ORIGINAL
Applied Social Psychology Annual
forthcoming

Research on the Social Impacts of Robotics: Issues and Some Evidence

Paul S. Goodman Linda Argote

Carnegie-Mellon University

January 1984

Support for this research was provided by the Graduate School of Industrial Administration, the Robotics Institute, and the Program on the Social Impacts of Robotic and Information Technology at Carnegie-Mellon University.

The goal of our research program described in this chapter is to contribute to knowledge about how the introduction of robots affects individuals and the structure, functioning, and effectiveness of organizations. Robots are being used in increasing numbers in offices and factories throughout the world. Current estimates put the number of robots in use in the United States in 1983 at about 7,000 (Hunt & Hunt, 1983). Little is known, however, about how individual workers react to the introduction of robots or about how the use of robots affects organizational structures, processes, and outcomes. The research described here is an attempt to fill this gap in our knowledge about the social impact of robotics.

Other researchers have examined some aspects of the social impact of robotics. Previous work has typically examined more macroeconomic questions, such as how the use of robots is likely to affect employment levels (Ayres & Miller, 1983; Hunt & Hunt, 1983), or questions with a human factors orientation, such as what is the optimal division of tasks between robots and humans (Parsons & Kearsley, 1982). Our focus is more micro and more psychological: we examine how individual workers react to the introduction of robots and how organizations change when robots are introduced.

This chapter begins with a discussion of how we became interested in social issues surrounding robotics. We then discuss how we carved specific research questions out of the general problem area and how we began a research program to address these questions. Next, we summarize the methods and results of a recently completed field study about how workers react to a robot. The paper ends with a discussion of general issues, both conceptual and methodological, that arise in studying the impact of new technologies on individuals and organizations.

What Is a Robot?

One of the first issues that we confronted in our research was understanding and defining what a robot was. We spent a lot of time learning about the technical capabilities of robots and dimensions along which they varied. Many people at Carnegie-Mellon from Computer Science, Engineering, and the Robotics Institute helped us learn about robots and other forms of advanced manufacturing technology. We learned that there were two characteristics, multiple task capability and programmability, that differentiated robots (albeit not perfectly) from other forms of automation. The functions of most robots in U.S. factories today are to transfer material and to do certain processes such as welding, drilling, and spray painting. These are called level I or first-generation robots.

Researchers are currently working on level II or second-generation robots. The definition of a second-generation robot is a device that links intelligence to the capacity to do work. By intelligence, robotics researchers mean anything that can sense its environment, take information and understand it, and then plan and execute tasks to achieve a given goal. Examples of level II robots currently being developed include a roof-bolter robot that can go into an underground coal mine and secure its roof and an expert system that can make decisions about production scheduling on the factory floor. Thus, the abilities to sense and think are more developed in level II than in level I robots.

Another issue we confronted in thinking about our research project was: "Are robots different from other forms of technology?" In thinking about this question, we realized that there were at least two factors that might lead workers to view robots as qualitatively different from other forms of automation. For one thing, workers have been exposed to robots with glorified capabilities in films and television shows. In addition, in many current applications, a robot directly takes the place of a human worker. We think these two factors combine to make the introduction of a robot a more salient and probably more threatening event for workers than the introduction of another form of advanced technology.

Becoming Interested in the Problem

Our involvement in this research topic was stimulated by several factors. Part of the appeal of robotics as a research topic stemmed from the fact that it sounded like a new and exciting area that would build on previous research we each had done. Paul had done work on assessing the effectiveness of organizational changes (Goodman, 1979) and on theories of the institutionalization of change (Goodman, 1982). Linda had done research on the appropriateness of various organizational structures (Argote, 1982). Our expertise and interest in these areas seemed a good base for assessing the effects of introducing robots on individuals and organizations. The idea also appealed to us because we thought that the social issues surrounding robotics were important ones and that we could probably make a contribution to understanding them more fully.

Carnegie-Mellon University encourages and rewards an interdisciplinary approach to problems. In 1981 a large Robotics Institute was established to develop robotic technologies. While the primary work at the Institute is technological in nature, there is clear recognition of the need to look at the social consequences of robotics. We believe this "culture" at our University facilitated our entry and our continued work on robotics.

Framing the Problem

The Literature

While our research project appeared to be one of the first that examined the effects of introducing a robot on individuals and organizations, we thought that we could anticipate what some of these effects might be by reviewing literature on other forms of technological change. Accordingly, we reviewed literature on how individuals react to the introduction of electronic office equipment (Zuboff, 1982; Bikson & Gutek, 1983), to the introduction of computers (Mumford & Banks, 1967); and to the introduction of earlier forms of factory automation including numerical control machines (Mann & Hoffman, 1960; Whyte, 1961; Williams & Williams, 1964).

Several themes emerged from that literature that helped us in our study of robotics. One theme was the importance of the compatibility between an organization's technology, its structure, and its members (Emery & Trist, 1973). There are numerous examples in the literature of changes that had unintended negative consequences for the organization including, for example, decreased worker satisfaction, increased absenteeism, higher accident rates, and decreased productivity (Trist & Bamforth, 1951; Mirvis & Berg, 1977). These negative consequences have largely been attributed to the failure of the organization to take into account characteristics of its social system in designing and implementing technological change.

