

The Social Life of Information

John Seely Brown and Paul Duguid

HARVARD BUSINESS SCHOOL PRESS
BOSTON, MASSACHUSETTS

2000

INTRODUCTION

Tunneling Ahead

LIVING IN THE INFORMATION AGE can occasionally feel like being driven by someone with tunnel vision. This unfortunate disability cuts off the peripheral visual field, allowing sufferers to see where they want to go, but little besides. For drivers, this full attention to the straight ahead may seem impressively direct. For passengers who can see what the driver does not it makes for a worrisome journey. The car may skirt the mirrors of passing cars, threaten the hands and ankles of nearby cyclists, roll close to the toes of pedestrians, occasionally scrape walls and gateposts, and once in a while take a serious sideswipe that jeopardizes the journey.

Similarly, some of the people driving us all hard into the future on the back of new technologies appear to assume that if we all focus hard enough on information, then we will get where we want to go most directly. This central focus inevitably pushes aside all the fuzzy stuff that lies around the edges—context, background, history, common knowledge, social resources. But this stuff around the edges is not as irrelevant as it may seem. It provides valuable balance and perspective. It holds alternatives,

offers breadth of vision, and indicates choices. It helps clarify purpose and support meaning. Indeed, ultimately it is only with the help of what lies beyond it that any sense can be made of the information that absorbs so much attention.

Think of the way people interact. Talk may deliver information—something that can be recorded, transcribed, digitized, and shipped in packets. But as you talk, listeners set what you say in a much larger context. Your appearance, your age, your accent, your background, and the setting all contribute to what they understand. Few people are like Sherlock Holmes, able to spot an ex-Indian Army soldier now working as a billiard marker across a street. But we are all remarkably good at picking up clues and cues that underwrite or undermine what a speaker says. Con artists have to work for their living.

Except, perhaps, in the digital world. It is no surprise, really, that cyberspace has become famous for “identity experiments” and con games.¹ The world of information is often so thin, the cues and clues so few, that in many cases it’s easy to pose, even as an ex-Indian Army soldier now working as a billiard marker, and get away with it. In the tight restrictions of the information channel, without the corroboration that broader context offers (or refuses), the powerful detective skills that everyone relies on have little room to work.²

TUNNEL DESIGN

Ignoring the clues that lie beyond information doesn’t only lead to a narrow world of deception. It leads to a world of what we think of as *tunnel design*—a kind of purblind design of which, in the end, we are all victims. In this world we are often expected to live on a strict information-only diet. Indeed, it’s a world that usually

addresses worries about information by simply offering more. Yet when only information is on offer, more often means less.

Tunnel design, we suspect, produces technologies that, in Edward Tenner’s phrase, “bite back.”³ These are technologies that create as many problems as they solve. Many of the unintended consequences of design that Tenner describes and that make new technologies so frustrating arise from neglecting resources that lie outside the tight focus of information.

If badly designed technologies bite back, the good ones often fight back. Well-designed technologies, that is, refuse to retreat meekly in the face of tunnel design. Futurology is littered with the obituaries of tools that nonetheless continue a robust and healthy life. One of our own favorite examples is the hinge. This seems to be written out of every futuristic movie in favor of the sliding door; yet, it not only hangs on but is vital to many laptops and cell phones.⁴ Even the typewriter still finds a place in many “infomated” offices.⁵ These are trivial examples. But obituaries are now regularly written not just for tools, but for well-established practices, organizations, and institutions, not all of which seem to be fading away.

The dogged endurance of those that refuse to fade tends to produce opposing responses. Those with tunnel vision condemn the foolishness of humanity for clinging to the past. Those exasperated by tunnel design tend to cheer the downfall of new technology as if it were never likely to come to any good. There is, we believe, a better, more productive approach. It is for the new to learn from the old. Tools fight back when they offer people worthwhile resources that may be lost if they are swept away. So rather than condemning humanity as foolish, primitive, or stubborn for sticking with the old and rejecting the new, it seems better to stop and ask why.

The answer often turns out to be that the new alternatives have underestimated their targets. Paradoxically, tunnel design often takes aim at the surface of life. There it undoubtedly scores lots of worthwhile hits. But such successes can make designers blind to the difficulty of more serious challenges—primarily, to the resourcefulness that helps embed certain ways of doing things deep in our lives. Generations of confident videophones, conferencing tools, and technologies for tele-presence are still far from capturing the essence of a firm handshake or a straight look in the eye.

To say this is not to plead for a return to the buggy cart, the steam engine, or the vinyl record. It is to plead for attention—attention to stubbornness, to what will not budge, to the things that people fight for. So it's to plead for design that takes into account resources that people care about. Such design, we are confident, produces tools that people care about—a kind of tool that seems, despite modern inventiveness, in remarkably short supply. (Take a quick look over the computer applications you have bought, borrowed, or downloaded over the past five years and see how many you would actually fight for.)

Issues about the breadth or narrowness of design are not, we should stress, issues for designers alone. Increasingly, we all live in a heavily designed world and need to understand both the strengths and the limitations of the designs offered to us. In particular, we all need to be able to deal with the hype that accompanies new technological designs. In the digital world, moreover, many of the distinctions between designers and users are becoming blurred. We are all, to some extent, designers now. Many questions about design are thus becoming questions for us all. It is important, then, to understand our own limitations as designers, too, and to know where to look for resources.

FIGHT TO THE FINISH

So where might the resources that we believe are overlooked lie? Despite our metaphor of tunnel vision, our sense of the neglected periphery is not limited to the visual periphery of physical objects.⁶ It also embraces what we think of as the *social periphery*, the communities, organizations, and institutions that frame human activities. These, though vital to how we all live and work, are too often missing from the design stylebooks of the information age. Or, if they are included, they come as targets, not resources for design. It is to help draw attention to these hard-to-see (and hard-to-describe) resources that we gave our book the title it has. Attending too closely to information overlooks the social context that helps people understand what that information might mean and why it matters.

Well-honed resources for living tend to fall out of sight, lost to the immediate demands of daily life. New technologies, by contrast, draw and deserve a lot of attention. Consequently, it's easy to overlook the resources that lie beyond the immediacy of the information tunnel, even when they are quite substantial.

For example, a colleague at the University of California, Berkeley, recently sang the praises of the digital world to us. He can now, he told us, through the university's digital library, get direct access to information. He felt, he said, in touch with information as never before.

His enthusiasm was justifiable. New technologies have changed the life of researchers tremendously. Nonetheless, his enthusiasm had screened out an enormous array of people, organizations, and institutions involved in this "direct" touch. The university, the library, publishers, editors, referees, authors, the computer and infrastructure designers, the cataloguers and

library collections managers, right down to the students making their way through college by shelving and unshelving books, copying and scanning, and delivering files across the campus had no place in his story. When they do their job well, they do it more or less invisibly. But without them, there would have been no story.

Similarly, those who talk about having direct, unmediated access to the news sometimes sound equally oblivious to how news is made. They sound as if to find the “real” news on Russia, for example, they would expect to pick up the phone and get Boris Yeltsin on the other end of the line. But it requires a profoundly naïve belief in “disintermediation” (see chapter 1) to assume that all the links that fall between Yeltsin and the news are somehow interference in the information channel. Rather, it is in these steps—from sources to reporters to editors and news organizations—that news is made. Without them, again, there would be no story. Nonetheless, when information takes center stage and lights dim on the periphery, it’s easy to forget these necessary intermediaries. But while these may be invisible, they are not inconsequential.

Occasionally, the naïve neglect of the world beyond information reminds us of an old Popeye cartoon. In this, if we remember right, Popeye finds Sweet Pea, the baby, crawling along the ledge outside the bedroom, oblivious to the danger below. While Popeye frantically slips, slides, and dangles way above the street trying to reach the baby, Sweet Pea crawls giggling to the end of the ledge and off. But one of those swinging girders that obligingly come by in the cartoon world leads her on undisturbed. Popeye meanwhile, much disturbed, falls down the stairs, trips on the sidewalk, and collides with lampposts trying to keep up. From the edge of the swinging girder,

Sweet Pea next crawls onto the roof of a conveniently passing truck. The truck nearly runs Popeye over. And so the chase goes on until Sweet Pea manages to end happily back in the bedroom, all the while oblivious to the extraordinary resources that have provided support and protection. In Sweet Pea’s mind, Sweet Pea did it all. Popeye, meanwhile, has met the sharp end of each of these resources.

Some digital champions can appear a little like Sweet Pea, oblivious to the resources of the world that support them. Others, more worryingly, are actively hostile. They see society’s resources as constraints on information and yearn to be free of them. Material objects appear as the unwanted antithesis to bits, communities and organizations as enemies to individualism, institutions as “second wave.”

It’s not hard to understand this antagonism. For the most part people notice these supports primarily when something goes wrong. When everything works well, we are all a little like Sweet Pea, most ignorant of what makes us most secure. When things go wrong, we are more like Popeye, feeling that the stairs betrayed us, the lamppost attacked us, and that the world would be a better place without such things.

We hope that in the course of this book we can steer a path between blindness and bruises. We want to draw attention to the resources people use in the belief that what are resources for people are, by extension, resources for the design of useful tools. Tools that ignore these resources will be, in great part, tools that are at best ignored, at worst a burden on those who have to use them.

To make this case, we look at plans, prognostications, and designs that have been much applauded in theory yet significantly ignored in practice. We include prognostications about, for example, the world of information, digital agents, the home

office, the paperless society, the virtual firm, and the digital university. From here we try to explain why so many confident predictions remain just that, predictions. Too often, we conclude, the light at the end of an information tunnel is merely the gleam in a visionary's eye. The way forward is paradoxically to look not ahead, but to look around.

CONVERSATIONS

The essays that follow attempt to address the social context of information from a variety of different perspectives. They have grown out of long-running conversations that the two of us have conducted in public and in private over the last dozen years or so—conversations in which we found that, though the topic shifted, the way in which we addressed it had common threads.

These threads allow us to talk from a single standpoint about the limits of infopunditry (chapter 1), the challenges of software agents (chapter 2), the social character of work and learning—and the limits of management theory (chapters 3–5), resources for innovation (chapter 6), unnoticed aspects of the document and their implications for design more generally (chapter 7), and the future of institutions, in particular the university (chapter 8). Despite what we see as the common threads, each chapter, we hope, can be read on its own.

These essays, like the conversations they reflect, try to avoid the blinkered euphoria of the infoenthusiast. Though we are enthusiasts for technologies, we focus predominantly elsewhere—on society, practice, and institutions. We look this way because we know that while information and its technologies cannot solve all society's problems, society and social resources can solve many of

the problems of both information and technology. But it is society and social resources to which many designs are blinkered, if not blind.

Though we point to social resources for solutions, this is, we must admit, more a book of questions than answers. This may not seem a very constructive approach. We are, after all, writing in a genre where people often seem to identify the major problems of the world and then announce that the solution is running on the machine of a colleague in the office next door. We are enormously proud to be associated with our colleagues at the Xerox Palo Alto Research Center (PARC) and of the work they do. And we do point to their work quite often. In general, however, we have tried to avoid short-circuiting the loop between problem and solution. Our aim, though associated as we are with Xerox we are a little hesitant to say it, has been to keep the conversation going and even change its direction a little.

C H A P T E R **I**

Limits to Information

On an average weekday the New York Times contains more information than any contemporary of Shakespeare's would have acquired in a lifetime.

—ANONYMOUS (and ubiquitous)

Every year, better methods are being devised to quantify information and distill it into quadrillions of atomistic packets of data.

—BILL GATES

By 2047 . . . all information about physical objects, including humans, buildings, processes and organizations, will be online. This is both desirable and inevitable.

—GORDON BELL AND JIM GRAY

This is the datafication of shared knowledge.

—TOM PHILLIPS, Deja News¹

IT NOW SEEMS a curiously innocent time, though not that long ago, when the lack of information appeared to be one of society's fundamental problems. Theorists talked about humanity's "bounded rationality" and the difficulty of making decisions in conditions of limited or imperfect information. Chronic information shortages threatened work, education, research, innovation, and economic decision making—whether at the level of government policy, business strategy, or household shopping. The one thing we all apparently needed was more information.

