Persuasive Encounters: Ethnography in the Corporation

Based on a paper prepared for the National Meetings of the Society for Applied Anthropology
Portland, OR, March 19-23, 2003

Brigitte Jordan and Brinda Dalal
Palo Alto Research Center

Unpublished Draft
Last Changes: Saturday, December 24, 2005 by gj

Abstract In corporate settings, ethnographic methods are routinely challenged by managers who confront ethnographers with a set of typical objections that question the validity and effectiveness of ethnographically based findings and recommendations. This paper offers a series of steps towards 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 world views 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 in order to re-situate our practice.[1]

Keywords: corporate ethnography, workpractice analysis, workscapes,

Introduction

In this paper we make an attempt to deal with the daily reality of ethnographers doing work in a corporate world where they persistently face questions about the validity and effectiveness of their methods and results from their managers. Sometimes those questions seek information, sometimes they are hostile, sometimes they come in the form of a sniper attack or a deprecatory

[1] We expect to find our audience both in academic and corporate settings: in the first, students and teachers of what has been variously called business, industrial or corporate anthropology; and in the second, corporate ethnographers, middle managers and corporate decision-makers. While ethnographic methods have come to be widely accepted in companies large and small, this paper is also aimed at those working where the techniques are less well known to management. The issues we discuss are increasingly lively in corporate operations, engineering, and industrial research labs, many of which are struggling (as are we) with methodological questions and the legitimacy of ethnography. What is distinct in the corporate domain is the potential realm of influence: ethnographers are hired not just for their academic research skills, but are responsible for customer engagements that understand and shape change itself.

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 into a productive response, a persuasive (en)counter? Managers, especially inexperienced ones, often feel that they are exposing themselves to unknown risk. They may have nonverbalized questions, sometimes as basic as: Does this work at all? Am I running the risk of exposing myself and my people to criticism that I can avoid with more true and tried methods? And, the always fundamental one: What is the financial impact? They also may be struggling with questions about having full-time ethnographers on their staff.

In the following pages we discuss a set of typical challenges to ethnographic work that tend to arise under such conditions, and propose effective strategies for dealing with them. For the sake of making a clear argument, we sometimes paint extreme positions, but we do recognize that there are always other ways of looking at things and other solutions to the issues we encounter so routinely. We will draw examples from the wide range of our ethnographic experience, that is to say from product design, technology development, innovation, and organizational change practice, but also realize that regardless of the specific domain we may be discussing, these problems pop up in all kinds of ethnographic work. In any case, our methods are grounded in ethnographic practice, our own and that of our colleagues.

1. "This takes too long!"

A commonly heard objection to ethnographic workpractice analysis is that "it takes too long." This is frequently raised by managers who have become accustomed to base 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.[2]

The truth is that sometimes ethnographic workpractice 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

[2] As a matter of fact, when anthropologists studied tribal communities, the tradition was a full year of fieldwork. There were good reasons for that, most importantly the necessity to study the agricultural cycle over a year, 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.

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 there was 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. And then we wrote a four-page memo to management with recommendations that ranged from improving employee voice in requirements analysis, 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 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 we had originally been hired for. The Center Manager even suggested that our report provided more insight than that of the architects and was certainly 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 something about current work practices that might impact the adoption of a new tool, then that could probably be done in the space of days or weeks. If, on the other hand, we are talking about a systematic assessment of how a business division works, or about creating an environment for learning and innovation, or 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. Something to remember as well is that 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. Often ethnography is part of a larger project, and then 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 issue that falls into the temporal category has to do with the staging of fieldwork. While ethnography has come to be more accepted, and sometimes even enthusiastically embraced in corporate circles in the last few years, the common myth is that fieldwork is something you do when all the technical problems have been solved. In a typical negotiation recently with an executive team from an external company we were discussing the question of how to design new technologies for them that would be loved, used, and hopefully 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 a rather authoritative manner that fieldwork was appropriate for the later phases of the project when prototypes would be tested on users, but "obviously" not for the first where, he said, the design team would be dealing with "purely technical" issues.

This "do ethnography later" mentality is widespread, persistent, and powerful. It is a very common managerial and engineering misconception. We responded that ethnography is appropriate any time that user and practice data could have an impact on design (cf. 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 and consistently 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 ("wrong jungle!").

2. "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 thinking about a brainstorming session to get the information he wants? Does she believe a focus group would get her the data she needs? It is useful to generate a discussion at that point to clarify why he wants data in the first place and what kinds of data would be most useful for that purpose.

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, he may well be working on the wrong design or be implementing something that isn't needed, or not fitted into, ongoing work practice and work culture or customer requirements.