Another theme from the technological change literature was that changes in technology often affect characteristics of the jobs of individual workers (Billings, Klimoski, & Breaugh, 1977). In general, jobs with certain characteristics such as variety, significance, control, and feedback have been found to be more satisfying and motivating than jobs without these characteristics (Blauner, 1964; Hackman & Lawler, 1971; Hackman & Oldham, 1975). If a technological change such as increased automation affects these job characteristics, one would expect changes in the satisfaction and motivation of employees. In actual studies of technological change in the factory, automation has been found to increase the extent to which jobs were mentally demanding (Whyte, 1961), and to lead to workers' feeling more responsibility (Mann & Hoffman, 1960), less control (Blauner, 1964), and a greater sense of pressure (Mann & Hoffman, 1960; Whyte, 1961).

A third theme was that technological change often affects interaction patterns on the job. For example, Whyte (1961) reported that increased automation decreased the opportunities workers had to interact with their coworkers. Williams & Williams (1964) found that introducing numerical control machines required more coordination activities between support and production personnel.

Another theme from the change literature was that the manner in which a change is introduced affects workers' acceptance of and commitment to the change (Griener, 1967; Beer, 1980). For example, in their classic study, Coch & French (1948) found that, under certain conditions, participation in the design of change was associated with higher productivity, lower turnover, and fewer acts of aggression against the company.

The Research Questions

After much discussion between the two of us, we carved four research questions from the larger problem area of the social impact of robotics. We selected these questions because they emerged as important themes from our literature review and because the questions are of practical interest to managers, government officials, and union leaders. These questions were:

(1) How do individual workers react to the introduction of robots?

(2) How do the structures of organizations change when robots are introduced?

(3) What effects do the use of robots and workers' reactions to them have on organizational outcomes including absenteeism, accidents, turnover, and productivity?

(4) Which strategies for introducing technological change are most effective?

Conducting the Field Study

The Design

The design was essentially a before and after design, with a quasi-control group. We interviewed production workers in a factory department, where a robot was being introduced, three months before and three months after the robot was put on line. Since we believed that the introduction of the robot would have effects reaching beyond the immediate department where it was introduced, we also interviewed individuals from other departments at the plant. Individuals were selected if we thought their jobs were likely to be affected by the robot (e.g., engineers) or if we thought they could provide us with another perspective on how well the robot was working out (e.g., representatives from quality control). Thus, we interviewed first- and second-line supervisors and managers and representatives from engineering, maintenance, quality control, production scheduling, and personnel relations. We also interviewed production workers in an adjacent department (our control group).

We developed different interview schedules to interview this diverse group of individuals. In particular, we developed interview schedules for employees in the department where the robot was introduced, employees in an adjacent department, support staff, and supervisory staff, and we developed a supplemental interview schedule for the robot operators. We talked to a total of about 60 people during each of our main visits to the plant. We also observed the workplace during the introduction of the robot and administered a satisfaction questionnaire to production workers.

Most of our attention was focused on the department where the robot was introduced. At Time 1, 37 employees from this department were interviewed; 25 were interviewed at Time 2. Average interview time was 35 minutes for most employees; the interview with the robot operators took 45-50 minutes. The interview included topics such as workers' views on robots, on consequences of robots, and on the introduction of robots as well as descriptions of workers' job activitites and interactions at work.

Site Management

Our contact with the site began with a conversation with a faculty member at Carnegie-Mellon who was working with engineers at the site to develop robots with particular technical capabilities. The faculty member invited us to a meeting where we could meet representatives from the site and sell them the idea of a study on the social impact of introducing robots. Our proposal was accepted by the site and we were invited to a meeting at the plant to talk about the study. During the meeting, we toured the plant and talked about the purpose and methodology of the study with several managers. We promised to treat respondents' answers confidentially and to give the company a feedback report about their organization.

Collecting Data

Our first visit to the plant was designed to collect background information about the plant, its structure, technology, and culture and about how the work force had been prepared for the robot.

Our next visit to the plant, our first major data collection (Time 1), was both fascinating and exhausting. In order to interview people across three shifts, for example, we began collecting data from the third shift around midnight and worked until around 4:00 a.m. We returned the next day around noon to collect data from the first shift and remained at the plant until 8:00 p.m. to interview the second shift. Everyone we came into contact with at the plant during our first visit was very cooperative. This had changed slightly by our second visit, Time 2. One of the first-line supervisors was very concerned during our second visit that the time employees spent talking to us would detract from the department's ability to meet production goals. In retrospect, we might have been able to prevent this by anticipating the supervisor's concerns and by getting management to explicitly introduce some slack into the department's production goals. Other than the less-than-enthusiastic attitude of one of the supervisors, the second phase of our data collection went smoothly.