So it's not surprising that infoenthusiasts exult in the simple volume of information that technology now makes available. They count the bits, bytes, and packets enthusiastically. They cheer the disaggregation of knowledge into data (and provide a new word—*datafication*—to describe it). As the lumps break down and the bits pile up, words like *quadrillion*, *terabyte*, and *megaflop* have become the measure of value.

Despite the cheers, however, for many people famine has quickly turned to glut. Concern about access to information has given way to concern about coping with the amounts to which we do have access. The Internet is rightly championed as a major information resource. Yet a little time in the nether regions of the Web can make you feel like the SETI researchers at the University of California, Berkeley, searching through an unstoppable flood of meaningless information from outer space for signs of intelligent life.²

With the information spigot barely turned on—the effect has seemed more like breaching a dam than turning a tap—controlling the flow has quickly become the critical issue. Where once there seemed too little to swim in, now it's hard to stay afloat. The "third wave" has rapidly grown into a tsunami.³

Faced by cheery enthusiasts, many less optimistic people resemble the poor swimmer in Stevie Smith's poem, lamenting that

*I was much too far out all my life
And not waving, but drowning.*

Yet still raw information by the quadrillion seems to fascinate.

COULD LESS BE MORE?

Of course, it's easy to get foolishly romantic about the pleasures of the "simpler" times. Few people really want to abandon information technology. Hours spent in a bank line, when the ATM in the supermarket can do the job in seconds, have little charm. Lose your papers in a less-developed country and trudge, as locals must do all the time, from line to line, from form to form, from office to office and you quickly realize that life without information technology, like life without modern sanitation, may seem simpler and even more "authentic," but for those who have to live it, it is not necessarily easier or more pleasant.

Even those people who continue to resist computers, faxes, e-mail, personal digital assistants, let alone the Internet and the World Wide Web, can hardly avoid taking advantage of the embedded microchips and invisible processors that make phones easier to use, cars safer to drive, appliances more reliable, utilities more predictable, toys and games more enjoyable, and the trains run on time. Though any of these technologies can undoubtedly be infuriating, most people who complain want improvements, not to go back to life without them.⁴

Nonetheless, there is little reason for complacency. Information technology has been wonderfully successful in many ways.

But those successes have extended its ambition without necessarily broadening its outlook. Information is still the tool for all tasks. Consequently, living and working in the midst of information resources like the Internet and the World Wide Web can resemble watching a firefighter attempt to extinguish a fire with napalm. If your Web page is hard to understand, link to another. If a "help" system gets overburdened, add a "help on using help." If your answer isn't here, then click on through another 1,000 pages. Problems with information? Add more.

Life at Xerox has made us sensitive to this sort of trap. As the old flip cards that provided instructions on copiers became increasingly difficult to navigate, it was once suggested that a second set be added to explain the first set. No doubt, had this happened, there would have been a third a few years later, then a fourth, and soon a whole laundry line of cards explaining other cards.

The power and speed of information technology can make this trap both hard to see and hard to escape. When information burdens start to loom, many of the standard responses fall into a category we call "Moore's Law" solutions. The law, an important one, is named after Gordon Moore, one of the founders of the chip maker Intel. He predicted that the computer power available on a chip would approximately double every eighteen months. This law has held up for the past decade and looks like it will continue to do so for the next.⁵ (It's this law that can make it hard to buy a computer. Whenever you buy, you always know that within eighteen months the same capabilities will be available at half the price.)

But while the law is insightful, Moore's Law solutions are usually less so. They take it on faith that more power will somehow solve the very problems that they have helped to create. Time alone, such solutions seem to say, with the inevitable cycles of the

Law, will solve the problem. More information, better processing, improved data mining, faster connections, wider bandwidth, stronger cryptography—these are the answers. Instead of thinking hard, we are encouraged simply to "embrace dumb power."⁶

More power may be helpful. To the same degree, it is likely to be more problematic, too. So as information technology tunnels deeper and deeper into everyday life, it's time to think not simply in terms of the next quadrillion packets or the next megaflop of processing power, but to look instead to things that lie beyond information.

DROWNING AND DIDN'T KNOW IT

If, as one of our opening quotations suggests, "all information about physical objects, including humans, buildings, processes and organizations, will be online," it's sometimes hard to fathom what there is beyond information to talk about.

Let us begin by taking a cue from MIT's Nicholas Negroponte. His handbook for the information age, *Being Digital*, encouraged everyone to think about the differences between atoms, a fundamental unit of matter, and bits, the fundamental unit of information.⁷ Here was a provocative and useful thought experiment in contrasts. Moreover, it can be useful to consider possible similarities between the two as well.

Consider, for example, the industrial revolution, the information revolution's role model. It was a period in which society learned how to process, sort, rearrange, recombine, and transport atoms in unprecedented fashion. Yet people didn't complain that they were drowning in atoms. They didn't worry about "atom overload." Because, of course, while the world may be composed of atoms, people don't perceive it that way. They

perceive it as buses and books and tables and chairs, buildings and coffee mugs, laptops and cell phones, and so forth. Similarly, while information may come to us in quadrillions of bits, we don't consider it that way. The information reflected in bits comes to us, for example, as stories, documents, diagrams, pictures, or narratives, as knowledge and meaning, and in communities, organizations, and institutions.⁸

The difficulty of looking to these various forms through which information has conventionally come to us, however, is that infocentric visions tend to dismiss them as irrelevant. Infoenthusiasts insist, for example, not only that information technology will see the end of documents, break narratives into hypertext, and reduce knowledge to data, but that such things as organizations and institutions are little more than relics of a discredited old regime.

Indeed, the rise of the information age has brought about a good deal of "endism." New technology is widely predicted to bring about, among other things,

the end of the press, television, and mass media

the end of brokers and other intermediaries

the end of firms, bureaucracies, and similar organizations

the end of universities

the end of politics

the end of government

the end of cities and regions

the end of the nation-state

There's no doubt that in all these categories particular institutions and particular organizations are under pressure and many will not survive long. There's nothing sacred here. But it's one thing to argue that many "second wave" tools, institutions, and organizations will not survive the onset of the "third wave." It's another to argue that in the "third wave" there is no need for social institutions and organizations at all.

The strong claim seems to be that in the new world individuals can hack it alone with only information by their side. Everyone will return to frontier life, living in the undifferentiated global village.⁹ Here such things as organizations and institutions are only in the way. Consequently, where we see solutions to information's burdens, others see only burdens on information.

ORIGIN MYTHS

From all the talk about electronic frontiers, global villages, and such things as electronic cottages, it's clear that the romanticism about the past we talked about earlier is not limited to technophobes.¹⁰ Villages and cottages, after all, are curious survivors from the old world applied to the conditions of the new. They remind us that the information age, highly rationalist though it seems, is easily trapped by its own myths. One of the most interesting may be its origin myth, which is a myth of separation.

Historians frequently trace the beginnings of the information age not to the Internet, the computer, or even the telephone, but to the telegraph. With the telegraph, the speed of information essentially separated itself from the speed of human travel. People traveled at the speed of the train. Information began to travel at close to the speed of light. In some versions of this origin story (which tends to forget that fire and smoke had long

been used to convey messages over a distance at the speed of light), information takes on not only a speed of its own, but a life of its own. (It is even capable, in some formulations, of “wanting” to be free.)¹¹ And some scholars contend that with the computer, this decisive separation entered a second phase. Information technologies became capable not simply of transmitting and storing information, but of producing information independent of human intervention.¹²

No one doubts the importance of Samuel Morse’s invention. But with the all-but-death of the telegraph and the final laying to rest in 1999 of Morse code, it might be time to celebrate less speed and separation and more the ways information and society intertwine. Similarly, it’s important not to overlook the significance of information’s power to breed upon itself. But it might be time to retreat from exuberance (or depression) at the volume of information and to consider its value more carefully.¹³ The ends of information, after all, are human ends. The logic of information must ultimately be the logic of humanity. For all information’s independence and extent, it is people, in their communities, organizations, and institutions, who ultimately decide what it all means and why it matters.

Yet it can be easy for a logic of information to push aside the more practical logic of humanity. For example, by focusing on a logic of information, it was easy for *Business Week* in 1975 to predict that the “paperless office” was close. Five years later, one futurist was firmly insisting that “making paper copies of anything” was “primitive.”¹⁴ Yet printers and copiers were running faster and faster for longer and longer periods over the following decade. Moreover, in the middle of the decade, the fax rose to become an essential paper-based piece of office equipment. Inevitably, this too was seen as a breach of good taste. Another

analyst snorted that the merely useful fax “is a serious blemish on the information landscape, a step backward, whose ramifications will be felt for a long time.”¹⁵

But the fax holds on. Rather like the pencil—whose departure was predicted in 1938 by the *New York Times* in the face of ever more sophisticated typewriters—the fax, the copier, and paper documents refuse to be dismissed.¹⁶ People find them useful. Paper, as we argue in chapter 7, has wonderful properties—properties that lie beyond information, helping people work, communicate, and think together.

If only a logic of information, rather than the logic of humanity, is taken into account, then all these other aspects remain invisible. And futurists, while raging against the illogic of humankind and the primitive preferences that lead it astray, will continue to tell us where we ought to go. By taking more account of people and a little less of information, they might instead tell us where we are going, which would be more difficult but also more helpful.

HAMMERING INFORMATION

Caught in the headlights of infologic, it occasionally feels as though we have met the man with the proverbial hammer to whom everything looks like a nail. If you have a problem, define it in terms of information and you have the answer. This redefining is a critical strategy not only for futurology, but also for design. In particular, it allows people to slip quickly from questions to answers.

If indeed Morse did launch the information age, he at least had the modesty to do it with a famously open-ended question. “What,” he piously asked in his first message, “hath God

wrought?" Now, "we have answers," or "solutions" or "all the answers you need" (11,000 according to Oracle's Web site). Similarly, IBM claims that a single computer can contain "answers to all the questions you ever had."¹⁷ So if Morse were to ask his question again today, he would no doubt be offered an answer beginning "http://www. . . ."

True, Microsoft advertises itself with a question: "Where do you want to go today?" But that is itself a revealing question. It suggests that Microsoft has the answers. Further, Microsoft's pictures of people sitting eagerly at computers also suggest that whatever the question, the answer lies in digital, computer-ready information. For though it asks where you want to go, Microsoft isn't offering to take you anywhere. (The question, after all, would be quite different if Microsoft's Washington neighbor Boeing had asked it.) Atoms are not expected to move, only bits. No doubt to the regret of the airlines, the ad curiously redefines "go" as "stay." Stay where you are, it suggests, and technology will bring virtually anything you want to you in the comfort of your own home. (Bill Gates himself intriguingly refers to the computer as a "passport.")¹⁸ Information offers to satisfy your wanderlust without the need to wander from the keyboard.¹⁹

REFINING, OR MERELY REDEFINING?

In the end, Microsoft's view of your wants is plausible so long as whatever you do and whatever you want translates into information—and whatever gets left behind doesn't matter. From this viewpoint, value lies in information, which technology can refine away from the raw and uninteresting husk of the physical world.

Thus you don't need to look far these days to find much that is familiar in the world redefined as information. Books are

portrayed as information containers, libraries as information warehouses, universities as information providers, and learning as information absorption. Organizations are depicted as information coordinators, meetings as information consolidators, talk as information exchange, markets as information-driven stimulus and response.

This desire to see things in information's light no doubt drives what we think of as "infoprefixation." *Info* gives new life to a lot of old words in compounds such as *infotainment*, *infomatics*, *infomating*, and *infomediary*. It also gives new promise to a lot of new companies, from InfoAmerica to InfoUSA, hoping to indicate that their business is information. Adding *info* or something similar to your name doesn't simply add to but multiplies your market value.

