

Persuasive Encounters: Ethnography in the Corporation

BRIGITTE JORDAN
BRINDA DALAL
Palo Alto Research Center

In corporate settings, ethnographic methods are challenged routinely by managers who confront ethnographers with a set of typical objections that question the validity and effectiveness of ethnographic methods, findings, and recommendations. This article offers a series of steps toward overcoming this impasse by laying out a set of arguments for legitimizing ethnographic work. We discuss ways of responding to a variety of problematic encounters, involving some relatively quick answers to challenges of that sort but also acknowledging that the different worldviews of managers and ethnographers can be reconciled only in a long-term educational effort. In the last analysis, embedding ethnography in corporations is an exercise in culture change that almost always relies on rephrasing questions and reformulating metaphors to resituate our practice.

Keywords: workscapes; ethnographic methods; work-practice analysis; corporate anthropology

In this article, we address the daily reality of ethnographers doing work in a corporate world in which they persistently face questions about the validity and effectiveness of their methods and results from managers. Managers, especially inexperienced ones, often feel that they are exposing themselves to unknown risk. They may have nonverbalized questions, sometimes as basic as the following: Does this work at all? Am I running the risk of exposing myself and my people to criticism that I can avoid with other methods? and the always ubiquitous question, What is the financial impact? Sometimes

This article is based on a paper we delivered at the 2003 meetings of the Society for Applied Anthropology. We thank the many colleagues who have commented on subsequent versions, especially our colleagues at the Palo Alto Research Center (PARC) and the former Institute for Research on Learning (IRL). If we could, we would name all of them as coauthors to acknowledge the magnitude of our debt. For specific comments on the present manuscript, we particularly thank Françoise Brun-Cottan, Elizabeth Churchill, Robert Irwin, Les Nelson, Langdon Morris, Steve Smoliar, Peggy Szymanski, Alessandra Talamo, Jack Whalen, and Marilyn Whalen and acknowledge that we have shamelessly incorporated their suggestions. As always, this article should be seen as a collaborative piece of work.

Field Methods, Vol. 18, No. 4, November 2006 1–24
DOI: 10.1177/1525822X06292976
© 2006 Sage Publications

these questions seek information; sometimes they are hostile—coming in the form of a sniper attack or a deprecatory comment thrown out in a meeting where it is inappropriate to respond on the spot. What do you do in such situations? How can you turn a conversation of that sort not only into a productive response, but a persuasive encounter?

In the following pages, we discuss a set of typical challenges to corporate ethnographic work and propose effective strategies for dealing with them. Our methods are grounded in ethnographic practice—our own and that of our colleagues. So, although we may sometimes paint extreme positions, we recognize that there are always other ways of looking at things and other solutions to the issues we encounter routinely. We will draw examples from the wide range of our ethnographic experience: product design, technology development, innovation, and organizational change practice. However, regardless of the specific domain, these problems may be common in all kinds of ethnographic work. The issues we discuss are increasingly lively in corporate operations, engineering, and industrial research labs, many of which are struggling with methodological questions and the legitimacy of ethnography. What is distinct in the corporate domain is the potential realm of influence: Ethnographers not only are hired for their academic research skills but are responsible for customer engagements that understand and shape change itself.

We expect to find our audience both in academic and corporate settings: In the first are students and teachers of what variously has been called business, industrial, or corporate anthropology; in the second are corporate ethnographers, middle managers, and corporate decision makers. Although ethnographic methods have come to be accepted widely in companies large and small, this article also is aimed at those working where the techniques are less well known to management.

“THIS TAKES TOO LONG!”

A commonly heard objection to ethnographic work-practice analysis is that “it takes too long.” This is frequently raised by those who have become accustomed to basing their decisions on market research reports, and especially, the results of focus groups or surveys, both of which produce concrete output in short order, while ethnographers may argue for studying a given site for many months.¹

The truth is that sometimes, ethnographic work-practice analysis can generate results in surprisingly short periods. For example, when it was necessary to move a call center with several hundred employees from a high-prestige downtown location in Dallas to a new rural facility, we were asked to provide

suggestions for the redesign of an already existing warehouse. Our original charge had been to help the call center transform itself into a learning organization, but with the unexpected order to move, management rearranged priorities. Out of the blue, the center manager wanted to know what we could tell him that would be useful for getting the new facilities ready. When we pointed out that we had barely begun fieldwork, he said, with desperation in his voice, "I know, I know. But can't you tell us *something*?"

So we took a look at the warehouse. It was truly dismal. No light, low roofs, and a chicken farm across the street. We also spent some days at the downtown location, looking at current conditions, talking to the employees, noting their concerns. We attended a meeting with the architects during which they did a requirements analysis. Then we wrote a four-page memo to management with recommendations that ranged from improving employee voice in requirements analysis to strengthening communication with employees about timetables and responsibilities for the move, proposing landscaping to deal with the chicken coop problem, designing "home bases" for work teams, improving the interior design of the warehouse with skylights and informal gathering places, and relieving the monotony of rows of cubicles by dedicating a piece of unused space to an indoor garden atrium. Management was delighted with our report, which, incidentally, convinced them that we might be able to help them solve the learning problems for which we originally had been hired. The center manager even suggested that our report provided more insight than that of the architects and was a lot cheaper.

Obviously, such quick results may happen. Or not. A much more realistic expectation is that the longer the field engagement lasts, the deeper will be the insights it produces. If it has to do with, say, wanting to know how current work practices might affect the adoption of a new tool, then that might be done in days or weeks. If, on the other hand, we are talking about a systematic assessment of how a business division works, about creating an environment for learning and innovation, or about understanding how a factory works before, during, and after automation, then we are certainly talking about months, and ideally, an involvement that extends beyond a year. The extent to which ethnographers are already familiar with a given company, either as employees or because of prior project work, will reduce the time necessary to become familiar with the local *modus operandi* and build the necessary collaborative relationships at the field site. When ethnography is part of a larger project, results depend on the scope and purpose of the ethnographic component within the engagement. In general, a longer time frame allows us to attack more systemic problems, and this is critical when examining the impact of change in organizations. What is important to realize (and difficult to convey to corporate partners) is that

extended periods of fieldwork generate cumulative results that are impossible to achieve with time compression.