Characteristics of the Site

The primary technological processes at the plant we studied involved forging and machining metal products. The work force at the plant, which numbered about 1,000, was nonunion, predominantly blue collar, and fairly stable. Relationships between labor and management appeared to be good.

The robot was introduced in a department in which the basic operations were milling and the grinding bar stock. There were approximately ten different operations in the department. Forty people worked across the department's three shifts. The robot, placed at the beginning of the work flow, loaded and unloaded two milling machines. One person operated the robot on each shift. No one lost his or her job as a result of the robot's use.

What We Learned from the Study

Our findings are discussed here in terms of how the robot affected workers' beliefs, activities, and interactions and the overall organization. The results are for the department where the robot was introduced unless otherwise stated. Our findings are described in greater detail in Argote, Goodman, and Schkade (1983).

Beliefs about Robots

We were interested in our respondents' beliefs about what a robot was. So we asked them how they would describe a robot to a friend. The concepts mentioned most frequently were: mechanical man, preprogrammed machine, something that loads machines, increases productivity, or reduces manual work. These concepts fall into three categories: general descriptions (mechanical man), functions (loads machines), and consequences (reduces manual work). The frequency of responses in the different categories did not change significantly between Time 1 and Time 2. However, we found a significant increase in the number of concepts mentioned by each individual over time. This suggests that our respondents' conceptions of a robot became more complex as they gained experience with the robot.

Another issue that interested us was how our respondents learned about robots. The movie Star Wars and television shows depicting humanlike robots were mentioned most frequently as the source from which our respondents learned about robots. These humanlike robots in the media probably contributed to the tendency we observed at the plant for workers to anthropomorphize the robot. Workers on each shift named the robot and endowed it with human qualities.

Beliefs about the Effects of Robots

We were also interested in our respondents' perceptions at Time 1 about how the robot would affect them or their department and in our respondents' perceptions at Time 2 about how the robot actually had affected them or their department. Hence, we presented our respondents with an outcome (e.g., the chances of an accident) and asked them at Time 1 whether the robot would increase, decrease, or have no effect on that outcome. Similarly, we asked them at Time 2 whether the robot actually had increased, decreased, or had no effect on the outcome. Our findings were that a majority of the workers at Time 2 felt that robots increased productivity but did not have much effect on the quality of output, the amount of downtime, or the number of people who work in the department.

We were also interested in testing whether the number of responses in the increase, decrease, or no effect categories changed from Time 1 to Time 2. Our dependent variable, the numbers of responses in the different categories, was categorical. Hence, we began to learn about techniques appropriate for categorical dependent variables, such as probit analysis (McKelvey & Zavoina, 1975). This was the first time we used these techniques in our research and we believe one of the first times the techniques were used in this type of organizational behavior study. We found that workers were significantly more likely at Time 2 than at Time 1 to say that the robot increased the chances of an accident, increased costs, and lowered the quality of the output. Thus, workers in the department where the robot was introduced became less optimistic over time about the effects of the robot. Workers in the adjacent department, our quasi-control group, also became less optimistic over time about the effects of the robot; however, the adjacent workers were more optimistic than workers in the department with the robot about its effects.

We did not have access to company records data to examine the effect of the robot on objective measures of productivity and quality. While perceptual measures or ratings of productivity often correlate positively with objective measures of performance (Georgopoulos, 1965), the correlations are not perfect. In our work in progress, we are examining the effects of introducing robots on objective measures of productivity. We will also examine the correspondence between our objective measures and respondents' perceptions of the effects of the robot.

Beliefs about Introducing Change

We were also interested in the effectiveness of the strategies the company had used to communicate about the robot to the work force. The company had used a fairly comprehensive set of strategies to introduce the robot into the plant. These strategies included an open house in which the operation of the robot was demonstrated, talks given by the plant manager, discussions with first-line supervisors, and notices posted around the plant. We asked our respondents at Time 1 whether they learned about the robot from a particular source at the plant and how much that source increased their understanding of the robot. We learned that the demonstration was seen by workers as the most effective communication method used at this site. More generally, we found discrepancies between what the workers wanted to learn about the robot and the information that management provided. This discrepancy might be due to the fact that this was the first robot installation in the plant. There were ambiguities about how the job of robot operator was to be rated and about who was to be the operator that management was somewhat reticent to talk about.

In addition, we asked our respondents questions about their participation in the robot introduction. Specifically, we asked them how much influence or involvement they actually had on decisions about: (a) whether the robot would be introduced; (b) where it would be placed; and (c) who would run it. We also asked them how much influence they thought they should have had on these decisions.

Our respondents reported that they had had no influence on these three decisions, but somewhat surprisingly, they did not think that they should be involved in decisions about where the robot was placed. Our respondents felt that it was management's job to make such a decision. Our respondents thought that they should have had a little influence on decisions about whether the robot would be introduced and who should run it, a finding that contrasts with a general view in the organizational field about the desirability of participation. Our experience at this plant suggests that it is important to do a fine-grained analysis of the types of decisions that employees want to participate in and that employees may prefer not to participate in all of the decisions associated with introducing a robot.