Undoubtedly, information is critical to every part of life. Nevertheless, some of the attempts to squeeze everything into an information perspective recall the work of the Greek mythological bandit Procrustes. He stretched travelers who were too short and cut off the legs of those who were too long until all fitted his bed. And we suspect that the stretching and cutting done to meet the requirements of the infobed distorts much that is critically human. Can it really be useful, after all, to address people as information processors or to redefine complex human issues such as trust as "simply information?"²⁰

6-D VISION

Overreliance on information leads to what we think of as "6-D vision." Unfortunately, this is not necessarily twice as good as the ordinary 3-D kind. Indeed, in many cases it is not as good, relying as it does on a one-dimensional, infocentric view.

The *D* in our 6-*D* notion stands for the *de-* or *dis-* in such futurist-favored words as

demassification

decentralization

denationalization

despacialization

disintermediation

*disaggregation*²¹

These are said to represent forces that, unleashed by information technology, will break society down into its fundamental constituents, principally individuals and information. (As we scan the *Ds*, it sometimes feels as though the only things that will hold up against this irresistible decomposition are the futurists' increasingly long words.)

We should say at once that none of these *D*-visions is inherently mistaken or uninteresting. Each provides a powerful lens on an increasingly complicated world. They help expose and explain important trends and pressures in society. Nonetheless, the *Ds* too easily suggest a linear direction to society—parallel movements from complex to simple, from group to individual, from personal knowledge to ubiquitous information, or more generally from composite to unit.

Yet it does not feel that modern life is moving in one direction, particularly in the direction from complex to simple. To most of us, society seems capable of moving in almost any direction, and often in the direction of chaos rather than simplicity. Indeed, many shifts that the 6-*Ds* reveal are not the first step in

an unresisting downward spiral from complex to simple. Rather, they are parts of profound and often dramatic shifts in society's dynamic equilibrium, taking society from one kind of complex arrangement to another, as a quick review of a few *Ds* will suggest.

Dimensions of the Ds

Much talk about disaggregation and demassification readily assumes that the new economy will be a place of ever-smaller firms, light, agile, and unencumbered. It was once commonplace, for example, to compare the old Goliath, GM, against the new David, Microsoft. As Microsoft's market capitalization passed GM's, the latter had some 600,000 employees and the former barely 25,000. The difference is stark. Not, though, stark enough to step from here to what the business writers Larry Downes and Chunka Mui call the "Law of Diminishing Firms." After all, it's GM that's shrinking. Microsoft continues to grow while other high-tech start-ups compete for the title of "fastest growing ever."²²

Downes and Mui draw on the theory of the firm proposed by the Nobel Prize-winning economist Ronald Coase. Coase developed the notion of *transaction costs*. These are the costs of using the marketplace, of searching, evaluating, contracting, and enforcing. When it is cheaper to do these as an organization than as an individual, organizations will form. Conversely, as transaction costs fall, this glue dissolves and firms and organizations break apart. Ultimately, the theory suggests, if transaction costs become low enough, there will be no formal organizations, but only individuals in market relations. And, Downes and Mui argue, information technology is relentlessly driving down these costs.

Though he produced elegant economic theory, Coase had strong empirical leanings. He developed his theory of transaction costs in the 1930s to bridge the gap between theoretical accounts of the marketplace and what he saw in the actual marketplace—particularly when he traveled in the United States. There, business was dominated by huge and still-growing firms. These defied the purity and simplicity of the theoretical predictions, which envisaged markets comprising primarily individual entrepreneurs.²³

In honor of Coase's empiricism, it's important to look around now. When we began work on this book, Justice Department lawyers opened their case against Microsoft, accusing it of monopolistic practices. David now resembles Goliath. At the same time, other Justice Department lawyers were testifying that 1998 would be the first two-trillion-dollar year for mergers. Seven of the ten largest mergers in history had occurred in the first six months alone. We began keeping a list of firms involved. These included Amoco, AT&T, Bankers Trust, BMW, British Petroleum, Chrysler, Citibank, Deutsche Bank, Exxon, Ford, IBM, MCI, Mercedes, Mobil, Travelers, and many more.

Nor were these large firms buying up minnows. They were buying up each other. Ninety years after the era of trust busting, oil, banking, and tobacco, the initial targets, were all consolidating again.²⁴ As the *Economist* put it, after Exxon's merger with Mobil followed British Petroleum's purchase of Amoco: "Big Oil is Dead. Long Live Enormous Oil."²⁵

Whatever else was apparent, we soon realized that whenever the book came out, any list of ours would be profoundly out of date. The only successful strategy in such conditions would be to imitate the great comic novelist of the eighteenth century, Laurence Sterne, who faced with an impossible description inserted a blank page into his manuscript and told the readers to

take up their own pen and do it for themselves. As we were revising the manuscript, the two behemoths of the information age, AT&T and Microsoft, began their own extraordinary mating dance. That we found well beyond the reach of our pens.

Undoubtedly, several of the mergers we mentioned may represent the last struggles of dinosaurs to protect their ecological niche before other forces destroy it. Investment and even retail banking, for example, may have particularly precarious futures.

But massification is not occurring in dying "second wave" sectors alone. Many mergers have involved firms based in the "third wave" information sectors. Here mergers often involve not so much dinosaurs as phoenixes rising from the ashes of old business models. These might include AT&T's absorption of TCI and Time-Warner's of Turner Broadcasting. They surely do include Internet-driven combinations such as MCI's merger with WorldCom, IBM's takeover of Lotus, and AT&T's purchase of IBM's Global Network. Meanwhile, firms wholly within the new economy, such as AOL, Microsoft, Amazon, and eBay, go on regular shopping sprees for other companies.

Elsewhere in the information sector, Sir John Daniel, vice-chancellor of Britain's Open University, points to the rise of the "mega-university." Daniel presides over some 160,000 students, but his school hardly qualifies as "mega" in a field in which the largest—China's TV University System—has 580,000 students in degree programs. According to Daniel's figures, two universities break the half-million mark, one exceeds one-third of a million, and three are approaching a quarter million.²⁶ These are all "distance" universities, using a variety of information technologies to reach their students. So no simple demassification here either. Similarly, the concentration of the media in recent years challenges any simple idea of media demassification.²⁷

It doesn't feel then as if firms are shrinking under an iron law. Rather, it feels more as if, as the economist Paul Krugman puts it, "We've gone from an economy where most people worked in manufacturing—in fairly large companies that were producing manufactured goods and engaged in things like transportation—to an economy where most people work for fairly large companies producing services."²⁸

The resilience of the large organization is not all that surprising. Given that information technologies are particularly good at taking advantage of large networks, the information economy in certain circumstances actually favors the aggregated, massified firm.²⁹ These are firms that can or have knit diverse networks together, as AOL hopes to do with its purchase of Netscape or as Microsoft hopes to do with the insertion of Windows into television set-top boxes. Consequently, the small, agile firm with big ideas and little money is less likely to be the viable start-up of legend. (As a recent article in *Red Herring* put it, referring to the famous garage-based start-ups of Silicon Valley, the "garage door slams shut.")³⁰ And any that do start up in the traditional way are likely to be snatched up by the giants of the industry.

So, while stories abound about the new "niche" markets exploited through the Internet, the examples often come not from niche firms, but from large ones with well-established networks. The paradoxical phrase "mass customizing" suggests that fortune favors the latter. It is possible, for example, to have jeans cut to your personal dimensions. But it is quite probably Levi's that will do it for you. Here the strategy for customized goods relies on a large firm with a large market and a highly standardized product. So the demassification of production relies on the

massification of markets and consumption. The Henry Ford of the new economy would tell us that we can all have jeans made to measure, so long as they are Levi's.

Finally, firms are not merely taking power from one another. They are accumulating power that once lay elsewhere. The political scientist Saskia Sassen traces the decline of the nation-state not to the sweeping effects of demassification and disaggregation, but to the rise of powerful, concentrated transnational corporations. The new economic citizen of the world, in her view, is not the individual in the global village but the transnational corporation, often so formidable that it has "power over individual governments."³¹ The state and the firm, then, are not falling together along a single trajectory. At least in some areas, one is rising at the other's expense.

In sum, as people try to plot the effects of technology, it's important to understand that information technologies represent powerful forces at work in society. These forces are also remarkably complex. Consequently, while some sectors show disaggregation and demassification, others show the opposite. On the evidence of the 6-Ds, attempts to explain outcomes in terms of information alone miss the way these forces combine and conflict.

So while it might seem reasonable to propose a law of increasing, not diminishing, firms, that too would be a mistake. It would merely replace one linear argument with another. It's not so much the actual direction that worries us about infocentrism and the 6-Ds as the assumption of a single direction. The landscape is more complex. Infocentricity represents it as disarmingly simple. The direction of organizational change is especially hard to discern. The 6-Ds present it as a foregone conclusion.

More Dimensions

Similarly, despite talk of disintermediation and decentralization, the forces involved are less predictable and unidirectional than a quick glance might suggest.³² First, the evidence for disintermediation is far from clear. Organizations, as we shall see, are not necessarily becoming flatter. And second, where it does occur, disintermediation doesn't necessarily do away with intermediaries. Often it merely puts intermediation into fewer hands with a larger grasp. The struggle to be one of those few explains several of the takeovers that we mentioned above. It also explains the "browser wars" between Netscape and Microsoft, the courtship of AT&T and Microsoft, and the continuing struggle for dominance between Internet Service Providers (ISPs). Each of these examples points not to the dwindling significance but to the continuing importance of mediation on the 'Net (as does the new term *infomediary*, another case of infoprefixation). Moreover this kind of limited disintermediation often leads to a *centralization* of control. These two Ds, then, are often pulling not together, but against one another.

NOT FLATTER. Francis Fukuyama and Abram Shulsky conducted a RAND study in 1997 into the relationship between disintermediation, flat organizations, and centralization on behalf of the army. They began by studying the private sector. Here they give little hope for any direct link between information technology and flatter organizations. Indeed, like us, they believe that the conventional argument that information technology (IT) will lead to flatter organizations is an infocentric one

[that] focuses on a single, if very important, function of middle management: the aggregation, filtering, and transmission of information.

*It is of course precisely with respect to this function that the advances in IT suggest that flattening is desirable, since IT facilitates the automation of much of this work. On the other hand, middle management serves other functions as well.*³³

If managers are primarily information processors, then information-processing equipment might replace them, and organizations will be flatter. If, on the other hand, there is more to management than information processing, then linear predictions about disintermediation within firms are too simple.

Empirical evidence suggests such predictions are indeed oversimplified. Despite the talk of increasingly flatter and leaner organizations, Paul Attewell, a workplace sociologist, argues that "administrative overhead, far from being curtailed by the introduction of office automation and subsequent information technologies, has increased steadily across a broad range of industries."³⁴ Attewell's data from the U.S. Bureau of Labor Statistics show that the growth of nonproduction employees in manufacturing and the growth of managerial employment as a percentage of the nation's workforce has risen steadily as the workplace has been infomated.

NOR MORE EGALITARIAN. Fukuyama and Shulsky also argue that in instances where information technology has led to disintermediation, this has not necessarily produced decentralization. "Despite talk about modern computer technology being necessarily democratizing," they argue, "a number of important productivity-enhancing applications of information technology over the past decade or two have involved highly centralized data systems that are successful because all their parts conform to a single architecture dictated from the top."³⁵ Among the successful examples they give are Wal-Mart and FedEx, both of which have famously centralized decision making.