Another temporal issue is the staging of fieldwork. Ethnography has come to be more accepted, and sometimes, even enthusiastically embraced in corporate circles in the past few years, but the common myth is that fieldwork is something you do when all the technical problems have been solved. This “do ethnography later” mentality is widespread, persistent, and powerful. It is a very common managerial and engineering misconception. In a typical negotiation recently with an executive team from a foreign company, we were discussing the question of how to design new technologies for them that would be loved, used, and possibly bought by large numbers of their customers. Their project manager outlined several phases of research and development. Participants agreed that new technologies will need to be designed with a deep understanding of consumers. Then, one of their design engineers proposed initiating early investigations of consumer behavior and preferences through data mining, pattern matching, and sensor data analysis. When we suggested ethnography as another important method for gaining insights into users’ preferences and “pain points,” the presenter responded in an authoritative voice that fieldwork was appropriate for the later phases of the project when prototypes would be tested on users, but “obviously” not for the initial phases in which, he said, the design team would be dealing with “purely technical” issues.

We responded that ethnography is appropriate any time that user and practice data could have an impact on design (see Rogers and Bellotti 1997). We suggested that a common problem in these kinds of research-and-development projects is that ethnographic fieldwork gets started too late to have the optimal effect on design, but it was not clear that the point actually sank in. In this, as in many other cases of preconception, it is important to systematically counteract those notions at every chance we get. And that is a long-term educational effort. On a more pragmatic level, what may be more immediately effective is to point out the large amounts of money that can be saved by not going down the wrong path.

“THIS COSTS TOO MUCH!”

One way to tackle objections to the cost of a project is to investigate what alternatives the manager has in mind. Is he or she thinking about a brainstorming session to get the information he or she wants? Does she or he believe a focus group would provide the needed data? It is useful to generate a discussion at that point to clarify what kinds of questions the manager needs

to answer, what kinds of data would be most useful for that purpose, and the best ways to obtain them.

Another answer to this objection requires taking a look at what the manager is spending money on now. The likely cost of the ethnographic component may be a very small slice of the total project development and implementation cost, yet without it, there is a danger of working on the wrong design or implementing something that is not needed or is not applicable to ongoing work practice and work culture or customer requirements.

Concerns about cost also might be related to concerns about project management. In our experience, managers are genuinely concerned about how to distribute roles and responsibilities, resources, time, and labor effectively. They might say, "It's better to build something now; we can always tweak it later," or "What should my developers be doing while you are gathering requirements?" Often, these reflect a linear but common view of the product-development process that assumes that making architectural or design changes later in the prototype-development process will be less costly than having ethnographers do fieldwork up front.

Where technology development is involved, it is useful to point out the terrific costs that arise when such projects have to be abandoned or when a roll-out does not take. Ideally, engineers and ethnographers will visit field sites together to understand and translate how technology can support everyday activity and working practice (Button and Dourish 1996). By helping abort an ill-conceived program early or by supporting effective implementation and acceptance by employees, work-practice analysis may be crucial to the success of a project or the avoidance of catastrophic failure.

Finally, the issue of cost may sometimes be identified with the issue of scale. One vice president said, "How do we take what you do and scale it across the organization?" (implying that it was too expensive to hire more anthropologists). As a result of such concerns, many corporations have started attempts to transfer ethnographic expertise to their own employees through internal training programs. Much of the work carried out by the Institute for Research on Learning in the 1990s focused on this effort (Bishop et al. 1994; Aronson et al. 1995; Jordan 1996a, 1996b).

One illustration of how ethnographic expertise can be transferred involved a Fortune 500 company. Our team of ethnographers and designers was asked to develop a training program on ethnographic methods for a global sales and service organization. During the training workshops, salespeople learned about approaches such as participant observation and how to elicit and make more visible the practices from the workers' point of view.² After the training period, salespeople began to incorporate what they learned into the sales cycle. The tenor of their conversations with customers

shifted from one that was largely product oriented to one that focuses on understanding and resolving the needs of their customers. The training was deemed a success. It was integrated into the sales cycle at a national level, and later, even adapted for a large consultancy engagement. More importantly, the organization developed a range of ways in which the actual “cost” of fieldwork—the time that company employees spent at customer sites—was either charged to the customer or absorbed by the company.

Whenever the issue of cost comes up, it is best to go back to a common understanding with our corporate colleagues about what ethnographic work-practice analysis is supposed to accomplish and then discuss what happens if ethnographic work is not carried out. To do this, it is critical to identify early on the scope of managers’ concerns and to address each appropriately. We have found it especially helpful to use examples from past projects in which we can track the time and resources that were spent on designing change programs built on insufficient understanding of the realities on the ground. From there, it may be possible to propose ways in which the whole team can be engaged in better understanding the subtleties of their customer’s business, which are crucial to a successful adoption of new technologies and organizational changes. To sum up, concerns about costs can be addressed productively at both the tactical and strategic levels.

**“DON’T BOTHER. WE CAN DO THIS FASTER
AND CHEAPER WITH MARKET RESEARCH
AND FOCUS GROUPS!”**

Sometimes, managers perceive ethnographic research and market research to be competitive alternatives. We would argue that they are, in fact, complementary. Market research is primarily concerned with making business decisions and forecasting the size of the market. Ethnography is concerned with design decisions that are based on a true understanding of users’ needs. For market research, you have to have knowledge of the product or service that is to be marketed to make predictions about how well it will sell. Ethnography can handle the more diffuse problem of how ill-defined products and products that are still unrealized can be designed in such a way that they will fit into the day-to-day environment of future users. For an ethnographic investigation, you need to be able to identify the user population, but you do not need to have a clearly defined product or know what kinds of questions you want to ask. Ethnography is a discovery science, not a validating one (see Whalen and Whalen 2004).