Effects on the Operator's Job

We wanted to capture in detail how operators' jobs changed with the introduction of the robot. We developed a special interview schedule to assess changes in the job activities and interaction patterns of the robot operators.

The robot provided material handling functions for two milling machines. The human operator was responsible for the two milling machines and the robot. Introducing the robot removed the materials handling activity from the operator's job and added a new activity, robot operator. When we asked the operators about the differences between their jobs before and after the robot introduction, they commented that operating the robot: (a) increased the number of job activities; (b) required more monitoring than doing activities; (c) required more skills (e.g., programming); and (d) implied more responsibility.

What are some of the consequences of these changes for the workers? From our qualitative interviews with the robot operators, we learned that they experienced more stress or pressure:

> It's nerve racking . . . There are lots of details . . . There is more stress now . . . We have more responsibility . . . They want the robot to run and we have to keep it going . . . That's hard because it's still relatively new.

Another more subtle source of stress arose from workers' comparing themselves to the robot. During our first visit to the plant, workers speculated about whether an operator who was particularly quick would be able to beat the robot. By our second visit, though, workers seemed resigned to the robot always being able to outproduce a human worker. The reason for greater productivity was simple: robots do not take breaks or even go to lunch!! Another potential source of stress was the incompatibility between activities required by the job and preferences of the worker. One of the operators in our sample mentioned that he found observing and monitoring activities more boring than manual activities and that he actually preferred doing manual activities.

It would be premature at this point to speculate whether this increased stress was good or bad for the individual or the organization. Studies have shown that increased stress is associated with increased turnover and absenteeism (Porter & Steers, 1973) and that stress can lead to both increments and decrements in performance (McGrath, 1976).

Our analysis revealed some interesting discrepancies between the operators' perceptions of their job and objective realities at Time 2. Workers reported that they had more activities to perform but an objective count showed fewer activities. Operators also reported that they had less control over their work at Time 2, yet at both Time 1 and Time 2 they had control over all their equipment--they could start and stop the machinery and override the robot. We are examining these discrepancies in our current work.

Effects on Interaction Patterns

Introducing a robot can change interaction patterns in the workplace, and we believe that these changes can have psychological and behavioral consequences. For example, if the new technology breaks up existing social interactions and isolates the worker, we expect increases in alienation and resistance to the new technology.

The robot operators reported at Time 2 that they had less opportunity to talk with people on the job than they had had before the robot was introduced. Introducing the robot did not change the basic work flow in the department, however. All the workers, including the robot operator, were located in the same area and participated in the same part of the work flow. The major changes in interactions occurred between support personnel from engineering and maintenance departments and the operators of the manufacturing cell. There was more frequent contact among engineering, maintenance, and the robot operators.

Effects on the Organization

Introducing the robot also affected the department because it required re-evaluating and reclassifying the operator's job. Introducing the robot had eliminated certain job activities and added other activities, and the question was whether the net change indicated that the job should be upgraded. Management did upgrade the job, but there was agreement among the workers that the grade and associated pay for the operator's job were still too low. There was no evidence that the robot had any effect on other department policies, procedures, or formal coordination mechanisms.

Research Dilemmas

In the previous section, we outlined the findings of our first field study of the impact of robots in the factory. These results are important because they represent, to our knowledge, the first systematic evaluation of the impact of robots on individuals and organizations. With this foray into the factory of the future, a number of dilemmas emerged that shaped our research plans. We call them research dilemmas because they are problems not easily resolved by current social, organizational, or psychological theory or measurement techniques. To some extent, they require new approaches and thus present new research opportunities. In some ways, we see these research dilemmas as unique to the study of new technologies. That is, given the relative infancy of this research area and the nature of current robotic installations, there are a set of special problems about theory, sampling, and instrumentation that need to be confronted. While any of these issues may be relevant for field research, it is the constellation of dilemmas presented below that are common to studies of new technologies.

Lack of Developed Theory

When we initiated our study on the impact of robots on the individual and the organization, our first thoughts turned to selecting a theory that would guide our research efforts. Our literature review, unfortunately, did not identify any well-developed theoretical perspective. This lack could be due to the general state of organizational psychology or to something inherent in our research problem. While the field of organizational psychology is still relatively new, there are some fairly well-defined theoretical paradigms in such areas as work attitudes and work motivation. Therefore, the problem or dilemma seems to be caused by some of the inherent characteristics of our research questions rather than being inherent in the field. In particular, the research questions are at different levels of analysis and require perspectives from different disciplines. For example, concepts in social perception may help us understand how workers view robots. Motivational theory may help us understand the type and level of employee response generated by the new technology. Research on stress may help us understand sources of stress associated with operating new technologies. Engineering perspectives may be important in helping us define or dimensionalize the new technology. Perspectives from economics or finance may be important in assessing the costs and benefits of robotics and other advanced manufacturing technologies.