These two are merely recent examples of a clear historical trend whereby information technology centralizes authority. Harold Innis, an early communications theorist, noted how the international telegraph and telephone lines linking European capitals to their overseas colonies radically reduced the independence of overseas administrators. Previously, messages took so long to travel that most decisions had to be made locally. With rapid communication, they could be centralized. Similarly, histories of transnational firms suggest that with the appearance of the telegraph, overseas partners, once both financially and executively autonomous, were quickly absorbed by the "home" office.³⁶

Less innocent than infoenthusiasts, commanders in the U.S. Navy understood the potential of information technology to disempower when they resisted the introduction of Marconi's ship-to-shore radio.³⁷ They realized that, once orders could be sent to them on-board ship, they would lose their independence of action. (Their resistance recalls a story of the famous British admiral Lord Nelson, who "turned a blind eye" to his telescope at the Battle of Copenhagen to avoid seeing his commander's signal to disengage.)³⁸

In contemplating assumptions about the decentralizing role of information technology, Shoshona Zuboff, a professor at Harvard Business School, confessed to becoming much more pessimistic in the decade since she wrote her pathbreaking book on the infomated workplace, *In the Age of the Smart Machine*: "The paradise of shared knowledge and a more egalitarian working environment," she notes, "just isn't happening. Knowledge isn't really shared because management doesn't want to share authority and power."³⁹

Of course this need not be the outcome. As Zuboff argues,

it's a problem of management, not technology.⁴⁰ Smaller organizations, less management, greater individual freedom, less centralization, more autonomy, better organization, and similar desirable goals—these arguments suggest—will not emerge spontaneously from information's abundance and the relentless power of the 6-Ds. Rather, that abundance is presenting us with new and complex problems that another few cycles of Moore's Law or "a few strokes of the keyboard" will not magically overcome.⁴¹ The tight focus on information, with the implicit assumption that if we look after information everything else will fall into place, is ultimately a sort of social and moral blindness.

THE MYTH OF INFORMATION

6-D vision, while giving a clear and compelling view of the influence of the 'Net and its effects on everything from the firm to the nation, achieves its clarity by oversimplifying the forces at work. First, it isolates information and the informational aspects of life and discounts all else. This makes it blind to other forces at work in society. Second, as our colleague Geoffrey Nunberg has argued, such predictions tend to take the most rapid point of change and to extrapolate from there into the future, without noticing other forces that may be regrouping.⁴²

This sort of reductive focus is a common feature of futurology. It accounts, for example, for all those confident predictions of the 1950s that by the turn of the century local and even domestic nuclear power stations would provide all the electricity needed at no cost. Not only did such predictions overlook some of the technological problems ahead, they also overlooked the social forces that confronted nuclear power with the rise of environmentalism. (Fifties futurism also managed to miss the

significance of feminism, civil rights, and student protest while continually pointing to the imminence of the videophone and the jet pack.)

We began this chapter with a brief look back to the industrial revolution. In many ways the train epitomized that earlier revolution. Its development was an extraordinary phenomenon, spreading from a 12-mile line in the 1830s to a network of nearly 25,000 miles in little more than a decade.⁴³ The railway captured the imagination not only of Britain, where it began, but of the world. Every society that could afford a railway, and some that couldn't, quickly built one. Standards were developed to allow for interconnections. Information brokers emerged to deal with the multiple systems involved.⁴⁴ The train also sparked an extraordinarily speculative bubble, with experienced and first-time investors putting millions of pounds and dollars into companies with literally no track record, no income, and little sign of profitability. Unsurprisingly, in popular imagination, both at the time and since, the train has presented itself as a driving force of social and economic revolution.

Economic and social historians have long argued, however, that the story of the industrial revolution cannot be told by looking at the train alone. Historians might as well whistle for all the effect they have had. The myth of the train is far more powerful.

Today, it's the myth of information that is overpowering richer explanations. To say this is not to belittle information and its technologies. These are making critical and unprecedented contributions to the changes society is experiencing. But it is clear that the causes of those changes include much more than information itself. So the myth significantly blinds society to the character of and forces behind those changes.

In particular, the myth tends to wage a continual war against aspects of society that play a critical role in shaping not only society, but information itself, making information useful and giving it value and meaning. It's hard to see what there is other than information when identity is reduced to "life on the screen," community thought of as the users of eBay.com, organization envisaged only as self-organization, and institutions merely demonized as "second wave."

We do not believe that society is relentlessly demassifying and disaggregating. Though we admit it would be much easier to understand if it were. The social forces that resist these decompositions, like them or not, are both robust and resourceful. They shaped the development of the railroad, determining where it ran, how it ran, and who ran it. And they will continue to shape the development of information networks. As we hope to show in the course of this book, to participate in that shaping and not merely to be shaped requires understanding such social organization, not just counting on (or counting up) information.⁴⁵

CHAPTER 5

Learning—in Theory and in Practice

Knowledge management is the use of technology to make information relevant and accessible wherever that information may reside. To do this effectively requires the appropriate application of the appropriate technology for the appropriate situation.

Knowledge management incorporates systematic processes of finding, selecting, organizing, and presenting information in a way that improves an employee's comprehension and use of business assets.¹

WE BEGAN THE LAST CHAPTER contemplating the trend from business process reengineering to knowledge management. There, we focused primarily on the limits of process, which we suggested was an infofriendly concept, but one that might be blind to other issues.

In this chapter, we take up the other half of the matter and consider knowledge and learning, again in relation to practice

and again as distinct from information. We do this with some trepidation. On the one hand, epistemology, the theory of knowledge, has formed the centerpiece of heavyweight philosophical arguments for millennia. On the other, knowledge management has many aspects of another lightweight fad. That enemy of lightweights, *The Economist*, has pronounced it no more than a buzzword. We may then, be trying to lift a gun too heavy to handle to aim at a target too insubstantial to matter.

Certainly much about knowledge's recent rise to prominence has the appearance of faddishness and evangelism. Look in much of the management literature of the late 1990s and you could easily believe that faltering business plans need only embrace knowledge to be saved. While it's often hard to tell what this embracing involves, buying more information technology seems a key indulgence.

Nonetheless, people are clearly taking up the idea of knowledge in one way or other. From within organizations come sounds of fighting between the IT (information technology) and HR (human resources) factions over who "owns" knowledge management.² Similarly, technology giants have entered a propaganda war over who best understands knowledge.³ Elsewhere, the management consultants are maneuvering for high ground in the knowledge stakes.

In the process, *knowledge* has gained sufficient momentum to push aside not only concepts like reengineering but also *information*, whose rule had previously looked so secure. To be, in Peter Drucker's term, a "knowledge worker" now seems much more respectable than being a mere "information worker," though for a while the latter seemed very much the thing to be. Similarly, pundits are pushing "information economy" and the venerable "information age" aside in the name of the more

voguish "knowledge economy" and "knowledge age." There's even a bit of alternative prefixation in such terms as *knobot*, which we talked about in chapter 2, where the buzz of bots and the buzz of knowledge meet.

Beyond its buzz, however, is there any bite to these uses of *knowledge*? When people talk about *knowledge*, are they just clinging to fashion (as many no doubt are), or might some be feeling their way, however intuitively, toward something that all the talk of information or of process lacks? Is there, we begin by asking, something that *knowledge* catches, but that *information* does not?

PERSONAL DISTINCTION

Twenty-five hundred years of unresolved epistemological debate from the Sophists to the present argue that we would be unwise to seek the difference by pouring over rigorous definitions. Moreover, whatever differences abstract definitions might clarify, persuasive redefinition now obscures. People are increasingly eager that their perfectly respectable cache of information be given the cachet of knowledge. Such redefinitions surreptitiously extend the overlapping area where *knowledge* and *information* appear as interchangeable terms.

Nevertheless, if we check the language of knowledge management at the door, there do appear to be some generally accepted distinctions between knowledge and information. Three strike us as particularly interesting.

First, knowledge usually entails a knower. That is, where people treat information as independent and more-or-less self-sufficient, they seem more inclined to associate knowledge with someone. In general, it sounds right to ask, "Where is that information?" but odd to ask, "Where's that knowledge?" as if

knowledge normally lay around waiting to be picked up. It seems more reasonable to ask, "Who knows that?"

Second, given this personal attachment, knowledge appears harder to detach than information. People treat information as a self-contained substance. It is something that people pick up, possess, pass around, put in a database, lose, find, write down, accumulate, count, compare, and so forth.⁴ Knowledge, by contrast, doesn't take as kindly to ideas of shipping, receiving, and quantification. It is hard to pick up and hard to transfer. You might expect, for example, someone to send you or point you to the information they have, but not to the knowledge they have.

Third, one reason knowledge may be so hard to give and receive is that knowledge seems to require more by way of assimilation. Knowledge is something we digest rather than merely hold. It entails the knower's understanding and some degree of commitment. Thus while one person often has conflicting information, he or she will not usually have conflicting knowledge. And while it seems quite reasonable to say, "I've got the information, but I don't understand it," it seems less reasonable to say, "I know, but I don't understand," or "I have the knowledge, but I can't see what it means." (Indeed, while conventional uses of *information* don't necessarily coincide with the specialist uses, as we noted earlier, "information theory" holds information to be independent of meaning.)⁵

WHERE IS THE KNOWER LOST IN THE INFORMATION?

Knowledge's personal attributes suggest that the shift toward knowledge may (or should) represent a shift toward people. Focusing on process, as we argued, draws attention away from people,

concentrating instead on disembodied processes and the information that drives them. Focusing on knowledge, by contrast, turns attention toward knowers. Increasingly, as the abundance of information overwhelms us all, we need not simply more information, but people to assimilate, understand, and make sense of it.

The markets of the knowledge economy suggest that this shift is already underway. Investment is no longer drawn, as postindustrial champions like to point out, to bricks and mortar and other forms of fixed capital. Nor does it pursue income streams. (In some of the newest knowledge organizations there is as yet barely enough income to puddle, let alone stream.) Instead, investors see value in people and their know-how—people with the ability to envisage and execute adventurous new business plans and to keep reenvisioning these to stay ahead of the competition.

So, while the modern world often appears increasingly impersonal, in those areas where knowledge really counts, people count more than ever. In this way, a true knowledge economy should distinguish itself not only from the industrial economy but also from an information economy. For though its champions like to present these two as distinct, the information economy, like the industrial economy, shows a marked indifference to people. The industrial economy, for example, treated them en masse as interchangeable parts—the factory "hands" of the nineteenth century. The information economy threatens to treat them as more or less interchangeable consumers and processors of information. Attending to knowledge, by contrast, returns attention to people, what they know, how they come to know it, and how they differ.

The importance of people as creators and carriers of knowledge is forcing organizations to realize that knowledge lies less in its databases than in its people. It's been said, for example, that if

NASA wanted to go to the moon again, it would have to start from scratch, having lost not the data, but the human expertise that took it there last time. Similarly, Tom Davenport and Larry Prusak argue that when Ford wanted to build on the success of the Taurus, the company found that the essence of that success had been lost with the loss of the people that created it. Their knowledge was not stored in information technologies. It left when they left.⁶

Mistaking knowledge and its sources for information and its sources can, then, be costly. In her book *Wellsprings of Knowledge*, Dorothy Leonard-Barton of Harvard Business School tells the story of one firm, ELP, taking over a rival, Grimes, primarily to capture Grimes's impressive intellectual capital. Only after it had paid generously for the business, machine tools, and research did ELP find that Grimes's real competitive advantage had lain in the operating knowledge of its line employees, all of whom had been let go.

Similarly, the sort of blind downsizing produced by business process reengineering has caused organizations to lose "collective memory."⁷ It's impossible to assess the value of such layoffs. But the business journalist Thomas Stewart estimated the cost of AT&T's last round as equivalent to an \$8 billion capital write-off.⁸ In all, the job of knowledge management cannot involve just the protection and exploitation of patents. It must include the cultivation of knowledgeable workers. Focusing on information, however, makes this kind of cultivation difficult.

KNOWN PROBLEMS

Curiously, if knowledge will go out of the door in the heads of people who have developed and worked with that knowledge, it seems reluctant to go out (or stay behind) in the heads of people

who have not been so involved. The CEO of the innovative steel manufacturer Chaparral Steel told Leonard-Barton that for this reason the firm has no problem with competitors touring their plant. Chaparral, he said, is willing to show just about everything "and we will be giving away nothing because they can't take it home with them."⁹ Unlike information, knowledge, as we said, is hard to detach.

While the challenge of detaching knowledge from some people and attaching it to others may protect some knowledge assets, it makes management of the knowledge much more difficult. The difficulty has revealed itself in, for example the struggle over "best practices." To maintain competitive edge, firms first search for the best practices either within their own or in their competitors' units. Once identified, they then transfer these to areas where practices are less good. The search part has led to a great deal of useful benchmarking. The transfer part, however, has proved much more awkward.