When 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. By helping to abort an ill-conceived program early or by supporting effective implementation and acceptance by employees, workpractice analysis may be crucial to the success of a project or the avoidance of catastrophic failure.

Frequently questions about cost are related to concerns about effective project management. Managers have expressed opinions such as, "It's better to build something now: we can always tweak it later" or, "What should the developers be doing while you are gathering requirements?" Often these reflect a linear but common view of the product development process where newer managers assume that making architectural or design changes later in the prototype development process will be less costly than having ethnographers do fieldwork upfront. In our experience, many managers are genuinely concerned about how to parse roles and responsibilities in ways that are cost effective with regard to resources, time and labor. Yet some of them have felt strongly that it was unnecessary for engineers to spend time with customers, maintaining that software developers "need answers, not details [about customers' work practices]" from ethnographers. This unnecessarily segregates ethnographers from their technical and administrative corporate counterparts.

Finally, the issue around cost may sometimes be identified with the issue of scale. One Vice President expressed it thus, "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 90's focused on that effort (Aronson et al. 1996, Bishop et al. 1994, Jordan 1996a, 1996b).

In a telling example of such an approach, a Fortune 500 company asked our team of ethnographers and designers to develop a training program on ethnographic methods for a global sales and service

organization. Sales people underwent workshops that exposed them to participant observation and ideas around how workpractices can be understood from the workers' point of view and made more visible. (For a deeper discussion on making work visible and the distinctions between visible/invisible work, see Suchman, 1995; Star, 1999; and Button and Dourish 1996.) After the training period, sales people 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 to be a success. It was integrated into the sales cycle at a national level and later even adapted for a large consultancy engagement. Equally importantly, the organization developed a range of ways in which the actual "cost" of fieldwork – the time that people spent at customer sites – was either charged to the customer or absorbed by the company. So what we would like to emphasize is that managers' concerns about costs can be addressed productively at both a tactical and strategic level.

To sum up, when the issue of cost comes up, it is best to go back to a common understanding with the our corporate counterparts about what ethnographic workpractice 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 where 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.

3. "Don't bother. We can do this faster and cheaper with market research and focus groups!"

In our experience it has sometimes been the case that managers perceive ethnographic research and market research as 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.

There is an excellent paper by David Gilmore in which he points out that one thing to avoid is a territorial war with market research. The best strategy is to point out 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 early conceptualization in the heads and labs of researchers and designers, through the stages where it "gets a life" in peoples' workscapes and lifescapes. 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 most commonly used method for gathering data in companies, not only about work practices but also about opinions and attitudes, about 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, commonly accepted methodology for collecting data about users or other groups of people of interest to the company. For another, you can outsource it. It gets done as a limited, usually one-shot effort, with a predictable schedule, and comparatively cheaply.

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

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 a fact well-known 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 *opinions* about the future and their future *intentions*. Decades ago, W.F. White's 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.

The business press is full of stories about the limitations of focus group research. For example, Tom Van Horn, CEO of Mercata Inc, an early dot.com company that offered customers deep discounts through group buying, talked in an article in Fast Company about the fact that the business had relied on focus groups to direct them toward the kinds of merchandise that people would buy. That research indicated that people would be drawn in by high-cost products. By the time they found out that that was misleading information, the company had burned through 89 million dollars of start-up money and dissolved. As analysts told them, in spite of expensive research and advertising, they were never able to "connect with the customer," that is, find out what customers *really* wanted (Wilson 2004). This article is mirrored by an increasing number of others in the popular business press (e.g. Kiley 2005), indicating a rising consciousness among managers that focus group data are not as unproblematic as previously thought.

Qu 1 **Say**	Qu 2 **Do**
Qu 3 **Think**	Qu 4 **Feel**

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 employ focus groups as part of their repertoire of methods. However, focus groups may not give our managers the data they are looking for. They do certain things well, but not others. Let's look at a two-by-two matrix that considers what people say, do, think, and feel.

Figure 1 about here

Data in quadrants two, three, and four are subtle, tacit, implicit and context dependent, and are not easily surfaced 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 <u>say</u> they want. So if that is what managers are looking for, this is the right methodology. Often, however, that is not what they want. They are really interested in what people actually <u>*do*</u> in the workplace, or really <u>*think*</u> about a product, or, maybe most importantly, <u>*feel*</u> 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. They 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, but it is exceedingly common and 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 field work. 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 (though that would require an extended field engagement).

4. "Couldn't I just go myself and watch for a while?"