In our first study, the approach we took to this theory problem was to draw on the systems view of organizations as articulated by organizational researchers such as Katz & Kahn (1978). This perspective forced us to consider the total organizational context when assessing the impact of robotics and accounts for our interest in support personnel around the plant and workers adjacent to the robot as well as those workers directly affected by the change. In addition, we adopted the interaction, activities, and sentiment paradigm used by Whyte (1961) and others in studies of technological change. While this perspective had some useful focusing and organizing value, it did not generate any precise hypotheses that could be used in our original or subsequent studies.

The dilemma, then, is that we have an interesting set of research questions but not a well-defined theory to use to approach these research questions. Indeed, there is unlikely to be a single theory for an inherently interdisciplinary research program such as determining the social impact of robotics. In the short run we are trying to develop a sequence of theories that deal with the following questions:

1) How does one objectively dimensionalize technology? What is the underlying theory or construct space of this new technology?

2) How do workers perceive this new technology? What can we draw from the literature of object and person perception to understand this process?

3) How do workers react to robots after some perceptions have been established? Some researchers have approached this question by using a person/job fit model where incongruency leads to stress and dysfunctional behaviors. We also need models that explain how and when new technologies will increase motivation and functional behaviors for the individual and organization.

4) How and under what conditons will the new technology lead to changes in organizational effectiveness indicators (e.g., productivity) and changes in organizational structures?

5) What effect does the manner in which the new technology is introduced have on workers' reactions to the technology?

While we think this research area is difficult because of the lack of developed theory, it also provides a unique opportunity. Rather than developing a general theory, we see the need to develop a series of linked theories focusing on the questions enumerated above. One of the exciting aspects of this work is that it requires connecting psychological theories with other theoretical perspectives.

<u>Selecting The Appropriate Time Frame</u>

The design of our first study was a before and after design with a quasi-control group. Major measurements occurred three months prior to the robot installation and three months after the installation. There was some limited on-site observation at the time of the change. Using this type of longitudinal design is necessary if one wants to assess the impact of robots on the worker and organization. However, selecting the appropriate time frame for the design is critical in determining what types of results are observed. If one is interested in the effect of robots on accidents and accidents are a relatively low frequency event, then selecting a short time frame may lead to significant underestimates on the accident variable. On the other hand, if one wants to assess the effect of robots on worker stress and adjustment to stress is a phenomenon that occurs over time, then the time frame selected for an analysis will affect any interpretation of one's results. For example, we reported increases in stress as a function of the robot installation. What we do not know is whether workers would adjust or accommodate to the new levels of stress and report lower levels of stress over time. We are looking at this in research currently underway.

The problem of selecting the appropriate time frame is affected by the lack of control that characterizes organizational field experiments. In a laboratory experiment, the researcher has control of when to take the before measures, manipulate the variables, and take after measures. In organizational experiments that level of control is not available. For example, in our study we did not have a clear before measure. There was knowledge and discussion within the plant that a robot was going to be introduced before we were able to collect the before measure. The same problem holds for the after measure, which is probably more crucial. We negotiated with the firm for rights to return three months after the installation to collect our first after measure and then to return later to conduct a second follow up. We were able to do the first after measure, but the firm decided not to permit the additional follow-up at the time we planned. The expressed reason was there had been a lot of changes in the firm (e.g. layoffs and restructuring) and there was some concern that our intervention through measurement might generate some negative reactions from the workers.

The issue inherent in selecting the appropriate time frame is not just that we, as researchers, have little control over the organization. We also suffer from a lack of theory about the process of organizational adjustment over time. Consider the following example from our study. A robot is introduced. Both the act of creating change and inherent aspects of the change (e.g., doing more cognitive than physical activities) can create stress. Over time, the impact of creating the change should decrease and the worker, if he is experiencing some level of stress, may engage in some processes to reduce this stress. The question is when will these effects occur. Unfortunately, there is little guidance in the change or stress literature to tell us when to take these after measurements. The dilemma, then, is that when we take the measurements affects the results. However, we have little control over when to take the measurements and little theoretical guidance. In other studies (Goodman, 1979) we have approached this problem by taking frequent measures over a long time period. This approach, while costly, enables one to capture more fully adjustment processes over time.

Small Sample Problem

While there has been a lot of publicity directed toward robotics, the installation of robots is still a low-frequency event. That means we are dealing with the small sample problem whether we are talking about a within-firm or a between-firm design.

In the within-firm case in the study we reported, the corporation that gave us access was actively involved in introducing robots. The particular plant where our study was conducted was a natural place for robotics installations. On one hand, the plant was technologically sophisticated and the idea of introducing new technology was more the norm than the exception. On the other hand, there were multiple work stations that lent themselves to robots (e.g., stations where there was a lot of lifting of heavy materials). But even in this environment, there was only one robot installation. At the onset of our study, we had a department of 40 employees but only 3 individuals were directly involved in operating the robot. Three is a small sample size when one wants to assess the impact of operating the robot on the worker. The sample size is too small to allow one to control for variables such as age, education, or previous experience with advanced technologies that might affect workers' reactions to the robot.