David Gilmore has written an excellent article in which he points out that one thing to avoid is a territorial war with market research. The best strategy is to explain that the two study the same thing (potential users) but pursue complementary goals (Gilmore 2002). Under the best of circumstances, ethnographic research and market research would be strongly allied, providing information for the projected life course of a product, from its early conceptualization in the heads and labs of researchers and designers through the stages in which it “gets a life” in peoples’ workspaces and lifespaces. Different stages would need different concentrations of effort from the two sides, but both should be involved in the decision-making process in product design.

Focus groups, together with informal brainstorming and surveys, are probably the method used most commonly for gathering data about opinions and attitudes, product features and brands, design specifications, pricing models, buying preferences, and problem solutions. They are very attractive to managers and for very good reasons. For one thing, they constitute a well-recognized methodology for collecting data about users or other groups of people of interest to the company. For another, you can out-source them as a limited one-shot effort with a predictable schedule. And they are comparatively inexpensive.³

Recently, however, market researchers themselves have become increasingly disillusioned with focus groups. In business journals, you now find articles entitled “Shoot the Focus Group” (Kiley 2005), the chief marketing officer at Yahoo says “I’m killing all our focus groups,” and the Marketing Science Institute itself recognizes that the focus group is no longer the pre-eminent methodological tool for qualitative inquiry (McCracken 2005).

A main problem with using focus groups as the primary means to understand users is that focus-group data are subject to all the limitations of semantic data, that is, data collected in response to questions or as parts of conversations. It is well known not only to ethnographers but also to lawyers and physicians that accounts of events and the actual events often do not coincide. The problem is not that people are deliberately deceptive but rather that memory is notoriously unreliable. There is good documentation that shows that about one-third to one-half of what people remember is factually incorrect (Bernard et al. 1984).

The picture is no rosier when we look at people’s intentions about the future and what they actually do. Decades ago, W. F. Whyte’s (1993) urban research showed that urban people say they want quiet plazas for their lunchtime break, yet when you watch what they actually do, they frequent busy places. Contemporary opinion polls show that people place high value on fuel economy, yet they drive gas-guzzling SUVs. The lesson: Do not equate attitudes with actions.

FIGURE 1
A Two-by-Two Data Matrix

	Qu 1	Qu 2	
	Say	Do	
	Qu 3	Qu 4	
Say/Do/Think/Feel	Think	Feel	Two-by-Two6

NOTE: This particular version was proposed by David Kelley of IDEO to the World Economics Forum in 2001, but many others of this type are in circulation.

The business press is full of stories about the limitations of focus-group research, indicating a rising consciousness among managers that focus-group data are not as unproblematic as previously thought. Famous examples of false positives that led to unsuccessful product launches include the disaster Pepsi Cola Company experienced with marketing Pepsi Edge on the basis of focus group data.

Quite frequently, managers are unaware of the distinction between focus groups and ethnographic research and think of them as interchangeable, a situation that is exacerbated by the fact that many ethnographers use focus groups as part of their repertoire of methods. However, focus groups may not give our managers the data they are looking for. Focus groups do certain things well but not others. Look at Figure 1, which shows a two-by-two matrix that considers what people say, do, think, and feel.

Data in quadrants two, three, and four are subtle, tacit, implicit, and context dependent and are not surfaced easily by guided talk with strangers in situations that are subject to the constraints of civilized conversation, as they are

in focus groups. Focus groups do better with quadrant one. They get at what people say they want. If that is what managers are looking for, this is the right methodology. Often, however, that is not what managers want. They are really interested in what people actually do in the workplace, or really think about a product, or maybe most importantly, how they feel about it. These managers need to be helped to make explicit whether they want to know what people say they do or what they actually do in the workplace, at home, in the supermarket, or the polling booth. Managers should not assume that people's expressed opinions reflect their actual practices.

The fit between what people say and what they do is an empirical question that always requires at least some observation. Failure to pay attention to the say/do distinction is rarely questioned in focus-group research, yet it is likely to produce data that are invalid in the technical sense (i.e., data that do not measure what we intend to measure).

Some approximation to do/think/feel can be achieved by providing appropriate stimuli during group sessions, such as examples of the product, video clips of use situations, role plays, and the like. Still, the most valid data are clearly obtained by studying those issues in context, within the situations in which they play out—precisely the strength of ethnographic fieldwork. This means that methods such as focus groups, surveys, formal interviews, and expert testimony may not be sufficient, since all of them pull people out of their actual context and ask them to explain tacit knowledge that is not normally or easily articulated. If the first quadrant is the realm of focus groups, the second and third definitely belong to ethnographic work-practice analysis, as, actually, does the fourth (although that would require an extended field engagement).

“COULDN'T I JUST GO MYSELF AND WATCH FOR A WHILE?”

Sometimes, managers are convinced they know what goes on in the workplace—otherwise, they wouldn't be managers. There may be skepticism that ethnography does not provide anything different than what can be discovered by simply speaking with people. The problem here is that ethnography looks and sounds straightforward, but as any card-carrying anthropologist knows, it is not. Some years ago, Diana Forsythe wrote a beautiful article on that topic, titled “It's Just Common Sense!” in which she discussed the “tendency for social and communicative work to be rendered invisible in technical settings” (Forsythe 2001:162). She pointed out that because ethnography looks easy, people assume that there is nothing to

it and that anybody can do it. Actually, it requires years of theoretically grounded training and practical experience and involves systematic data collection and rigorous analysis—almost all of which are invisible to the person casually observing an ethnographer at work.

What managers often do not understand is that one of the special skills trained ethnographers bring to the field is the ability to look for known patterns (sometimes based on theory, sometimes on prior experience) while at the same time being constantly alert to interesting observations that might herald a new pattern (or simply an idiosyncratic occurrence). This is a balancing act, attained only after years of training and experience that can make ethnographic work extremely powerful in complex situations in which “anything goes” and nobody has a good handle on what moves the system.

Closely allied to the idea that anybody can do ethnographic observations is another mind bug, the idea that if you want to know what people do, you could just ask them. It’s only a matter of common sense, right? Yet, anybody who has tried this has found that people’s memories get distorted (often in the direction of cultural biases) and that they have little access to the dynamics of their own work practices. Diana Forsythe pointed out that just as medical diagnosis is not just talking with patients, so ethnographic fieldwork is not just talking with people in the workplace. In actuality, ethnography often runs counter to common knowledge because it requires tapping into what people take for granted about their work, and thus, do not ordinarily discuss.