Mostly, managers are convinced they know what goes on in the workplace - otherwise they wouldn't be managers. So they think they should easily be able to check on a few extra little items. The problem here is that ethnography looks and sounds straightforward, but as any card-carrying anthropologist knows, it isn't. Some years ago, Diana Forsythe wrote a beautiful article on that topic with the title "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 since 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 is invisible to the person casually observing an ethnographer at work.

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 to patients, so ethnographic fieldwork is not *just* talking to people in the workplace. In actuality, ethnography often runs counter to what is common knowledge since it requires tapping into invisible structures and tacit knowledge.

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 over 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 vs. 95% preparation, analysis, synthesis and communication, most of which is invisible to outsiders.[3]

Part of the challenge therefore is how to successfully communicate the quality of work practice analysis as it unfolds over time, especially in situations where, by way of comparison, managers are more accustomed to assessing the amount and quality of code that's 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 evaluate 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?

Whether or not these instances reflect the disjunction between positivist and relativist paradigms (Forsythe 2001, Salomon 1991), 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 managers who are customer-focused, rather than purely technology focused. Managers responsible for such projects often consider the viewpoints and knowledge of customers as

[3] This invisibility is not restricted to corporations. As Giri points out, audits that are conducted in higher educational institutions often privilege the process - that is, the inherent "bias towards quick visible productivity" - over content (2000:178-179). Strathern explains further, "In teaching there must be a lapse of time - the [learning] process is not one of consumption but one of absorption and reformulation. In research, time must be set aside for all the ... activities that precede the genuine findings. Both require otherwise non-productive periods. Yet there is almost no language in the audit culture in which to talk about productive non-productivity." (Strathern 1997:318).

strategically important and may use our results to cross-pollinate ideas and decisions across business groups.

Secondly, 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 email 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 dollar 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!"

Thirdly, we try to involve managers in field visits before and after the technology is installed or the workpractices have changed at a customer site. If they don't see the realities on the ground, they often come to believe that they are faced with "naïve 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 we and they need to "change the maze, not the rat."

After participating in site visits and speaking directly with customers, 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 manages' opinions about ethnographic contributions to the project. They can help track what the early assumptions were of the customer and the technology *before* field data was 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

5. "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 can't be generalized.

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

What, then, are some strategies for making our findings applicable beyond the one site where we are doing fieldwork? Basically, we provide evidence that our 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, of hypotheses, of 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 their 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 workpractices in an outsourcing fieldsite that had just one employee who, of course, had to handle a gamut of production requests. When we documented the workarounds this employee had invented and employed on a routine basis, and checked them in several other field sites that were 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 multi-employee 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

[4] Of course, we would want to introduce a number of safeguards that would reassure us that we haven't run into the one atypical case in the whole company.

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 video tapes of work practices and the opportunity to co-analyze those with employee participants (Cefkin and Jordan 1994; Ruhleder and Jordan 1998), we always ask about typicality and combine this 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 which, when synthesized, demonstrate critical points of convergence and differentiation. Given that we are really 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 wide-spread 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 co-design session, some patterns need remarkably few probes before it is clear that, no, nobody else engages in this; or, yes, this is common practice.

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 work of the ethnographer 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. [6]

---

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

[6] What managers often don't understand is that one of the special skills trained ethnographers bring to the field is the ability to simultaneously look for known patterns (sometimes based on theory,

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 weren't exactly the same but exhibited the same phenomenon. For example, in an airlines operations room the workers had a curious habit of talking out loud, saying things to nobody in particular. Given that talking to yourself is highly censored in normal interaction, we wondered about the significance of this. It turned out that colleagues doing fieldwork in an operations room of the London metro had documented the same phenomenon, and so had workers in other kinds of control rooms where people were co-present 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.

6. "You can't quantify this!"

This is a challenge that is close to the preceding one about generalizing. Again, it may be useful to try to come to an understanding of the scope of the challenge. 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 (like interruptions, stalls, encouragement, collaboration) - all of these are easily retrieved from our records, especially from tape- and video recordings.

But frequently that is not the underlying issue for managers. What they really are asking for is a quantification of the effect our study-cum-intervention might have on performance and Return On Investment (ROI). This a much more difficult challenge because it requires admitting that, realistically, there is no scientific methodology that could specify that effect with precision. The reason for that is that organizations are not laboratories where you 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

---

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 where "anything goes" and nobody has a good handle on what moves the system.

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, etc.

Nevertheless, there are usually some measures we can cite in order to gauge the positive (or negative, for that matter) 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 dealing with providing 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 due, at least in part, to our intervention. Similarly, our support in the introduction of a new technology might lead to more rapid and thorough adoption than would otherwise 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 in order to be able to argue later on for ROI as well as for performance and productivity effects.