Robert Cole of the University of California at Berkeley's Haas Business School has investigated this difficulty in a recent study of best practice strategy. He looked at, among others, Hewlett-Packard's attempts to raise quality levels in its plants around the globe by identifying and circulating the best practices within the firm. Even internally, Cole showed, transfer was uncertain.¹⁰ Cole's findings seem to justify the now-famous lament of HP's chairman, Lew Platt, as he considered how much better the firm would be "if only we knew what we know at HP."¹¹

Although, as Cole emphasizes, HP works across continents and countries, failure to transfer practice is not simply a matter of national or linguistic boundaries. Best practices can have as much trouble traveling across town as they do across continents. As one winner of the prestigious Baldrige prize who grappled

with this problem told researchers in frustration, "We can have two plants right across the street from one another, and it's the damndest thing to get them to transfer best practices." Similarly, Jeff Papows, president of Lotus, whose Notes is a widely used tool for knowledge management, acknowledges that for all the power to communicate that Notes and similar groupware provide, "spreading the practice has not been easy."¹²

IN DEFENSE OF LEARNING

Circulating human knowledge, these experiences suggest, is not simply a matter of search and retrieval, as some views of knowledge management might have us believe. While knowledge is often not all that hard to search, it can be difficult to retrieve, if by *retrieve* people mean detach from one knower and attach to another.

So learning, the acquisition of knowledge, presents knowledge management with its central challenge. The defense of intellectual property, the sowing and harvesting of information, the exploitation of intellectual capital, and the benchmarking of competitors' intellectual assets are all important parts of the knowledge management game. But all of these are subordinate to the matter of learning. For it is learning that makes intellectual property, capital, and assets usable.

The difficulty of this central challenge, however, has been obscured by the redefinition that, as we noted earlier, infoenthusiasts tend to indulge. The definitions of knowledge management that began this chapter perform a familiar two-step. First, they define the core problem in terms of information, so that, second, they can put solutions in the province of information technology.¹³ Here, retrieval looks as easy as search.

If information retrieval were all that is required for such things as knowledge management or best practice, HP would have nothing to worry about. It has an abundance of very good information technology. The persistence of HP's problem, then, argues that knowledge management, knowledge, and learning involve more than information. In the rest of this chapter we try to understand what else is involved, looking primarily at knowledge and learning on the assumption that these need to be understood before knowledge management can be considered.

COMMUNITY SUPPORT

To understand learning in the context of knowledge management, let's begin by asking in what conditions do knowledge and best practice move. In chapter 4, we saw the reps sharing their knowledge, insights, and best practices quite effectively. These traveled first among the small group of coworkers and then, with the help of the Eureka database, across larger groups of company reps. To understand how these best practices travel, this example suggests, requires looking not simply from knowledge to information, but (as the idea of best practice might suggest) from knowledge to practice and groups of practitioners. For it is the reps' practice shared in collaborative communities that allowed them to share their knowledge.

As we saw, the reps formed themselves into a small community, united by their common practice, servicing machines. The members of this community spent a lot of time both working and talking over work together. In Orr's account, the talk and the work, the communication and the practice are inseparable. The talk made the work intelligible, and the work made the talk intelligible. As part of this common work-and-talk, creating,

learning, sharing, and using knowledge appear almost indivisible. Conversely, talk without the work, communication without practice is if not unintelligible, at least unusable. Become a member of a community, engage in its practices, and you can acquire and make use of its knowledge and information. Remain an outsider, and these will remain indigestible.

Two learning researchers, whose individual work we mentioned earlier, Jean Lave of the University of California, Berkeley, and Etienne Wenger, a consultant formerly of the Institute for Research on Learning, explain this sort of simultaneous working, learning, and communication in terms of both the practice and the community. Learning a practice, they argue, involves becoming a member of a "community of practice" and thereby understanding its work and its talk from the inside.¹⁴ Learning, from this point of view, is not simply a matter of acquiring information; it requires developing the disposition, demeanor, and outlook of the practitioners.

Like Orr's study of reps, Wenger's study of claims processing (see chapter 4) showed the importance of the group to both what people learn and how. Within the group, Wenger's study reveals, knowledge, traveling on the back of practice, was readily shared.

It may at first seem that group practice and community support are only appropriate for the tedium of "lowly" claims processing. They might seem to have little to do with the "higher" altitudes of knowledge work, where the image of the lone, Rodinesque "thinker" is more common. Yet the value of communities of practice to creating and sharing knowledge is as evident in the labs of particle physicists and biotechnologists as in the claims processing unit.¹⁵ The apprenticeship-like activity that Lave and Wenger describe is found not only on the shop floor, but throughout the highest reaches of education and

beyond. In the last years of graduate school or in internships, scientists, humanists, doctors, architects, or lawyers, after years of schoolroom training, learn their craft in the company of professional mentors. Here, they form learning communities capable of generating, sharing, and deploying highly esoteric knowledge.

Recently a computer engineer described a group that he led on a difficult project that, despite the difference in subject matter, resembles the groups of interdependent technicians and claims processors:

[It] was less than half a dozen people; and the group that did the software and hardware never did get to be more than about a dozen people. It was a tiny enough group that everyone knew everything that was going on, and there was very little structure . . . there were people who specifically saw their role as software, and they knew a lot about hardware anyway; and the hardware people all could program. There wasn't a great deal of internal difficulty. There's always a little, I don't think you can get even six people together without having some kind of a problem. . . . There was amazingly little argument or fighting.¹⁶

This description catches central properties of the community of practice. In particular, it notes how, in getting the job done, the people involved ignored divisions of rank and role to forge a single group around their shared task, with overlapping knowledge, relatively blurred boundaries, and a common working identity. The speaker in this case is Frank Heart of Bolt Beranek and Newman; the group's task, designing the core computers for what came to be the Internet.

In all, whether the task is deemed high or low, practice is an effective teacher and the community of practice an ideal learning environment.

TO BE OR NOT TO BE

Of course, whatever the strengths of communities of practice, people learn on their own, picking up information from numerous sources about numerous topics without ever becoming a “member.” We can learn something about Tibetan medicine or racing without needing to work with Tibetan doctors or become a Formula 1 driver. The critical words here, however, are *about* and *become*. They point to a distinction made by Jerome Bruner, a professor of psychology at New York University, between *learning about* and *learning to be*. Certainly, most of anyone’s knowledge might best be described as knowledge “about.” Many people learn *about* a lot of things—astrophysics, Australian Rules football, Madagascan lemurs, or baseball statistics. In the age of the Web, this learning about is easier than ever before.

But, picking up information about Madagascan lemurs in the comfort of our home doesn’t close the gap between us and Madagascan field zoologists. Learning to be requires more than just information. It requires the ability to engage in the practice in question.

Indeed, Bruner’s distinction highlights another, made by the philosopher Gilbert Ryle. He distinguishes “know that” from “know how.” Learning about involves the accumulation of “know that”: principally data, facts, or information. Learning about does not, however, produce the ability to put “know that” into use. This, Ryle argues, calls for “know how.” And “know how” does not come through accumulating information. (If it did, “know that” and “know how” would, in the end, be indistinguishable—build up enough “know that” and you would become a practitioner.) “We learn *how*,” Ryle argues, “by practice.” And, similarly, through practice, we learn to be.¹⁷

Ryle’s philosophical argument may have brought us dangerously near the realm of abstruse epistemology that we promised to avoid. But it helps explain why the same stream of information directed at different people doesn’t produce the same knowledge in each. If the people are engaged in different practices, if they are learning to be different kinds of people, then they will respond to the information in different ways. Practice shapes assimilation.

The practice of managing a baseball team, for example, is not the same as the practice of playing on a baseball team. The “know that” for each job may be fairly similar. Managers and players gather a lot of the same information. But the “know how” for the two (thus the way each makes use of their “know that”) is quite different. One’s practice is to manage; the other’s is to play. Similarly, while management theorists and managers may possess similar “know that,” their different practices keep them apart. The two can read the same books, magazines, and journals, but these don’t allow either to do the other’s job. A good management theorist may explain the practice of management well, but never make a good hands-on manager. Similarly, an excellent manager may prove an inept theoretician.¹⁸

LEARNING IN PRACTICE

Practice, then, both shapes and supports learning. We wouldn’t need to labor this point so heavily were it not that unenlightened teaching and training often pulls in the opposite direction. First, they tend to isolate people from the sorts of ongoing practice of work itself. And second, they focus heavily on information.

Nowhere is this isolation more true than in the workplace. Or perhaps we should say *than not* in the workplace. For while

many of the resources for learning to work lie in the workplace, training regularly takes people away from there, to learn the job in classrooms. The ideal of learning isolated from the “distractions” of work practice still influences many training regimens. So let us look briefly at a couple of examples that suggest some of the limits of the classroom and the resources of practice.

Limits to Going by the Book

The first example draws on research by two educational psychologists, George Miller and Patricia Gildea, into how children learn vocabulary. Miller and Gildea compared learning words in the everyday practice of conversation with trying to learn vocabulary from dictionaries.¹⁹

In the everyday case, they found that learning is startlingly fast and successful. By listening, talking, and reading, the average 17-year-old has learned vocabulary at a rate of 5,000 words per year (13 per day) for over 16 years. The children know both what these words mean and how to use them.

By contrast, learning words from abstract definitions and sentences from dictionaries is far slower and far less successful. Working this way, the children in the study acquired between 100 and 200 words per year. Moreover, much of what they learned turned out to be almost useless in practice. Despite their best efforts, looking up *relate*, *careful*, *remedy* and *stir up* in a dictionary led to sentences such as, “Me and my parents correlate, because without them I wouldn’t be here”; “I was meticulous about falling off the cliff”; “The redress for getting sick is staying in bed”; and “Mrs. Morrow stimulated the soup.”

Most of us have seen the workplace equivalent of this—the eager young intern with all the right information but none of

the practical knowledge that makes the job doable. Similarly, a lawyer friend of ours recalled how the first days at work were a nightmare because, despite all her excellent results in law school and on the law board exams, nothing in the classroom had prepared her for the realities of having a client on the other end of the telephone.

The Practical Value of Phone Cords

Another colleague, Jack Whalen, showed the power of practice in his study of learning in a service center taking the calls from customers and scheduling technicians.²⁰ Sending technicians to fix broken machines is an expensive undertaking. It is a waste if the problem does not really require a technician. So the people who take the calls can save the company money by diagnosing simple problems and telling the customer how to fix these for themselves. It makes customers happy, too. They don’t have to sit with a dead machine, waiting for a technician to bring it back to life.

The phone operators are not, of course, trained as technicians. In the past, however, they learned from the reps when the latter called in to pick up their next job. The reps would then explain how trivial the last one had been, and in the process the phone operators could learn a lot from these mentors. When they next took such a call, they could offer a solution. As a result of a change in communications technology, however, technicians no longer pick up their calls this way. Consequently, operators no longer pick up insights. Their opportunity for inherent learning has been lost.

The company has tried to replace this kind of learning with the more explicit support of a “case-based expert system.” This is an information-based system that prompts operators to ask the

customer a series of questions. The operator types the responses into the system, which then searches for a ready solution. This alternative has not worked well. As the reps found with "directive documentation," it can be surprisingly difficult to get a clear diagnosis and solution this way. Moreover, such a system doesn't help the operators understand what they are doing. And that lack of understanding undermines the customer's confidence. It's hard to put faith in people who are obviously reading instructions off a screen. As a result, customers will ask for a technician anyway, and so defeat the whole expert-system strategy.

To overcome these problems, the company contemplated new training courses with several weeks off site to better prepare new operators. Whalen and his fellow researchers took a slightly different route, however. They studied one service center and the quality of diagnosis its staff provided. There they found two operators who gave especially reliable answers. One, unsurprisingly, was an eight-year veteran of the service center with some college experience and a survivor from the days when reps served as mentors. The other, however, was someone with only a high-school diploma. She had been on the job barely four months.