Part of the challenge, therefore, is how to successfully communicate the quality of work-practice analysis as it unfolds through time, especially in situations in which managers are more accustomed to assessing the amount and quality of code that is written than the perspectives that ethnographers bring to a team. Once, when reviewing an ethnographer’s performance appraisal, a relatively new manager was stumped: “How can I judge what the quality of your work is—that you are doing ‘good analysis’ rather than ‘bad analysis’? How do I know that I wouldn’t have found out the same things myself at the customer site?” As a computer scientist, the manager was used to evaluating the work of engineers by looking directly at the software code to determine whether the code was well written and the prototype robust, whereas she felt that without domain knowledge of cultural anthropology, an evaluation would be difficult. Another individual who had been promoted as a first-time manager asked why one needed an ethnographer’s skills when he could just as easily read market research reports and put together customer requirements himself.

Unfortunately, there is no easy fix to the conviction that this is easy and anybody can do it. What is required is an educational effort that makes clear that what looks like “just talking” or “just hanging out with those guys” is

part of a rigorous methodology that worries about such things as validity and reliability and sample size and rival hypotheses. More critically, it entails accruing layers of experience through time by observing a range of phenomena across different sites. It may be overstating the case, but we like to tell our managers that competent ethnographic work is only 5% visible fieldwork versus 95% preparation, analysis, synthesis, and communication, most of which are invisible to outsiders.

Whether or not these differences of opinion reflect the disjunction between positivist and relativist paradigms (Salomon 1991; Forsythe 2001; Maxwell 2004), they embody significant lessons in the articulation of ethnographic work itself. From them, we have been learning the following.

First, whenever possible, we try to align with projects that are critically important to the company as well as with managers who are customer focused rather than purely technology focused. Managers responsible for such projects often consider the viewpoints and knowledge of customers as strategically important and may use our results to cross-pollinate ideas and decisions across business groups.

Second, we try to get people who understand the value of ethnographic work to vouch for our contributions and the results that they have personally seen. If this is in writing—whether in an informal e-mail or a more formal memo—so much the better. It is a powerful testimony to the value of ethnographic work to have a project manager for a \$200 million project send out a technical-specifications document with a cover note that reads, “Here is the information [we] provided to [the third-party developer]. It is filled with specific examples of how the work-practice study enabled us to develop the tech spec. Take your pick!”

Third, we try to involve managers in field visits before and after the technology is installed or the work practices have changed at a customer site. If managers do not see the realities on the ground, they often come to believe that they are faced with naive users who use an application incorrectly, resist a technological innovation, or simply carry out their work in the wrong way. When that is the case, field visits often help us to convince them that, as one manager said, we need to “change the maze, not the rat.”

After participating in site visits and speaking directly with customers or workers, managers are more fluent in picking up the nuances of work-practice studies, and more importantly, in representing ethnographic work to their own management or external customers. We also find it helpful to solicit managers’ opinions about ethnographic contributions to the project. These can help track what the early assumptions were about the customer and the technology before field data were collected and can compare those assumptions to actual findings from the field: “This is what we thought initially; this

is what we actually found; and here is what was incorporated into the final recommendations.”

“YOU CAN’T GENERALIZE FROM THIS!”

One of the most deep-reaching objections we encounter comes from managers’ legitimate concerns about enterprise-wide solutions. They often want what they call (without irony) “cookie-cutter solutions,” maintaining that global solutions should not be based on fieldwork with only one or two customers. They fear that the results of a field study cannot be generalized.

To some extent, this is true. But in these situations, it makes sense to find out what universe the individual actually needs to generalize to, which often turns out to be quite limited (e.g., all of the company’s call centers, or all call centers within the industry, maybe including competitors). This narrows the problem from a universal intractable one to one that we may well be able to deal with.

One basic strategy for making our findings applicable beyond the one site where we are doing fieldwork is to provide evidence that the findings apply beyond the field site. This evidence may come from a number of different sources, ranging from our own prior experience to tapping into local knowledge about the distribution of the phenomena of interest, doing literature triangulation, drawing on the ethnographic professional community to find similar cases, and carrying out “ethnographic probes”—brief additional investigations.

For an experienced fieldworker, there is hardly a situation that does not call up a similar one. When we begin work at a particular site, there is a very high likelihood that we have worked in similar sites before. As a consequence, we walk in with a set of hunches, hypotheses, and theories about what might be the case in the present situation, depending on how familiar we are with this type of workplace. For example, in a company that has a central headquarters where policy is made and a set of field organizations that work with clients on a day-by-day basis or carry out the actual product production, we have a pretty good hunch that there will be communication issues between headquarters and the field. Typically, in such situations, headquarters knows very little about the ways in which its directives actually impinge on the field organization, and field organizations have little opportunity to communicate with headquarters in a meaningful way. Documenting this situation in one field site could be made to generalize with very little work, at least to the whole company, if not beyond.⁴

In one such situation, we observed work practices in a field site that, atypically, had just one employee. This person had to handle a wide variety

of production requests. When we documented the workarounds she had invented and used on a routine basis and checked them in several other field sites that were also run by solitary employees, we had a good argument that there was a company-wide problem. We presented our observations of the single site and the testimony of the other solitary employees to the managers responsible for the rollout of programs designed for multiemployee sites. They immediately saw that what we described in detail for the one site was true company-wide and instituted changes in the rollout of new programs.

So, one way to argue for generalizability is to rely on company-internal testimony. We make it a routine feature of questions we ask at the field site to always inquire about typicality. We might say: "You just did this. Do most people here do it that way, too?" Or, "Who else does it that way? What about other parts of the company? How do they handle this problem?" Similarly, if we have videotapes of work practices and the opportunity to coanalyze those with employee participants (Cefkin and Jordan 1994; Ruhleder and Jordan 1998), we always ask about typicality and combine these data with what we observe in the field. Emerging lines of generalizability are then investigated and corroborated by interviewing multiple people within teams and across organizations.