We were also interested in the impact of robots on the structure of the organization (e.g., decision making and coordination mechanisms) and on organizational effectiveness indicators. For example, we expected that the robot would place greater demands on certain staff jobs (e.g., maintenance) and hence hypothesized increases in role stress and perhaps the emergence of new coordination mechanisms between staff and production personnel. Although there were traces of emerging stress, it was hard to verify. The single installation of a single robot did not create sufficient demands to fully test the hypothesis. The small sample size does not give us the range of variation we need to fully test some of the research questions.

An obvious strategy to counteract the small sample problem is to move to a between-plant design. For example, in the case of the corporation we are working with, assume that it has 50 plants with at least two robot installations. This would generate a sample of approximately 300 operators (2 installations x 50 plants x 3 operators), which would lend itself better to some of our research questions. However, the sample is likely to be very heterogeneous in terms of type of job, pay level, skill training. The plants are likely to be very different in terms of union status, stability of employment, region, and product. This means that the apparent sample size is really much smaller when we control for these job, occupation, and plant differences. In addition, the costs involved in getting access to, traveling to, and collecting data from a large number of plants is enormous given the small amount of information collected per plant. The point is that the between-plant design does not completely solve the small sample problem.

We labeled this section research dilemmas because there is no easy solution to the small sample problem. Our current strategy is to draw a purposive serial sample of single plants. We are sampling critical plant and environmental characteristics that provide the background for robot installations. Union-nonunion contracts and stable-unstable economic (and employment) environments are two obvious sampling criteria. We are also looking at job and technology characteristics. Some robots replace parts of existing jobs (e.g., lifting material) while others replace all the job activities (e.g., welding). Some robots are doing primarily simple activities (e.g., lifting) while others are doing more complicated activities (sensory functions). These technological differences as well as the basic components of the job (e.g., skill level, training required) are reflected in our sampling plan.

The results of this serial sampling will be an integrated set of case studies from different organizations involving different technologies. The task of the researcher will be to identify uniformities in this stream of data. The problem of the small sample size will not go away. For some of the research questions we're interested in, we will not be able to use the data analytic techniques that are prevalent in the organizational psychology literature. The question is whether mainline journals will find this type of research acceptable.

<u>Lack of Instrumentation</u>

An obvious derivation from the discussion of theory and the small sample problems is that we lack precise and sensitive measures that demonstrate reasonable psychometric properties for assessing the impact of new technologies. This lack of instrumentation is not simply because we are examining a new research problem. New technologies have always been evolving and a subject of inquiry. Robots are simply a new generation of technologies. We think the lack of instrumentation can be traced to the low level of theoretical elaboration about: (a) the concept of technology <u>per se</u>, and (b) the technology - job/person interface. In addition, in the case of robotics, we are confronted with the small sample problem, which was discussed earlier in more detail. In our study there were only three or four individuals that were directly affected by the robot, a relatively small sample for developing standardized instruments.

In our initial study, we used standard work attitude measures to capture any effects of the new technology. We found few, if any, differences between the experimental and control departments and between before and after on these standard attitude measures. That is probably because both groups were fairly satisfied with their jobs and the items, worded at a general level (e.g., "How satisfied are you with your supervisor?"), were not sensitive to some of the actual changes that were occurring. We also used some semi-structured items in an interview context that were much more useful in providing information about workers' perceptions of and beliefs about the consequences of introducing robots.

The broad strategy questions in the measurement area concern whether we should move toward a more idiographic versus a nomothetic approach, what types of methods and instruments we should use, and how we can deal with the psychometric properties of our measures. Our current position on measurement issues is shaped by the current low level of theoretical development as well as the problem of small samples. Our next series of studies on robots will be in different types of organizations with different types of jobs and robots. The subject pool will be small. Our approach will be idiographic in nature; the instruments will be designed around specific jobs and specific organizational settings. Of course, we will look at common issues (e.g., motivation, stress, participation) across the studies. The questions we ask, however, will be tailored to the particular site. Interviewing and observing will be the primary measurement tools. The reliability estimates will be derived from examining coding of the interviews and observational protocols.

Attrition Problem

A longitudinal design is appropriate for tracking the impact of robots on individual and organizational outcomes. Inherent in that design, however, is the problem of sample loss or attrition. This problem is particularly acute in the impact of new technologies where sample size is likely to be small. In the traditional laboratory study, the researcher has some control over the subjects over time, particularly if they are students. In the organizational setting, there are a variety of forms of attrition, none of which are controlled by the researcher. Absenteeism and illness are examples of temporary forms of attrition. While these workers eventually return to the organization, it may not be when the researcher is present. For example, 12 of the 37 employees from the department where the robot was introduced that we interviewed at Time 1 were not available for interviewing at Time 2. There did not appear to be any differences in absentee rates between the two departments during the time of our interviews or that the absentee rates were related to our study. For us, returning to the site at another time to pick up these workers was much too costly given our research budget. Further, access to the organization for the researcher is rarely open-ended.