The researchers noticed, however, that the newcomer had a desk opposite the veteran. There she could hear the veteran taking calls, asking questions, and giving advice. And she began to do the same. She had also noticed that he had acquired a variety of pamphlets and manuals, so she began to build up her own stock. Moreover, when she didn't understand the answers the veteran gave, she asked him to show her what he meant, using the service center's own copier.

So instead of training courses, the sociologists suggested restructuring the phone center. They sought to draw on its

reservoir of knowledge by putting all its operators in positions to learn from each other. By opening the place up to this collective knowledge, the redesign effectively created a small laboratory of what Whalen calls "indigenous sharing and collaborative learning." The new plan also asked technicians to come in and take calls intermittently. As a result, operators could learn from them once again.

From these changes, the operators were up to speed in about the time it took to plan a training course for them and in far less time than was set aside for actual training. Ultimately, Whalen concluded, given the amount and level of knowledge already available in the room, what the operators needed were not so much expert systems or new training courses, but "longer phone cords."²¹ (These allow an operator taking a call to slide over to the desk and the screen of a resourceful colleague who could provide the necessary help.) Both examples, the classroom and the workplace, indicate how the resources for learning lie not simply in information, but in the practice that allows people to make sense of and use that information and the practitioners who know how to use that information. Where in other circumstances knowledge is hard to move, in these circumstances it travels with remarkable ease.

PHILOSOPHICAL EXPLANATIONS

To venture cautiously again onto philosophical grounds, the distinction between explicit and implicit dimensions of knowledge can help illuminate why practice is so effective. It's possible, for example, to learn about negotiation strategies by reading books about negotiation. But strategy books don't make you into a good negotiator, any more than dictionaries make you into a speaker

or expert systems make you into an expert. To become a negotiator requires not only knowledge of strategy, but skill, experience, judgment, and discretion. These allow you to understand not just how a particular strategy is executed, but when to execute it. The two together make a negotiator, but the second comes only with practice.

The chemist and philosopher Michael Polanyi distinguished these two by talking about the explicit and the tacit dimensions of knowledge. The explicit dimension is like the strategy book. But it is relatively useless without the tacit dimension. This, Polanyi argues, allows people to see when to apply the explicit part.²²

To take another simple example of this sort of tacit “seeing,” consider dictionaries again. These are the guidebooks of language and particularly for spelling. But if you lack the tacit dimension required for spelling, shelves of dictionaries do you no good. For being able to use a dictionary (the explicit part) is not enough. You have to know when to use a dictionary. A good speller will say, “I just know that doesn’t look right.” This is the tacit part. Once it has done its work, you can turn to the explicit information in the dictionary. The problem for a bad speller, of course, is that if he or she lacks the tacit knowing that makes words look wrong, then a dictionary’s use is limited. In the end, paradoxically, you only learn to use a dictionary by learning to spell.

In making his distinction between explicit and tacit, Polanyi argues that no amount of explicit knowledge provides you with the implicit. They are two different dimensions of knowledge, and trying to reduce one to the other is a little like trying to reduce a two-dimensional drawing to one dimension. This claim of Polanyi’s resembles Ryle’s argument that “know that”

doesn’t produce “know how,” and Bruner’s that *learning about* doesn’t, on its own, allow you to *learn to be*. Information, all these arguments suggest, is on its own not enough to produce actionable knowledge. Practice too is required. And for practice, it’s best to look to a community of practitioners.

PRACTICAL IMPLICATIONS

*Teach these boys nothing but Facts. Facts alone are wanted in life.
Plant nothing less and root out everything else. You can only form
the minds of reasoning animals upon Facts. . . . Stick to Facts, Sir.*

—CHARLES DICKENS, *Hard Times*²³

The view of knowledge and practice we have put forward here has several implications for how to think about learning—and related issues such as spreading best practice, improving performance, or innovating—as well as training and teaching.

In the first place, it takes us beyond information. The idea of learning as the steady supply of facts or information, though parodied by Dickens 150 years ago, still prevails today. Each generation has its own fight against images of learners as wax to be molded, pitchers to be filled, and slates to be written on.

Literature about workplace learning is still laced with ideas of “absorptive capacity,” as if humans were information sponges. Indeed, the idea that learning is mere information absorption may be on the rise today because it allows for more redefinition. If we accept this view of learning, then it’s a short step to talking about such things as computers or bots learning, as if what they do is just what people do. Looking beyond information, as we have tried to do, provides a richer picture of learning. From this picture, the following features stand out for us.

Learning on Demand

Learning is usually treated as a supply-side matter, thought to follow teaching, training, or information delivery. But learning is much more demand driven. People learn in response to need. When people cannot see the need for what's being taught, they ignore it, reject it, or fail to assimilate it in any meaningful way. Conversely, when they have a need, then, if the resources for learning are available, people learn effectively and quickly.

In an essay we wrote about learning some years ago, we referred to this aspect of learning as "stolen knowledge." We based this idea on a short passage in the biography of the great Indian poet and Nobel Prize winner Rabindrath Tagore.²⁴ Talking of an instructor hired to teach him music, Tagore writes, "He determined to teach me music, and consequently no learning took place." Tagore found little to interest him in the tedious tasks he was given as practice for these involved not the authentic activity itself, but only a pale imitation. "Nevertheless," he continues, "I did pick up from him a certain amount of stolen knowledge."²⁵ This knowledge, Tagore reveals, he picked up by watching and listening to the musician when the latter played for his own and others' entertainment. Only then, when what was evident was the practice of musicianship and not dismembered teaching exercises, was Tagore able to see and appreciate the real practice at issue.

A demand-side view of this sort of knowledge theft suggests how important it is not to force-feed learning, but to encourage it, both provoking the need and making the resources available for people to "steal." We regard this as the paradoxical challenges of encouraging and legitimating theft. Organizations have become increasingly adept both at provoking and at responding

to changes in their clients' needs. They need to consider how to do this for their employees as well.

Social Learning

Despite the tendency to shut ourselves away and sit in Rodinesque isolation when we have to learn, learning is a remarkably social process. Social groups provide the resources for their members to learn. Other socially based resources are also quite effective.

For example, people who are judged unfit to learn to operate simple tools or who fail to master domestic appliances nevertheless learn to operate complex machines that present users with hazardous, changing environments and sophisticated technologies. We refer, of course, to the car. Technologically, cars are extremely sophisticated. But they are also extremely well integrated socially. As a result, learning becomes almost invisible. Consider, by contrast, the triumphal despair with which people frustratedly boast that they still can't program their VCR. The success of learner drivers—with or without instruction—should undoubtedly be the envy of many who design far less difficult consumer or workplace appliances.

The car and the VCR make an interesting contrast. Almost everyone in our society who learns to drive has already spent a great deal of time traveling in cars or buses, along roads and highways. New drivers begin formal instruction with an implicitly structured, social understanding of the task. Now consider the VCR. Most people can use their machine to play tapes. What they find difficult is recording, though that's not a much more complex task. The central distinction between these two functions is that one is often a social act, the other highly individual. You might invite a group over to watch a movie. You are unlikely to invite one over to watch you record.²⁶

Learning and Identity Shape One Another

Bruner, with his idea of learning to be, and Lave and Wenger, in their discussion of communities of practice, both stress how learning needs to be understood in relation to the development of human identity. In learning to be, in becoming a member of a community of practice, an individual is developing a social identity. In turn, the identity under development shapes what that person comes to know, how he or she assimilates knowledge and information. So, even when people are *learning about*, in Bruner's terms, the identity they are developing determines what they pay attention to and what they learn. What people learn about, then, is always refracted through who they are and what they are learning to be.²⁷

So information, while a critical part of learning, is only one among many forces at work. Information theory portrays information as a change registered in an otherwise steady state. It's a light flashing out on a dark hillside (to borrow an example from the philosopher Fred Dretske²⁸) or the splash of a pebble breaking the calm of a still lake. In either case, the result, as the anthropologist Gregory Bateson puts it neatly, is "a difference that makes a difference."²⁹

The importance of disturbance or change makes it almost inevitable that we focus on these. We notice the ripple and take the lake for granted. Yet clearly the lake shapes the ripple more than the ripple shapes the lake. Against a different background, the pebble would register a different change or perhaps, in Bateson's terms, make no difference at all. So to understand the whole interaction, it is as important to ask how the lake is formed as to ask how the pebble got there. It's this formation rather than information that we want to draw attention to,

though the development is almost imperceptible and the forces invisible in comparison to the drama and immediacy of the pebble.

It's not, to repeat once more, the information that creates that background. The background has to be in place for the information to register. The forces that shape the background are, rather, the tectonic social forces, always at work, within which and against which individuals configure their identity. These create not only grounds for reception, but grounds for interpretation, judgment, and understanding.

A Brief Note on the "Social"

We emphasize the social side of learning and identity with some caution. The economist Friedrich Hayek claims that *social* is a weasel word.³⁰ Moreover, people readily point out that they can learn a great deal sitting alone in an office or a library. And you don't have to go very far with the thesis that learning is significantly social to encounter the question "What about Robinson Crusoe?"³¹

Early economists liked to present Crusoe as an example of the *homo economicus*, the universal economic man, learning and working in splendid individual independence. And that's the idea behind this question. It took Karl Marx to point out, however, that Crusoe is not a universal. On his island (and in Defoe's mind), he is deeply rooted in the society from which he came:

*Our friend Robinson . . . having rescued a watch, ledger, and pen and ink from the wreck, commences, like a true-born Briton, to keep a set of books. His stock book contains a list of the objects of utility that belong to him, of the operations necessary for their production; and lastly, of the labour time that definite quantities of those objects have, on average, cost him.*³²

Robinson is not just a man in isolation, but a highly representative member of what Napoleon was to call a “nation of shopkeepers.”

It is, of course, not only the British who play to type, even when alone. The French philosopher Jean-Paul Sartre, in a famous passage, illustrates how that true-born Frenchman, the café waiter, though working alone, conforms his actions to society’s idea of what a waiter does:

He returns, trying to imitate in his walk the inflexible stiffness of some kind of automaton while carrying his tray with the recklessness of a tight rope walker by putting it in a perpetually unstable, perpetually broken equilibrium which he perpetually re-establishes by a light movement of the arm and hand. . . . We need not watch long before we can explain it: he is playing at being a waiter in a café [T]he waiter plays with his condition in order to realize it.³³

So while people do indeed learn alone, even when they are not stranded on desert islands or in small cafés, they are nonetheless always enmeshed in society, which saturates our environment, however much we might wish to escape it at times. Language, for example, is a social artifact, and as people learn their way into it, they are simultaneously inserting themselves into a variety of complex, interwoven social systems.

LEARNING DIVISIONS

Looking at learning as a demand-driven, identity forming, social act, it’s possible to see how learning binds people together. People with similar practices and similar resources develop similar identities—the identity of a technician, a chemist, a lepidopterist, a train spotter, an enologist, an archivist, a parking-lot attendant, a business historian, a model bus enthusiast, a real estate

developer, or a cancer sufferer. These practices in common (for hobbies and illnesses are practices too) allow people to form social networks along which knowledge about that practice can both travel rapidly and be assimilated readily.³⁴

For the same reason, however, members of these networks are to some degree divided or separated from people with different practices. It is not the different information they have that divides them. Indeed, they might have a lot of information in common. Rather, it is their different attitudes or dispositions toward that information—attitudes and dispositions shaped by practice and identity—that divide. Consequently, despite much in common, physicians are different from nurses, accountants from financial planners.³⁵

We see two types of work-related networks that, with the boundaries they inevitably create, are critical for understanding learning, work, and the movement of knowledge. First, there are the networks that link people to others whom they may never get to know but who work on similar practices. We call these “networks of practice.” Second, there are the more tight-knit groups formed, again through practice, by people working together on the same or similar tasks. These are what, following Lave and Wenger, we call “communities of practice.” Here we sketch the two briefly before elaborating their role in later chapters.