What this provides is a variety of perspectives and practices that, when synthesized, demonstrate critical points of convergence and differentiation. Given that we really are not only interested in the personal views and behaviors of an individual worker or manager (or field site, for that matter) but are looking for patterns, preferably widespread patterns, such questions come naturally and make a lot of sense to the employees (who, after all, have been assured that we are not seeking personal information but are searching for ways in which we can identify company-wide problems and help solve them). Interestingly, in a codesign session, some patterns need remarkably few probes before it is clear whether something is common practice or not.

Having laid out the patterns in one work site of a company, it may be possible to generalize about the whole very efficiently by making a few short visits to one or a few other sites and ascertaining that the pattern holds. How easy or complex this approach becomes depends on the degree of variability in the system. There may be more or less difference in how work sites within the same company, exposed to the same company policies and directives, carry out their work. The greater the variability, the longer the ethnographer will have to work to understand how many different patterns there are and how they vary.⁵

An obvious avenue for arguing for the general applicability of our fieldwork is to do a literature review to establish what other researchers have found in similar situations. For example, when we did an intensive, video-based

investigation of the effect of technology-generated delays on collaboration and trust in video-supported, geographically separated teams, complementary data from laboratory research in psychology suggested that our findings are widely generalizable (Ruhleder and Jordan 2001). A literature review is also a good way to increase the range of hypotheses that one brings into the field.

We have found a particularly good source of material for typicality arguments in the interchanges we have with colleagues. They may have done similar studies in workplaces that were not exactly the same but exhibited the same phenomenon. For example, in an airlines operations (ops) room, the workers had a curious habit of talking out loud, saying things to nobody in particular. Given that talking to yourself at that time was still highly censored in normal interaction (with cell phones, it has become acceptable now), we wondered about the significance of this. In the ops room, work stations were positioned facing the walls, so that operators were working back to back. We hypothesized that the positioning of their computers generated this unusual communication channel. It turned out that colleagues doing fieldwork in an ops room of the London metro had documented the same phenomenon, and so had workers in other kinds of control rooms in which operators were copresent but not in direct face-to-face communication (Heath and Luff 1991). We eventually were able to generalize our findings to many other workplaces of this kind. One could imagine that a new technology that gets developed for one of these settings with the purpose of improving the mutual availability of knowledge about the state of operations might be a good candidate for improving knowledge flow at other sites of the same type.

“YOU CAN’T QUANTIFY THIS!”

This is a challenge similar to the preceding one about generalizing. Again, it may be useful to try to understand what managers or colleagues wish to measure. There are many aspects of our work that are easily quantifiable from our systematic records. Questions about how long something takes, knowing precisely how many of something (persons, technologies and artifacts, documents) are present or used in a situation, how many of certain kinds of actions are performed, how often certain kinds of events occur (such as interruptions, stalls, encouragement, collaboration)—all of these are easily retrieved from our records, especially from tape and video recordings.

But frequently, what managers are really asking for is a quantification of the effect that our study-plus-intervention might have on performance and return on investment (ROI). This a much more difficult challenge because

it requires acknowledging that, realistically, there is no scientific methodology that could specify that effect with precision. The reason for that is that organizations are not laboratories in which one can set up an experiment with specified variables and watch the effect behind a one-way mirror. Rather, they are living, complex systems that are constantly undergoing change, where in addition to our intervention, many different internal and external variables affect the outcome. So even if, let's say, employee morale, customer satisfaction, and the bottom line improve after our intervention, there are always rival hypotheses that could explain that, such as changes in the market, competition, a new CEO, and so on.

Nevertheless, there are usually some measures we can cite that assess the effect of our work. These measures are always specific to the scope and goals of the particular engagement and need to be chosen with care. For example, working with systems engineers, we might be able to specify how many bugs were detected (and hopefully fixed) during our work, how many recommendations we made, how many of those were actually implemented, what new market opportunities we identified, and the like.

In addition, there are probably some hard numbers the company collects anyway that could reasonably be cited as at least influenced by our work. For example, in a project that provides new learning opportunities for employees, employee satisfaction scores might rise substantially, employee turnover might decline, and customer satisfaction might improve. Supported by employee and customer testimony, it might well be possible to argue that these quantitative improvements are, at least in part, because of our intervention. Similarly, our support in the introduction of a new technology might lead to more rapid and thorough adoption than otherwise would be the case. The usefulness of these kinds of data and the amount of energy that should be spent on collecting them is something that needs to be negotiated with sponsors up front and constantly revised as new possibilities arise in the course of ethnographic work. Unfortunately, the natural tendency is to think about these issues at the end of a project, when one tries to sum up the results. We cannot overemphasize the importance of establishing up front what data need to be collected to be able to argue later on for ROI as well as for performance and productivity effects.

“THIS ISN'T SCIENTIFIC!”

Unfortunately, when someone voices this kind of objection, it is usually an indication that he or she intends to categorically dismiss anything we might propose. This is a tough one. Countering this objection means persuading the

person that there are different ways of doing science, each of which is appropriate for different kinds of situations and all of which can produce valid kinds of insights. We need to argue that although a work-practice study is different from what he or she might have in mind, we are proposing a systematic, rigorous approach to data collection and analysis that produces more useful results in the situations of interest than what he or she might believe is the standard scientific method.

What most managers are likely to mean by the *scientific method* is the deductive, hypothesis-testing approach to research that most of us were trained in and that is still the dominant research methodology in the fields with which managers are likely to be familiar. Managers tend to aspire to the model dominant in the natural sciences that equates scientific research with randomized controlled trials, double-blinding, and quantitative analysis. We might call this approach the analytic paradigm (Salomon 1991).

The analytic paradigm is the approach of choice in situations in which we have well-developed theories regarding the phenomena in which we are interested. From these theories, we deduce specific hypotheses that are used to test relationships between discrete, well-defined, operationalizable variables. Statistical hypothesis testing then provides a systematic procedure for ruling out rival hypotheses. It is used to good advantage in the natural sciences and in such fields as experimental psychology, pharmacology, agriculture, or ergonomics, in which researchers know ahead of time what variables they want to test. The outcome of research here is a statement of statistical significance of the difference between the variables tested.