7. "This isn't scientific!"

Unfortunately, when a manager voices this kind of objection, it is usually an indication that he intends to take the high road and categorically dismiss anything we might propose. This is a tough one. Countering this objection means we'll need to persuade him 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 while a workpractice study is different from what he 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 to him than what he 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 most of the fields managers are likely to be familiar with. Managers tend to aspire to the model dominant in the natural sciences that tends to equate scientific research with randomized controlled trials, double-blinding and quantitative analysis. We might call this approach *the analytic paradigm* (Solomon 1999).

The analytic paradigm is the approach of choice in situations where we have well-developed theories regarding the phenomena in which we are interested. From these theories we deduce specific hypotheses which 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, where 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 work places that he is in charge of. He may see quickly that most situations of interest in workplace research are very messy, dynamic situations where little can be held constant. Here trustworthy theories and well-defined variables are hard to pin down. He 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 point 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 13th 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 like 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 essentially looks at how different parts of a system work together. The

researcher here is not an unobtrusive outside observer, but is an active participant, immersed in the situation. The central question is not: What is the relationship between these two variables, but how does this system work. There is no fixed research protocol, but rather we look for patterns of behavior, spatial arrangements, patterns of tool use, of collaboration, of communication, of organizational structuring. We check out where these patterns hold. We ask ourselves questions like: how typical is this; who else does it this way; what alternatives are there to solve this problem -- and that leads you 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 the real world. Sometimes we have strong intuitions based on past experience and we check those out - but always we look for patterns, and for exceptions to those patterns, and then for the reasons for those exceptions which 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 does become 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.

8. "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 be an indicator that expectations of managers and ethnographers are not be completely aligned or, at least, are not explicit enough. It takes many forms, depending on the scope and nature of the interaction between ethnographer/researchers on the 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 in sections two and six. When the question is not actually about the financial bottom line, they may be asking if we can help them increase the productivity of their organization, or they may be pressing for customer requirements (in the case of product designers), or software engineers may ask how they can use ethnographic information for writing better software. While the question about the bottom line is almost always lurking in the background, it is important to keep in mind that some people in every company are actually concerned about getting results that will really benefit their customers.

No matter what lies behind the question, we need to ask ourselves (and them) for some discussion of what results *they* 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've been in situations where we have been asked to help them achieve "a 20% increase in productivity by year's end", though whether that meant reduction in working hours, increase in product output, improvements in quality, or something else, was by no means clear. Nor was it clear how those increases would be measured. In virtually all cases, companies already have a variety of metrics in place and while those may not be optimal, it is important to know about them and, if appropriate, use them. If this is possible, they may already be doing a piece of our work for us. So the first step is always to find out what they mean by productivity or customer satisfaction or whatever the main issue is that they are concerned about; the second is to ascertain how they measure that now, or come to a shared understanding of how it could be measured as a result of our work; the third is to engage our counterparts in a discussion of the limitations of before and after studies in order to avoid problems later on when the expected improvements don't materialize because the market has changed, the customer base has shrunk, a competing product has come on the scene, or some other development has generated a rise (or decline) in outcome measures irrespective of our intervention.

What is particularly difficult is when there is no clear answer to the results question 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? After extensive interviews, we discovered that engineers were loathe 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, they thought that engineers had viewed claims in a technical rather than legal manner. Furthermore, when we examined the daily tasks of hardware and software engineers in the laboratory, we discovered that communication between hardware and software engineers was meager. As a consequence similar 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. Based on our recommendations, attorneys were assigned to the laboratory much earlier in the investigative process.

Moreover, 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.

In cases where the "results" of an ethnographic analysis are hard to predict, it is probably best to openly acknowledge that while at the same time citing successes with similar projects.

At the same time, 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' expectations run to 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.

While managers probably expect a report, what they really need from us is help in generating a change in the workpractices 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 upon and don't (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 are 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 video tapes 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 also empowered the workers in that they gave them a shared view on what they were doing and how they were dealing 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 naïve

(but interested) learners, so that these interactions are less feedback sessions where we report authoritative findings than collaborative working sessions where 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 (2x2x2x2) 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 her choice) meeting with *two* people from the project team. They commit *two* hours to a discussion of a particular, implicative finding from the report. The discussion centers around what possible change could be made that would lead to better outcomes and what steps the hosts are going to take in the next *two* weeks to move those towards reality. The group meets again after two weeks, assesses progress, deals with issues that have come up, identifies success or failure and what they have learned from that, and goes on to repeat the process with another intervention suggested in the report. The big advantage of this method is that it generates often lively discussion during the intervening two weeks in both pairs, the executives as well as the writers of the report who almost always use this time to deepen their understanding of the issues.