Networks of Practice

While the name “networks of practice” helps us to emphasize what we see as the common denominator of these groups—practice—elsewhere they go by the name of “occupational groups” or “social worlds.”³⁶ People in such networks have practice and knowledge in common. Nevertheless, most of the members are

unknown to one other. Indeed, the links between the members of such networks are usually more indirect than direct—newsletters, Web sites, Bulletin boards, listservs, and so forth keep them in touch and aware of one another.³⁷ Members coordinate and communicate through third parties or indirectly. Coordination and communication are, as a result, quite explicit.³⁸

The 25,000 reps working for Xerox make up, in theory, such a network. They could in principle be linked through such things as the Eureka database (though it is in fact not worldwide) or corporate newsletters aimed at reps. Their common practice makes these links viable, allowing them to assimilate these communications in more-or-less similar ways. By extension, the network could also include technicians in other companies doing the same sort of work, though here the connections would be weaker, grounds for common understanding more sparse.

Networks of this sort are notable for their reach—a reach now extended and fortified by information technology. Information can travel across vast networks with great speed and to large numbers but nonetheless be assimilated in much the same way by whomever receives it. By contrast, there is relatively little reciprocity across such network; that is, network members don't interact with one another directly to any significant degree. When reach dominates reciprocity like this, it produces very loosely coupled systems.³⁹ Collectively, such social systems don't take action and produce little knowledge. They can, though, share information relating to the members' common practices quite efficiently.

Communities of Practice

Lave and Wenger's notion of communities of practice, which we mentioned earlier, focuses on subsections of these larger networks

of practice. These subsections stand in contrast to the network as a whole in several ways. They are relatively tight-knit groups of people who know each other and work together directly. They are usually face-to-face communities that continually negotiate with, communicate with, and coordinate with each other directly in the course of work. And this negotiation, communication, and coordination is highly implicit, part of work practice, and, in the case of the reps, work chat.⁴⁰

While part of the network, groups like this cultivate their own style, their own sense of taste, judgment, and appropriateness, their own slang and in-terms. These things can distinguish members of one community within a network from others. In networks of scholars, for example, while all may be from one field, it's often easy to guess who trained together in a particular lab or school by their style and approach.

In these groups, the demands of direct coordination inevitably limit reach. You can only work closely with so many people. On the other hand, reciprocity is strong. People are able to affect one another and the group as a whole directly. Changes can propagate easily. Coordination is tight. Ideas and knowledge may be distributed across the group, not held individually. These groups allow for highly productive and creative work to develop collaboratively.

UNDERSTANDING DIVISION

The divisions marked by the external boundaries of these groups have significant implications for the development of organizations, technologies, and indeed of societies as a whole. Yet they are divisions that discussions of such developments easily overlook.

For example, discussions of the emerging “network society” suggest that society is becoming a single, uniform entity. The network stretches indefinitely, linking the individuals that stand at each node to one another and providing them with common information. Communities, organizations, nations, and the like disappear (victims of the 6-Ds discussed in chapter 1). The network is all, configuring itself more or less as the vaunted global village.

From the perspective of practice, rather than of process or information, a rather different picture emerges. From this viewpoint, any global network has a highly varied topography. While the whole may ultimately be global, within it there are networks of practice with lines of reach that are extensive but nonetheless bounded by practice. And there are communities of practice, with dense connections of both reach and reciprocity, which again put limits on extent. These two, networks and communities, produce areas marked by common identity and coordinated practice within any larger network. And as a consequence of these areas, information does not travel uniformly throughout the network. It travels according to the local topography.

Curiously, organization theory suffers from similarly homogenizing vision. It has been fashionable of late to talk of workplace culture or organizational culture as if these made organizations internally uniform. But divisions created by practice produce significant variation here as well. Within organizations as without, connections are dense in some places and thin in others. Sometimes these networks extend across the boundaries of the organization. Elsewhere, they may confront discontinuities within, where meaningful communication breaks down.

Business process reengineering, in particular, ignores divisions created by different practices. Indeed, Hammer and Champy's *Reengineering the Corporation* seeks to supersede the

division of labor that the economist Adam Smith saw as central to capitalist production.⁴¹ Consequently, business process reengineering fails to understand the internally varied terrain of organizations and its fractures and divisions.

Sim Sitkin of the Fuqua School of Business at Duke University and a colleague revealed similar blindness in “Total Quality Management.” Managers of one large organization, Sitkin and his colleague found, attempted to implement a single scheme uniformly across a company. The approach overlooked the different ideas of quality that different practices develop. Pursuing a common goal in the face of these divisions made the different groups feel that they were being judged by the standards of others, and in consequence, fear their work would either be undervalued or unrecognized. Rather than spreading quality, the researchers concluded, the scheme only spread dissent and distrust.⁴²

Failure to read the topography may be at its most damaging as people try to predict the effects of new information technologies on organization. One of the remarkable things about these technologies is their reach. Consequently, they are well suited to support, develop, and even strengthen the networks of practice we have described. As these networks readily span the borders of organizations, their increasing strength will affect those organizations. Whether networks will grow at the expense of organizations is a question needing further research, not the linear assumptions of infoenthusiasts. It seems improbable that they will simply dissolve organization any more than, as we claimed in chapter 1, they will necessarily damage local communities, which remain robust.

New technologies may, though, spread these communities out more than before. The growing reciprocity available on the

'Net, while probably underused at the moment, is helping people separated by space maintain their dense interrelations. Yet for the sort of implicit communication, negotiation, and collective improvisation that we have described as part of practice, learning, and knowledge sharing, it's clear that there are advantages to working together, however well people may be connected by technology. Indeed, one of the most powerful uses of information technology seems to be to support people who do work together directly and to allow them to schedule efficient face-to-face encounters. Looking too closely at the progression from atoms to bits may miss the role the bits play in allowing us to reinforce the valuable aspects of the world of atoms. Critical movements in the knowledge economy may go not just from atoms to bits, but from atoms to bits and back again.

CHAPTER 6

Innovating Organization, Husbanding Knowledge

Electronic commerce is the single greatest change in business since the invention of money. . . . On the Internet there's perfect information. . . . Location doesn't protect you any more. It's a truly global economy.

—National Public Radio Business News¹

IF THE RECESSION OF the early 1990s disrupted U.S. business, as we noted in chapter 4, the long expansion that has followed has not, as might have been expected, brought calm. Far from it. As the recovery got under way, the World Wide Web arose out of nowhere, popularizing the Internet, populating cyberspace, and bringing on its back the turmoil of "e-commerce" and the challenges of "globalization."

At first e-commerce was seen as a matter for hobbyists only. (The personal computer was greeted in much the same way.) Economic reality soon changed that opinion. In 1995, Web sales

firms and formal institutions. We suspect it will again. Those who insist that neither regions nor institutions are significant, that distance and organizations are dead, will have trouble coping with this sort of innovation. To play with boundaries—of firms, networks, communities, regions, and institutions—as innovation increasingly seems to demand, requires first acknowledging them.

These issues about resources and constraints, structure and spontaneity, reach and reciprocity, and about dominant, residual, and emergent institutions are ones that affect design—from the design of digital appliances to the design of transnational organizations and institutions. In the conditions in which we live, we are all, as we suggested at the beginning, both producers and consumers of designed objects. So such questions and their implications affect us all, presenting themselves to us with greater speed and greater urgency than ever before. As we said at the outset, we do not have solutions to offer. We only know that solutions will be much harder to find if we drive at the problems with tunnel vision—if, to repeat our long list, peripheries and margins, practices and communities, organizations and institutions are left out or swept out of consideration.

Notes

A Note on URLs

As much of this book concerns the popular portrayal and popular perception of information technology, we have deliberately drawn on popular, rather than scholarly sources wherever possible. We have used newspapers, magazines, radio broadcasts, and, most important, the World Wide Web. Moreover, wherever we could, we have also provided URLs (Uniform Resource Locators) for print material that appears on the Web as well.

In directing readers to these Internet sources, we face the problems that every seasoned Internet user will recognize. Publishing is a slow, deliberative process. The 'Net, by contrast, is lightning fast, its content often shifting or disappearing. (We discuss this contrast in chapter 7.) Trying to pin down references in these conditions is remarkably difficult.

Nonetheless, where we have used the Web as a source we have provided URLs indicating that source. Before sending our manuscript to the press in July 1999, we attempted to revisit each of these URLs. Where we were successful, we have given the date (1999, July 21). Often, inevitably, we were not successful. Several pages we cite or quote have since been taken down, moved, changed. In such cases, we give the date when we last successfully visited that site.

Still, some URLs with the July date will not necessarily lead to the page cited for a variety of reasons. Frames technology, for instance, does not identify different pages separately. Some sites will maintain a URL but change the page's content. And several

sites (Wired.com and CNET.com among them) will not let you go directly to individual pages. On these sites, you must go to their “front door” and, generally, use their search engine.

Acknowledgments

1. See Brown and Duguid, 1996a, 1996b.

Introduction: Tunneling Ahead

1. Terkle, 1995.
2. Cutting off such resources can make a significant difference. Consider, for example, the historic Nixon–Kennedy debates of 1960. The majority of those who listened to the radio, who heard only the voices, apparently inclined to judge Nixon the winner. Those who watched television, who saw Kennedy’s ease and Nixon’s unease, leaned the other way.
3. Tenner, 1997.
4. Duguid, 1996b.
5. The term “infomated” comes from Zuboff, 1988.
6. Nor does widening the “bandwidth”—a standard strategy of adding more information when things go wrong—offer the solution. (See our notion of “Moore’s Law” solutions in chapter 1).

Chapter 1: Limits to Information

1. Gates, 1995, p. 21; Gordon Bell and James N. Gray, quoted in Nardi and O’Day, 1999; Phillips, quoted in Napoli, 1999.
2. SETI stands for Search for Extra Terrestrial Intelligence. For more information, see <http://seti.ssl.berkeley.edu> [1999, July 21].
3. The notion of the “third wave” comes from the futurist Alvin Toffler’s (1980) book of the same name.

4. Indeed, the degree to which we complain about such technologies when they do go wrong indicates the extent to which we have all become dependent on them to go right. The apprehension about the Y2K bug, which some predict may allow tiny embedded chips to bring huge social institutions to a halt, is itself a sign of how deeply enmeshed these things are in our lives.
5. Sometime around 2012, it has been predicted, Moore’s Law will come up against the physical limitations of current microchip components, though by then solid-state components may well have been replaced.
6. Kelly, 1997.
7. Negroponte, 1995. John Tukey coined the term *bit*. It stands for “binary digit.”
8. It’s worth remembering that formal information theory, while it holds the bit as a central concern, is indifferent to meaning. Claude Shannon, who with Warren Weaver laid the foundations of modern information theory, is quite clear about this: “the semantic aspects of communication are irrelevant to the engineering aspects” (Shannon and Weaver, 1964, p. 8).
9. The pervasive image of the new open frontier only works if we forget the presence on the old frontier of the U.S. Army, the Church of Latter-Day Saints, and myriad other organizations and institutions large and small (let alone its native inhabitants) that shaped the frontier.
10. For “electronic frontier,” see Rheingold, 1993; for “global village,” see McLuhan, 1962; and for “electronic cottages,” see Toffler, 1980, and chapter 3.
11. The phrase “information wants to be free” is usually attributed to the author Bruce Sterling.
12. Castells, 1996.
13. The SETI project (see note 2), after all, acknowledges that the universe has long been capable of producing raw information. Finding intelligence in its midst is a different matter.
14. *Business Week*, 30 June 1975, p. 48.; Toffler, 1980, p. 205.