At this point, a manager who is actually listening may realize that the analytic paradigm is not well suited to study what happens in the complex workplaces of which he or she is in charge. He or she may see quickly that most situations of interest in workplace research are very messy, dynamic situations in which little can be held constant. Here, trustworthy theories and well-defined variables are hard to pin down. The manager knows deep down that we deal not with laboratory situations but with dynamic, organism-like systems that do not remain in equilibrium long enough to be investigated but are in continual flux, undergoing constant self-reorganization even when there is no systematic, planned-from-the-outside change effort.

An alternative that can be proposed at that juncture is to work from an inductive systemic paradigm (Salomon 1991), an approach that has been used for centuries by naturalists, ethologists, and other investigators of natural dynamic systems—from Vesalius, who decided in the thirteenth century not to consult books with Greek and Latin theories but to actually observe how the human body works, to Galileo and Copernicus, who abandoned metaphysical dogma and observed the patterns in the sky, to naturalists such as

Charles Darwin, who observed finches on the Galapagos Islands, and Jane Goodall, who revolutionized our thinking about primates by observing chimpanzees in the wild.

The systemic approach looks at how different parts of a system work together. The researcher here is not an unobtrusive outside observer but an active participant, immersed in the situation. The central question is not What is the relationship between these two variables? The central question is How does this system work? There is no predetermined research protocol, but rather, we look for patterns of behavior, patterns of spatial arrangements, tool use, collaboration, communication, and organizational structuring. We check out where these patterns hold. We ask ourselves questions such as how typical this is, who else does it this way, and what alternatives there are to solve this problem—and that leads to a very rich characterization of how the system works rather than to a statement of the relationship between variables (Whalen et al. 2004).

So, in the work we do, we proceed empirically. We watch and listen. We look at what people say and do in space and time. Sometimes, we have strong intuitions based on past experience that we investigate—but we always look for patterns, for exceptions to those patterns, and then for the reasons for those exceptions that allow us to state the pattern more strongly. And we attempt to build generalizations, findings that hold across several venues.

Some of us would call this science. But to many of us, the is-it-science question is irrelevant. It becomes relevant, however, when it is an issue for our managers. In that case, pointing to the difference between the analytic and the systemic paradigms, between the deductive and inductive approaches, becomes worthwhile.

“WHAT KINDS OF RESULTS CAN YOU GIVE ME?”

The question of what kinds of results we can deliver always deserves to be taken seriously because it may indicate that expectations of managers and ethnographers are not entirely aligned, or at least, are not explicit enough. The question takes many forms, depending on the scope and nature of the interaction between ethnographers and researchers on one hand and company executives, engineers, and project managers on the other.

Often, the potential funders' inquiry is actually about ROI, an issue we dealt with above. Alternatively, we may be asked if we can help increase the productivity of the organization. Product designers may press for customer requirements, and software engineers may ask how they can use ethnographic information for writing better software.

No matter what lies behind the question, we need to ask ourselves and our colleagues to spell out what results we are looking for. For example, when managers ask ethnographers to help them improve productivity, the first issue that needs clarification is what they mean by *productivity*. We have been in situations in which we have been asked to help a client achieve “a double-digit increase in productivity by year’s end,” although whether that meant reduction in working hours, increase in product output, improvements in quality, or something else was by no means obvious. Nor was it clear how those increases would be measured. In virtually all cases, companies already have a variety of metrics in place, and although those may not be optimal, it is important to know about them, and if appropriate, to use them.

The first step, then, is always to find out what is meant by *productivity* or *customer satisfaction* or whatever managers are concerned about; the second is to ascertain how they currently measure that; the third is to come to a shared understanding of how it could be measured within the context of our work; and the fourth is to honestly discuss with our counterparts the fact that outcomes may be influenced by factors beyond the ones we consider in our project, such as changes in the marketplace, shrinking customer bases, or new competitive products.

Situations in which there is no clear answer to the results question are particularly difficult because results may be unpredictable (as is often the case in ethnographic investigations). Here is a case in point. The manager of a new intellectual-property laboratory was deeply concerned about the performance of the lab and asked the in-house ethnographer reporting to him to help increase the efficiency of the organization as a whole and to identify any bottlenecks in the laboratory’s work processes. This was not an easy task. The manager needed regular status updates for his monthly reports and was worried that engineers were not uploading enough material into the databases for patent investigations. Did this mean that engineers were having difficulties with their work or not working hard enough? Where were things breaking down? We discovered after extensive interviews that engineers were loath to put work-in-progress results into the database until attorneys had rendered their final opinion on patent claims. Why? Because attorneys often rejected test results or requested new tests because, as they said, engineers had viewed claims in a technical rather than legal manner. Based on our recommendations, attorneys were assigned to the laboratory much earlier in the investigative process, which dramatically improved productivity.

Furthermore, when we examined the daily activities of hardware and software engineers in the laboratory, we discovered that communication between hardware and software engineers was meager. As a consequence,

tests and analyses were being repeated unnecessarily. In patent litigation, time is indeed of the essence, and delays could mean the loss of millions of dollars. So we suggested ways in which teams could organize, track, refine, and share their extensive search results such that databases could be used to assist rather than hinder work practices. These changes were widely adopted and not only increased the number of submissions but sparked new collaborations and an effective reorganization of the teams themselves.

The point is that none of these results could have been anticipated when the project began. In cases in which the results of an ethnographic analysis are hard to predict, it is best to openly acknowledge that while at the same time citing successes with similar projects.

Nevertheless, it is important not to read too much into the “what kinds of results can you give me” question. Now and again, people only want to know when and in what form they can expect to see the outcome of our investigations. Typically, managers expect to see a written final report and/or a PowerPoint presentation that delivers the results, but experience has taught us that this may not be optimal either for them or for us.

Although managers probably expect a report, what they really need is help in generating a change in the work practices and process flows that have prevented optimal results so far. But reports tend to end up on some executive’s office credenza, and PowerPoints tend to get cascaded down the company, after which people go back to business as usual and behavior as before.