Another form of attrition is more permanent in nature. Workers may move to a different part of the organization or leave the organization. This reality also affected our study. By the time we planned to collect our second wave of follow-up data (Time 3), all the robot operators we had interviewed earlier had left the company or the department. (The company was less receptive to this third data collection, so it did not occur.)

The problem is that organizations are dynamic entities where constant change in the work force is the rule rather than the exception. Given the small sample size problem, attrition will severely limit certain types of analyses over time. In cases where all the respondents leave, there is little one can do. In cases where some of the sample is lost, one can use statistical techniques not generally used by psychologists to assess the effect of the attrition on the representativeness of the sample.

For example, in our study, 12 of the 37 employees we interviewed concerning their beliefs about the robot at Time 1 were not available for interviewing at Time 2. The question was whether our sample at Time 2 was representative of our population at Time 1.

Our first instinct for addressing this question was to do a standard χ^2 test based on the multinomial rule to see if the frequency of employees in different categories (e.g., first, second, or third shift) at Time 2 differed from the frequency of employees in the various categories at Time 1. However, we were concerned that our data did not meet the assumptions of the χ^2 test. In particular, our Time 2 sample was drawn without replacement from the population of employees we interviewed at Time 1; hence our Time 2 sample was not independent of our Time 1 population. Conversations with two economists at Carnegie-Mellon, Dennis Epple and Lars Hansen, confirmed our suspicions that it was not appropriate to use a standard χ^2 test. The standard χ^2 test was based on the multinomial rule, which assumed that sampling was done with replacement or that the sample space was very large so that the sampling plan did not make a difference.

We began to read about sampling distributions and learned that the hypergéometric distribution was the appropriate distribution for situations where one was sampling from a finite space without replacement (Hays, 1973). Therefore, we used the hypergeometric distribution to test whether our Time 2 sample was representative of our Time 1 population. We derived from our Time 2 sample maximum likelihood estimates of the frequencies of individuals in various categories (e.g., first, second, or third shift) in the population most likely to have generated our sample (see Thiel, 1971, for an overview of maximum likelihood techniques). We then computed a χ^2 test statistic based on the likelihood ratio (rather than on the multinomial rule that is the basis for the standard χ^2 test). The likelihood ratio compared the probability of drawing our Time 2 sample from our Time 1 population to the probability of drawing our Time 2 sample from the population most likely to have generated it.

The results of these analyses are presented in Table 1. The χ^2 values were not large enough to reject at more than moderate levels of significance ($p < .25$) the hypothesis that our Time 2 sample was a random sample drawn without replacement from the population of employees we interviewed at Time 1. Thus, our sample at Time 2 appears to be representative of our population at Time 1 on these variables.

Current and Future Research

The research dilemmas delineated above shape to a great extent the nature of our current and future work on robotics. We have begun to develop a sample of heterogeneous firms, with different robotic installations. Selecting a sample in organizational settings is, of course, quite different from drawing on a university subject pool. We have to convince the firms it is in their interest to participate in our research program. Managers in these firms are often more interested in technological information than in behavioral issues. A further complication is that we want to identify those firms that intend to introduce robots, so we can conduct a before and after study. Despite these problems, we are beginning to generate a sample of 5 to 10 firms that differ in terms of unionization, economic climate, type of jobs, and the nature of the robot installation. Thus, our strategy is to generate a sample of prototypical organizations that reflect the range of conditions under which robots are introduced. To the extent that we capture this range, we should be able to make meaningful statements about the social impact of robotics.

The unique properties of each organization will determine whether we should focus primarily on the impact of robotics on the individual, on the organization, or on various effectiveness indicators (e.g. productivity, accidents). For example, if the firm does not have an information system designed to measure productivity around a particular installation, it is unlikely we could consider that research problem. Developing information systems in organizations is often not permitted and clearly is a costly endeavor. Once the problems are identified, theory development would follow. The goal would be to develop carefully delineated models that deal with specific problems such as the process by which workers develop representations of robots or adjust to changes in their work. Of course, some of this theory development can occur before any site selection, but given the unique characteristics of organizations, it will be important to incorporate the contextual aspects of each organization into the model development.

Given the small sample problem and the unique characteristics of each organization, our strategy is to develop idiographic measurement procedures. We think measurement in this research area should reflect the unique aspects of the organization, the robot installation, and the individual worker. We are currently developing generic interview and observational instruments that can be tailored to particular settings.

We have begun collecting data at one of these firms. Its context is very different from our first study site: it is unionized, has several robot installations, is located in the north, and is in an area that is depressed economically. The introduction of robots in this plant will lead to more worker displacements than in our first site. From our first interview wave, we can tentatively identify similarities and differences between this site and our first site. In terms of similarities, discrepancies occurred between what management communicated about the change and what employees wanted to hear in both plants. In terms of differences, employees' views of robots in the new site are more work-based than media-based and in some cases more negative.

After we have collected data from this second and other organizations, our task will be to put the pieces together, find uniformities, and understand differences. In a sense, we will be developing a theory of contexts. That is, we will specify how our contextual variables condition our results. For example, ideally, we should be able to make statements about conditions under which workers are likely to react positively or negatively to robots, conditions under which robot operators experience stress, conditions under which organizational structures change, and conditions under which participation works. Finally, we envision that we will shift our attention to new generations of robots and new types of technology such as expert information systems.