Conclusion: Changing the System

To get back to the title of this paper: how can we move from the frustrations of forever having to counter our managers' reservations, towards 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, rather than forever responding to objections.

There are various strategies to make inroads on entrenched positions. For example, we've 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 subvert an existing one) by introducing articles that include successful applications of ethnographic methods. Looking at articles that *misapply* quantitative or 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 pulls 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 where ethnography does not deliver the right data, when 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 substantially contribute to a change of attitude in the company.

So 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, Co-Experiencing; of educating managers, technologists and administrators and getting educated ourselves. It's 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. Still, 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, 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 didn't - 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 towards 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!

References

Aronson, Meredith, Libby Bishop, Melissa Cefkin, Brigitte Jordan, Nancy Lawrence, Lindy Sullivan, Connie Preston and Julia 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. (June)

Bernard, H. Russell, P. Killworth, David Kronenfeld and Lee Sailer
1984 The Problem of Informant Accuracy: The Validity of Retrospective Data. Annual Review of Anthropology 13:495-517.

Bishop, Libby, Melissa Cefkin, William Clancey, Brigitte Jordan, Julia Oesterle and Beth 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, Graham and Paul Dourish.
1996 Technomethodology: Paradoxes and Possibilities." Proceedings of CHI 1996:19-26.

Cefkin, Melissa and Brigitte 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.

Forsythe, Diana
2001 Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence. Stanford: Stanford University Press.

Gilmore, David
2002 Understanding and Overcoming Resistance to Ethnographic Design Research. Interactions, May and June 2002:29-35.

Heath, Christian and Paul 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, The Netherlands.

Jordan, Brigitte
1996a Ethnographic Workplace Studies and Computer Supported Cooperative Work. Pp. 17-42 in: The Design of Computer-Supported Cooperative Work and Groupware Systems. Dan Shapiro, Michael Tauber and Roland Traunmüller, eds. Amsterdam, The Netherlands: North Holland/Elsevier Science.
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 Sao Paulo, Brazil. Itu, Brazil, August 7-9.
www.lifescapes.org\Papers\0212_from_training_to_learning.htm

Kilesy, David
2005 Shoot The Focus Group. Business Week November 14.

Rogers, Yvonne and Victoria Bellotti
1997 Grounding Blue-Sky Research: How Can Ethnography Help? Interactions 4(3): 58-63. New York: ACM Press.

Ruhleder, Karen and Brigitte Jordan
2001 Co-Constructing Non-Mutual Realities: Delay-Generated Trouble in Distributed Interaction. Journal of Computer Supported Cooperative Work 10:1:113-138. www.lifescapes.org\Papers\delay_for_jcscw_2001.htm

Salomon, Gavriel
1991 Transcending the Qualitative-Quantitative Debate: The Analytic and Systemic Approaches to Educational Research. Educational Researcher 20:6:10-18.

Schuler, Douglas, and Aki Namioka, eds.
1993 Participatory design: principles and practices. Hillsdale, NJ: Erlbaum Publishers.

Star, Susan Leigh and Anselm 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.

Strathern, Marilyn, ed.
1997 "Improving Ratings": Audit in the British University System. In European Review 5(3):305-321.

Suchman, Lucy
1995 Making Work Visible. Communications of the ACM 38:9:56-64.

Whalen, Marilyn, Jack Whalen, Robert Moore, Geoff Raymond, Margaret Szymanski, and Erik Vinkhuyzen
2004 Studying Workscapes. In: Discourse and Technology: Multimodal Discourse Analysis. Philip LeVine and Ron Scollon, eds. Pp. 208-229. Washington D.C.: Georgetown University Press.

Wilson, Amy
2000 Put your Customers in Charge. Fast Company, November 2000. http://www.fastcompany.com/magazine/40/mistake.html

Acknowledgments

This article is based on a paper we delivered at the 2003 meetings of the Society for Applied Anthropology. We thank the many colleagues that 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 co-authors to acknowledge the magnitude of our debt. For specific comments on the present manuscript we particularly thank (list) and acknowledge that we have shamelessly incorporated their suggestions. As always, this article should be seen as a collaborative piece of work.

Figure 1
Say/Do/Think/Feel Two-by-Two[7]

Qu 1 **Say**	Qu 2 **Do**
Qu 3 **Think**	Qu4 **Feel**

[7] 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.