We have developed two strategies that increase the chances that the actionable results we deliver are actually acted on and do not (only) end up on the credenza. The first consists of working sessions with all levels of stakeholders throughout the project, preferably starting early on. Depending on the company and the relationships already established, these may be very short and informal or more elaborate, maybe scheduled once a week or so but always carried out in a collaborative vein. We have found it very useful, for example, to invite frontline workers to weekly video pizza lunches where we would ask them to review with us pieces of the videotapes we had shot during the week that we had found interesting for some reason or other (Cefkin and Jordan 1994). Not only did we gain deep insights into what work processes and company strategies look like from the point of view of the workers, but we also built close relationships with them, which helped immensely in grasping the local realities.

In addition, these sessions empowered the workers in that they gave them a shared view on what they were doing and how they dealt with some of the problems that arose on the floor. Even early on, it usually makes sense to ask people at various levels of the organization: “Here is what we

observe. Is that what you see? How common is that? Is it a problem? Is this a practice to be protected as changes are being made?" In other words, we present ourselves as more or less naive (but interested) learners, so that these interactions are less feedback sessions in which we report authoritative findings than collaborative working sessions in which a common understanding of what the issues are is built up.

The second strategy gets around the archiving of our reports, the *credenza* phenomenon, by doing some version of the famous "Two-by-Two-by-Two-by-Two (2×2×2×2) Method" that was originally developed at IRL, the Institute for Research on Learning. It involves two influential people from the host company (maybe an executive with a peer or subordinate of his or her choice) meeting with two people from the project team. They commit 2 hours to a discussion of a particular, implicative finding from the report. The discussion centers on what possible change could be made that would lead to better outcomes and what steps the hosts are going to take in the next 2 weeks to move those toward reality. The group meets again after 2 weeks, assesses progress, deals with issues that have come up, identifies success or failure and what it has learned from that, and repeats the process with another intervention suggested in the report. The advantage of this method is that it generates lively discussion among the executives as well as the writers of the report during the intervening 2 weeks.

CONCLUSION: CHANGING THE SYSTEM

To get back to the title of this article, how can we move from the frustrations of forever having to counter our managers' reservations toward productive encounters with them that capitalize on collaboration and mutual appreciation of our strengths and weaknesses? A commitment to this goal would require a long-range plan to change the climate of the company and to actively shape fundamental attitudes and expectations.

There are various strategies for making inroads on entrenched positions. For example, we have had some success with giving company-internal workshops on ethnographic methods that get participants to verbalize ideas about when ethnography might be useful for solving company problems. This takes the issue away from the loaded interaction with an individual manager's personal and professional interests and builds wider support. One might also offer a reading group on technology topics (or join an existing one) by introducing articles that include successful applications of ethnographic methods. Looking at articles that misapply quantitative, hypothesis-testing,

or data-mining approaches naturally generates an interest in alternatives. The important point is not to get into territorial battles by insisting on the superiority of ethnographic analysis but to always pull the discussion back to the adequacy of different methods for particular situations and to get people to see ethnographic work as one of a number of alternatives. What this also implies, of course, is that we are clear that there are many issues in corporate life in which ethnography does not deliver the right data, and other kinds of research methods are more appropriate.

One major difference between workplace ethnography and traditional ethnographic research is that we now routinely work as part of interdisciplinary teams. It is those team members who can become the most effective advocates of the benefits of ethnographic approaches in the company as a whole. And where managers are attuned to hear findings directly from technical experts, the systems developers, customer account managers, market analysts, and computer scientists we work with may contribute substantially to a change of attitude in the company.

There are many ways to change the climate at a company. Admittedly, all of them are hard and slow. All of them are about C-questions: collaboration, compromise, coexperiencing; educating managers, technologists, and administrators and getting educated ourselves. It is a long-term effort. Some of us see it as trench warfare. Others of us prefer to see it as more or less friendly but always respectful, persuasive encounters. Both views have their consequences in what they generate on the other side. On the day-to-day level, this work is challenging, to be sure, but also deeply satisfying. One thing we anthropologists know is that culture change takes time.

A systemic change strategy requires a commitment to learn and to educate on both sides and considerable inventiveness in identifying and undermining the existing barriers. It requires initiative as well as taking advantage of opportunities as they arise to talk about a particularly useful ethnographic study, to introduce visiting colleagues who do ethnographic research, to point to the advantages ethnographic work would have conferred in projects that failed, and to put ethnographic methodology into a framework within which its contributions (and inadequacies) can be highlighted. To consistently and insistently point out what difference an ethnographic study would have made in projects that should have included an ethnographic component but did not—such is the long-range work of getting to know each other. And let us again emphasize that we, the ethnographers, need to take an active role and exhibit a positive attitude toward understanding the realities of the managerial and technical work with which our efforts have to integrate. As they say, it's a two-way street!

NOTES

1. When anthropologists studied tribal communities, the tradition was a full year of field-work. There were good reasons for that—most importantly, the necessity to observe the agricultural cycle, preferably with a return some years hence. In today's fast-paced world, the activity cycles of workplaces and communities of practice tend to be much shorter, justifying briefer periods in the field.

2. For a deeper discussion on making work visible, see Suchman (1995), Button and Dourish (1996), and Star and Strauss (1999).

3. It is actually a misperception that outsourced focus groups are low cost. In fact, focus groups are quite costly when done by professionals qualified to conduct them. In addition to planning and moderating the panels, focus groups also involve recruiting people to fit appropriate sample requirements; creating questions about product features, competition, and brand identity; preparing collaterals such as discussion guides, product concept descriptions, and videos; and committing a substantial effort to analysis and presentation of findings. This, if done rigorously, is not cheap.

4. Of course, we would want to introduce a number of safeguards that would reassure us that we have not run into the one atypical case in the whole company.