Table 1

Comparison Between Employees Interviewed at Time 1 and Time 2 on Key Characteristics

	Number of Employees Time One	Number of Employees Time Two	Maximum Likelihood Estimates	χ^2	df	p
Shift 1	18	10	15	2.28	2	$p < .50$
2	12	10	15			
3	7	5	7			
Grade 7	3	1	1	3.83	3	$p < .50$
8	9	8	11			
9	19	14	20			
10	3	2	2			
Tenure at Plant (PT)						
PT $\leq$ 7 years	10	9	14	3.03	2	$p < .25$
7 years $<$ PT $\leq$ 8 years	16	9	13			
PT $>$ 8 years	11	7	10			
Tenure on job (JT)						
JT $\leq$ 1 year	8	6	9	0.81	3	$p < .90$
1 year $<$ JT $\leq$ 3 years	9	5	7			
3 years $<$ JT $\leq$ 5 years	8	6	9			
JT $>$ 5 years	12	8	12			

References

Argote, L. (1982). Input uncertainty and organizational coordination in hospital emergency units. Administrative Science Quarterly, 27, 420-434.

Argote, L., Goodman, P. S., & Schkade, D. (1983). The human side of robotics: How workers react to a robot. Sloan Management Review, 24, 31-41.

Ayres, R. V., & Miller, S. M. (1983). Robotics: Applications and social implications. Cambridge, MA: Ballinger.

Beer, M. (1980). Organizational change and development. Santa Monica; CA: Goodyear.

Bikson, T. K., & Gutek, B. A. (1983). Advanced office systems: An empirical look at utilization and satisfaction. Santa Monica, CA: Rand.

Billings, S., Klimoski, R. J., & Breaugh, J. A. (1977). The impact of a change in technology on job characteristics: A quasi-experiment. Administrative Science Quarterly, 22, 318-339.

Blauner, R. (1964). Alienation and freedom. Chicago: University of Chicago Press.

Coch, L., & French, J. R. P. (1948). Overcoming resistance to change. Human Relations, 1, 512-532.

Emery, F. E., & Trist, E. L. (1973). Socio-technical systems. In F. Baker (Ed.), General systems approaches to complex organizations (pp. 249-260). Homewood, IL: Richard D. Irwin.

Georgopoulos, B. S. (1965). Normative structure variables and organizational behavior. Human Relations, 18, 155-169.

Goodman, P. S. (1979). Assessing organizational change: The Rushton quality of work experiment. New York: Wiley-Interscience.

Goodman, P. S. (1982). Change in organizations. San Francisco: Jossey-Bass.

Griener, L. E. (1967). Patterns of organization change. Harvard Business Review, 45, 119-128.

Hackman, J. R., & Lawler, E. E. (1971). Employee reactions to job characteristics. Journal of Applied Psychology Monograph, 55, 259-286.

Hackman, J. R., & Oldham, G. R. (1975). Development of the job diagnostic survey. Journal of Applied Psychology, 60, 159-170.

Hays, W. L. (1973). Statistics for the social sciences. New York: Holt, Rinehart and Winston.

Hunt, H. A., & Hunt, T. L. (1983). Human resource implications of robotics. Kalamazoo, MI: Upjohn Institute.

Katz, D., & Kahn, R. L. (1978). The social psychology of organizations (2nd ed.). New York: Wiley.

Mann, F. C., & Hoffman, L. R. (1960). Automation and the worker. New York: Holt.

McGrath, J. E. (1976). Stress and behavior in organizations. In M. Dunnette (Ed.), Handbook of industrial and organizational psychology. Chicago: Rand McNally, 1976.

McKelvey, R. D., & Zavoina, W. (1975). A statistical model for the analysis of ordinal level dependent variables. Journal of Mathematical Sociology, 4, 103-120.

Mirvis, P. H., & Berg, D. N. (Eds.) (1977). Failures in organization development and change. New York: Wiley.

Mumford, E., & Banks, D. (1967). The computer and the clerk. London: Routledge and Kegan Paul.

Parsons, H. M., & Kearsley, G. P. (1982). Robotics and human factors: Current status and future prospects. Human Factors, 24, 535-552.

Porter, L. W., & Steers, R. M. (1973). Organizational, work, and personal factors in employee turnover and absenteeism. Psychological Bulletin, 80, 151-176.

Thiel, H. (1971). Principles of econometrics. New York: John Wiley and Sons.

Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the Longwall method of coal-getting. Human Relations, 4, 1-38.

Whyte, W. F. (1961). Men at work. Homewood, IL: Richard D. Irwin.

Williams, L. K., & Williams, C. B. (1964). The impact of numerically controlled equipment on factory organization. California Management Review, 7, 25-34.

Zuboff, S. (1982). The new worlds of computer-mediated work. Harvard Business Review, 5, 142-152.