15. Negroponte, 1995, p. 187.
16. For the dismissal of the pencil, see Petrosky, 1990, p. 331.
17. See the IBM advertisement in the *New York Times*, 30 June 1999, sec. C, pp. 13–20.
18. Gates, 1995, p. 5.
19. The idea is certainly a tempting one to the weary business traveler and is echoed in the curious enigma of the laptop. Ads for these suggest that laptops can be so “loaded” with communications software that you can travel anywhere and remain a virtual presence in your own office. Yet in suggesting this possibility, they make you wonder why you need to travel at all.
20. Reagle, 1996.
21. Readers of Toffler’s (1980) *Third Wave* will recognize the first three terms here, particularly the first, *demassification*, to which Toffler adds three subtypes: demassification of media, production, and society. Notions of *disintermediation* and *decentralization* are features, for example, in the work of George Gilder or Kevin Kelly’s (1997) writing on the “new economy.” There are more “Ds” that could be added, such as Kevin Kelly’s *displacement* and *devolution*.
22. Downes and Mui, 1998.
23. Coase, 1937. Coase’s theory should be seen not so much as an attack on neoclassical individualism as an attempt to save it from itself. We return to transaction cost theory briefly in our discussion of the future of the firm in chapter 6. There we take a “knowledge based,” rather than transaction cost, view of the firm.
24. Among the targets of early, landmark trust cases were Northern Securities (1911), Standard Oil (1911), and American Tobacco (1911). In November 1998, Philip Morris acquired several brands from the Liggett corporation.
25. *The Economist*, 13 December 1997.
26. Daniel, 1996, table 3.2. We take up this topic again in chapter 8.
27. The concentration of media ownership was the subject of a

special edition of *The Nation* (March 13, 1997). The picture drawn there recalled Frank Norris’s famous image of the railroads as an “octopus,” but, as with any list we might produce, the *Nation*’s octopus is already well out of date.

28. Paul Krugman, quoted in Kelly, 1998.
29. The appearance of the fax provides a conventional example of network effects at work. At first, for most people, it was hardly worth owning a fax machine because so few others did. But with each new buyer, owning a fax became more useful and so more valuable. Each addition to the network meant that everyone had more people to fax to and receive faxes from. By the late 1980s, the network had grown to such a size that owning one became almost a necessity.

The fax spread on the back of nonproprietary standards that allowed many different firms to compete in the fax market, as they still do. But if someone had owned those standards, the results would have been quite different. Here the obvious example is not the fax, but the video recorder and the battle between VHS and Betamax. This battle was in effect a race between the two to establish network effects. In such races, a small lead can produce a “tipping” effect, where whoever gets ahead can quickly end up taking it all. In such conditions, firms will grow, not shrink.
30. Raik-Allen, 1998.
31. Sassen, 1996, p. 38.
32. With the rise of the Internet, many people argue that information can travel straight from its producer to its consumer, without need for any intermediaries. Kelly (1997), for example, envisages music going from musicians to listeners and then leaps to gourmet coffee from producers to drinkers. Undoubtedly disintermediation along these lines is happening; however, there’s a good deal of wishful endism here, too. Specific examples always seem to include unpopular professions that people want to see ended, with banking and real estate falling in behind travel agents and car salespeople. One curious example

of a job relentlessly pushed off the stage yet refusing to sing its own swan song is that of the humble meter reader. Most recently, Jeremy Rifkin (1995) sang a requiem for this job in 1995. Curiously, so did articles on the effects of computers in 1970 in *National Geographic*, and as far back as 1952 in *Fortune*.

33. Fukuyama and Shulsky, 1997, p. 63
34. Attewell, 1994, p. 36.
35. Fukuyama and Shulsky, 1997, p. 11.
36. Innis, 1991. For a case study of previously decentralized business organizations that were later centralized, see Duguid and Silva Lopes, 1999.
37. Fidler, 1997.
38. The same long arm of technology has led to the direct intervention of politics into battlefield planning; for example, though President Bush said he would leave decisions in the Gulf War to local commanders, the White House began to exercise veto power and control over the conduct of the war, after a smart bomb went astray and drew bad press.
39. Shoshona Zuboff, quoted in Lohr, 1996.
40. Meanwhile, Royal Dutch/Shell, one of the most well known and widely applauded decentralizers of the past decade, has announced that it will recentralize its "treasury." Decentralization had been too costly and inefficient. More generally, the *Economist* magazine's "Intelligence Unit" has noted a trend toward "shared services" in large corporations—and what they describe reads very much like recentralization.
41. Negroponte (1995) suggests that "a few strokes of the key-board" will close the generation gap (p. 203).
42. Nunberg, 1996. Coffman and Odlyzko (1998) show how many projections about the current and future growth of the 'Net extrapolate from a particularly rapid point in its development, failing to note that both before and after this point, the growth curve, though impressive, was far less steep.
43. Hobsbawm, 1977, pp. 61–62. The spread of the telephone was more dramatic, growing from 14 miles in 1845 in the United

States to 670,000 in 1886, but the infrastructure of the telephone was far easier to build than railway lines. The spread of the radio was more impressive yet.

44. See Campbell-Kelly and Aspray, 1996.
45. Wellman (1988) provides one of the few worthwhile studies of the effects of information technologies on social communities and networks.

Chapter 2: Agents and Angels

1. Distinguishing a computer from a human is the essence of the famous Turing test, developed by mathematician Alan Turing (1963). He argued that if you couldn't tell the difference then you could say the machine was intelligent. Shallow Red is not quite there yet. Indeed, the continuation of the exchange suggests that Shallow Red is still rather shallow (though pleasantly honest):

What are VSRs?

The botmaster has not provided me with a definition of "VSRs." Thank you for asking. The botmaster shall provide me with an appropriate definition.

Shallow Red could be found at <http://www.neurostudio.com> [1998, October 21]. It has since been upgraded to Red, though it is not yet quite as deep as that promotion might suggest.

2. The lists you get if you shop for other people or buy gifts online make it clear that tastes and tracks are different things.
3. The study, done at Yale, was reported on National Public Radio's *All Things Considered*, 18 November 1998.
4. Leonard, 1997, p. 186.
5. Hoggdon, 1997, p. 15.
6. Though less so, perhaps, than with the telegraph, which the president of the Magnetic Telegraph Co., W. M. Swain, suggested might be the fulfillment of scriptural prophecy.

35. Evolutionary economics looks closely at routine, as does Stinchcombe (1990) on whose ideas we lean here.
36. From the perspective of process, however, it may appear that members of an organization divide between the two categories. Most are expected to follow routine. Some—those in R&D, corporate strategy, and the like—are expected to improvise. But again, in practice, things don't break down so easily. People in the labs, for example, have to learn to tie their improvisations to company routine, otherwise their work detaches itself from the company and its goals. (We'll discuss this tension within innovation in chapter 6.)
37. This sort of processing is evident in the shoppers Lave (1988) studied. See chapter 3.
38. Hammer and Champy, 1993.
39. FAQ stands for "frequently asked questions."
40. Andrea Horneett, in preparation. See our discussion of Networks of Practice in chapter 6.
41. Granovetter, 1985, p. 502.
42. Ellen O'Connor, 1999.

Chapter 5: Learning—in Theory and in Practice

1. http://www.microsoft.com/exchange/community/lotus_km.asp [1998, December 20]; <http://www.microsoft.com/mcsp/coreso.htm> [1998, December 20].
2. Donkin, 1998.
3. On its Microsoft Exchange Web site, Microsoft formerly accused its competitor Lotus of merely rebranding old-fashioned groupware to claim relevance in the escalating knowledge wars (http://www.microsoft.com/exchange/community/lotus_km.asp [1998, December 20]). In response, Lotus admitted that it has relabeled its groupware, but argued that this merely showed that Lotus had always been concerned with knowledge: "But knowledge management is nothing new. It's a different label applied to something that Lotus has been working on for years" (<http://www.lotus.com/news/topstories.nsf/a1d792857da52f638525630f004e7ab8/ee02468b55e0824b85256598004a827d?OpenDocument> [1999, July 21]).
4. See chapter 7 for a further discussion of related "conduit" metaphors. As Reddy (1979) notes, these metaphors make it easy to overlook questions of meaning in discussions of information.
5. Shannon and Weaver, 1964.
6. Davenport and Prusak (1998) use the Ford example to talk about "corporate amnesia" (p. x).
7. Professor Amin Rajan, quoted in "Managing Knowledge Is Latest Business Buzz," 1998.
8. Stewart, 1997, p. 88.
9. Leonard-Barton, 1995, p. 7.
10. Cole, 1999, see chapter 8, "Adoption, Adaptation, and Reaction at Hewlett Packard."
11. This is how the story is commonly told. O'Dell and Grayson (1998) trace the remark to Junkins, CEO of Texas Instruments, who said, "If TI only knew what TI knows." It was then echoed by Pratt, chairman of HP, who said, "I wish we knew what we know at HP."
12. "Managing Knowledge Is Latest Business Buzz," 1998.
13. Lotus made the same shift to information but, adding in expertise to the equation, is a little more cautious about making technology the central issue: "Knowledge Management [is] the systematic leveraging of information and expertise to improve organizational innovation, responsiveness, productivity and competency" (<http://www.lotus.com/home.nsf/welcome/km> [1998, December 20]).
14. Lave and Wenger, 1993.
15. See, for example, Karin Knorr Cetina's (1999) study, *Epistemic Cultures*.
16. Heart interviewed in Gates, 1998.
17. Ryle, 1949.
18. The last example helps overcome conventional divisions between theory and practice. A management theorist is a

- practitioner. But his or her practice is the practice of theorizing, whereas a manager's is the practice of managing. Consequently, each has different "know how."
19. Miller and Gildea, 1987.
 20. Whalen and Vinkhuyzen, forthcoming.
 21. Jack Whalen, interview with authors, 16 June 1998.
 22. Polanyi, 1966.
 23. Dickens, 1969, p. 1. Part of the enjoyment of the passage is that Mr. Gradgrind's lecture, while trying to form the mind of a reasoning animal, contains not one fact, only opinion and imperatives.
 24. Bandyopadhyay, 1989.
 25. Quoted in Bandyopadhyay, 1989, p. 45.
 26. Too often the technology people work with, while designed for learning, produces this sort of damaging isolation, allowing people little insight into where their work goes, what it means, or why it matters. A couple of years ago a young architect wrote an article in *Architecture* complaining that spending her day in front of a CAD (computer aided design) machine prevented her from gaining any understanding of the practice of architecture and the path that leads from clients to buildings. All she executed were tasks, and in isolation they gave no real insight into how they had arisen or what role they played. *Architecture* 185 (1): 30.
 27. Orr understands the Xerox reps, for example, in terms of their collective development of the identity of a rep. The company, however, often addresses them as though the identity they are developing is that of a Xerox employee. The two may be quite distinct.
 28. Dretske, 1981.
 29. Bateson, 1972.
 30. Hayek, 1988, pp. 114–16. Hayek's claim primarily reflects the idea that everything is social, so adding the word *social* serves no real purpose. A recent book on Hannah Arendt, arguing that she used the word so capriciously that it could mean

- almost anything, was called *The Attack of the Blob: Hannah Arendt's Concept of the Social* (Pitkin, 1998). The word *blob* in the title stands for Arendt's use of *social*.
31. The first thing to remember, of course, is that Crusoe is a fictional character.
 32. Marx, 1947, p. 48.
 33. Sartre, 1957, quoted in Warnock, 1960, pp. 127–28.
 34. See Feenburg (1995) for an insightful discussion of the social character of illness.
 35. These examples are from Dorothy Leonard-Barton and Silvia Sensiper (1998) and Kenneth Arrow (1984), respectively.
 36. See van Maanen and Barley (1984) for "occupational communities"; Strauss (1978) and the following chapter for "social world."
 37. Listservs are e-mail lists that forward messages sent to a single address to everyone who subscribes to that list. Members of large lists rarely know who the other members are. Some lists are edited so that only messages approved by the editor are forwarded.
 38. This demand for explicit communication puts a strain on that communication. Creating the insights from which tips are written is part of the practice of Xerox reps. Turning these insights into written tips for the reps' database, however, is another practice that reps don't necessarily find easy.
 39. Weick, 1976.
 40. Here we are using *negotiation* in the *implicit* sense that we outlined in chapter 2.
 41. Hammer and Champy, 1993.
 42. Sitkin and Stickel, 1996.

Chapter 6: Innovating Organization, Husbanding Knowledge

1. Professor Arturo Pérez Reyes interviewed on National Public Radio's *Morning Edition*, 3 June 1999.
2. The \$436 million figure comes from the Internet Index 14