5. This is one reason why ethnographers are often reluctant to indicate a specific length of time for a field study.

REFERENCES

- Aronson, M., L. Bishop, M. Cefkin, B. Jordan, N. Lawrence, L. Sullivan, C. Preston, and J. Oesterle. 1995. Reflections on a journey of transformation: Learning, growth, and change at Xerox Business Services. IRL/XBS Systemic Assessment Project, Final Report. Palo Alto, CA: Institute for Research on Learning.
- Bernard, H. R., P. Killworth, D. Kronenfeld, and L. Sailer. 1984. The problem of informant accuracy: The validity of retrospective data. *Annual Review of Anthropology* 13:495–517.
- Bishop, L., M. Cefkin, W. Clancey, B. Jordan, J. Oesterle, and B. Tudor. 1994. A learning organization in the making: A report on the workpractice and design project at the Xerox CAC in Dallas. Project report. Palo Alto, CA: Institute for Research on Learning.
- Button, G., and P. Dourish. 1996. Technomethodology: Paradoxes and possibilities. *Proceedings of CHI 1996*:19–26.
- Cefkin, M., and B. Jordan. 1994. Video-based interaction analysis in the workplace: A tool for workers' appropriation of their own practice. Proceedings of PDC'94, the Participatory Design Conference. Chapel Hill, NC: Computer Professionals for Social Responsibility (CPSR).
- Forsythe, D. 2001. *Studying those who study us: An anthropologist in the world of artificial intelligence*. Stanford, CA: Stanford University Press.
- Gilmore, D. 2002. Understanding and overcoming resistance to ethnographic design research. *Interactions* IX (3): 29–35.
- Heath, C., and P. Luff. 1991. Collaborative activity and technological design: Task coordination in London underground control rooms. *Proceedings of the Second European Conference on Computer-Supported Cooperative Work*. Amsterdam: Kluwer.
- Jordan, B. 1996a. Ethnographic workplace studies and computer supported cooperative work. In *The design of computer-supported cooperative work and groupware systems*, edited by

- D. Shapiro, M. Tauber, and R. Traunmüller, 17–42. Amsterdam: North Holland/Elsevier Science. Accessed July 10, 2006, at www.lifescapes.org/Papers/ethnographic_workplace_studies.htm.
- . 1996b. From training to learning in the new economy. Symposium on Education and Employability, organized by RHODIA (Rhone Poulenc) and the Catholic University of São Paulo, Brazil. Itu, Brazil, August 7–9. Accessed July 10, 2006, at www.lifescapes.org/Papers/0212_from_training_to_learning.htm.
- Kiley, D. 2005. Shoot the focus group. *Business Week*, November 14, pp. 120–21.
- Maxwell, J. A. 2004. Using qualitative methods for causal explanation. *Field Methods* 16 (3): 243–64.
- McCracken, G. 2005. <http://www.cultureby.com/trilogy/>, accessed May 4 and 5, 2006.
- Rogers, Y., and V. Bellotti. 1997. Grounding blue-sky research: How can ethnography help? *Interactions* 4 (3): 58–63.
- Ruhleder, K., and B. Jordan. 1998. Video-Based Interaction Analysis (VBIA) in distributed settings: A tool for analyzing multiple-site, technology-supported interactions. *Proceedings of the Participatory Design Conference*, edited by R. Chatfield, S. Kuhn, M. Muller, 195–96. Seattle, Washington, November 12–14.
- . 2001. Co-constructing non-mutual realities: Delay-generated trouble in distributed interaction. *Journal of Computer Supported Cooperative Work* 10 (1): 113–38. Accessed July 10, 2006, at www.lifescapes.org/Papers/delay_for_jcscw_2001.htm.
- Salomon, G. 1991. Transcending the qualitative-quantitative debate: The analytic and systemic approaches to educational research. *Educational Researcher* 20 (6): 10–18.
- Star, S. L., and A. Strauss. 1999. Layers of silence, arenas of voice: The ecology of visible and invisible work. *Journal of Computer Supported Cooperative Work* 8 (9): 9–30.
- Suchman, L. 1995. Making work visible. *Communications of the ACM* 38 (9): 56–64.
- Whalen, M., J. Whalen, R. Moore, G. Raymond, M. Szymanski, and E. Vinkhuyzen. 2004. Studying workspaces. In *Discourse and technology: Multimodal discourse analysis*, edited by P. LeVine and R. Scollon, 208–29. Washington, DC: Georgetown University Press.
- Whalen, M., and J. Whalen. 2004. Studying workspaces. In *Discourse and technology: multimodal discourse analysis*, edited by Philip LeVine and Ron Scollon, 208–29. Washington, DC: Georgetown University Press.
- Whyte, W. F. 1993. *Street corner society: The social structure of an Italian slum*. 4th ed. Chicago: The University of Chicago Press.

BRIGITTE JORDAN has carried out ethnographic research for more than 30 years in academic and corporate settings. Now a freelance consultant, she also holds the position of consulting principal scientist at the Palo Alto Research Center, Palo Alto, California. Some of her publications include "Assessment as Practice: Notes on Measures, Tests and Targets" with Peter Putz (Human Organization, 2004); "Corporate Anthropology in Costa Rica: A Brief Research Report" (Anthropology News, 2003); and "Co-Constructing Non-Mutual Realities: Delay-Generated Trouble in Distributed Interaction" with Karen Ruhleder (Journal of Computer Supported Cooperative Work, 2001). Many of her publications can be found at her Web site, www.lifescapes.org. She can be reached at jordan@akamail.com.

BRINDA DALAL is a research scientist at the Palo Alto Research Center (PARC), Palo Alto, California. A corporate anthropologist, she has a deep interest and extensive experience in applying ethnographic skills to the design of technology innovations and

organizational change. She cofounded the Cleantech initiative at PARC. Brinda received a PhD and MPhil in social anthropology from the University of Cambridge. Before that, she worked with an organization in Bombay, India, focusing on microcredit, health, and low-cost housing schemes for urban migrants. Brinda may be contacted at bdalal@parc.com. She has presented “The Baker’s Dozen: The Presence of the Gift in Service Encounters” (Ethnographic Praxis in Industry conference, American Anthropological Association, 2005) and “What a To-Do: Studies of Task Management towards the Design of a Personal Task List Manager” (ACM Conference on Human Factors in Computing Systems, CHI2004, 2004).