemble the well-known picture that, looked at once, appears to show a vase, but looked at once again, turns into two people, face to face. The vase resembles well-defined and precisely structured process—easy to understand though hard to change. The faces reflect practice—always unfolding in unpredictable ways, full of promise and problems, just like a conversation. The manager's challenge is to keep both images in view simultaneously.

So, to come back to where we began, the swing from business process reengineering to knowledge management did represent a radical shift in focus. But the goal for managers is not to choose between the two. Rather, the goal is to find the right balance between them—one that can grow only more important as knowledge becomes the factor that distinguishes the successful companies from the failures. Indeed, as dot-com companies mature, they're starting to search for seasoned managers who can provide their inventive, explosive communities of practice with the structure of process—but who won't suffocate practice while they're at it. 

Reprint R00309

To place an order, call 1-800-988-0886.