NBER/Sloan Project Report

Industrial Technology and Productivity: Incorporating Learning from Plant Visits and Interviews into Economic Research

Papers presented at the Annual Meetings of the American Economic Association, January 2000

This volume has not been reviewed by the NBER Board of Directors. The opinions expressed in these papers are solely those of the authors. The volume includes an inventory of NBER Working Papers produced as part of the NBER Sloan Project on Industrial Productivity. More information on the project is available at **www.nber.org/sloan/**

CONTENTS

PREFACE: THE NBER-SLOAN PROJECT ON PRODUCTIVITY CHANGE

Martin Feldstein*

The papers presented at this session of the American Economic Association's annual meeting are based on research done at the NBER on the sources of productivity change in industrial companies. The project, supported by the Sloan Foundation, began in 1994. Professor Adam Jaffe and I have been the directors of the project with Adam taking primary responsibility for the project's operational management.

This investigation of the sources of productivity change is particularly timely because of the surprisingly rapid increases in productivity that have occurred in the past several years in the United States, particularly in manufacturing industries. Although an exogenous technological shift in communication and computing capability may be an important source of this productivity gain, the greater extent to which productivity has increased in the United States than abroad suggests that there are other factors about the industrial environment that are influencing the rate of increase of productivity.

An unusual feature of the project has been our emphasis on visiting companies and factories in order to see the production process and to talk with managements about the sources of productivity change. Since the beginning of the project, the NBER has supported some 50 individual studies done by NBER researchers in a wide range of different industries. Each of these projects has involved site based research in which the researcher observed the company in action, developed specific research ideas, and discussed those ideas with the business managers. A primary focus of the project has been to encourage such direct observation and discussion as an integral part of economic research.

While the specific research projects each involved between one and three researchers, some of the general visits to factories involved groups of ten to fifteen researchers. We have referred to these as our "Pin Factory" visits in deference to Adam Smith and his famous description in *The Wealth of Nations* of the production process in a pin factory. Smith was able to see how the division of labor increased productivity and went on to theorize about the gains from increased market size, expanded trade, improved infrastructure, and other things on the basis of this visit. We hoped that observing would be fruitful for us as well. Professor Susan Helper was particularly helpful in organizing these Pin Factory visits and in preparing us for what we would be seeing.

---

*President of the National Bureau of Economic Research, and Professor of Economics, Harvard University. These remarks were prepared for the session on the NBER-Sloan Productivity Project at the January 2000 meeting of the American Economic Association.

We found that the managements at the companies we visited were generally eager to show their facilities, to tell us about their management practices, and to explain why they did certain things and how their practices have changed over time. They were also generally open about answering our questions. We would discuss not just the production process itself but also things like the compensation systems used to motivate workers, the criteria on which incentives were based, the sources of ideas for new products, and new processes.

We also talked about why productivity-increasing changes were introduced. Among the reasons we heard were: pressure or requirements from industrial customers, the impact of a merger or of being acquired, and the pressures resulting from increased financial leverage or external competition. In short, the sources of change included a wide range of things that seemed obvious to us — after someone told us.

The group visits — our pin factory trips — persuaded many of the participants of the potential value of looking and of asking questions. That led to a substantial number of individual studies, by single researchers, or pairs of researchers who had never done site based research before. The papers in today's session and in this volume are about some of those studies. A full list of NBER Working Papers describing these studies appears on the NBER's web site (www.nber.org). A reference on the NBER home page directs users to the Sloan Foundation Productivity Project. The papers and their abstracts can then be read and downloaded.

When I have described this project to non-economists, they were invariably surprised that the process of visiting companies, looking at production, and asking questions is an unusual part of economic research. It seems like such a natural thing to do. But as economists all know, it is unusual. We economists are generally accustomed to getting our insights by reading economic literature, going to seminars, and thinking hard about problems. We elaborate these insights in more or less formal models and then sometimes test these theories with aggregate statistics or micro data.

But we rarely go and look and ask. I think that is a pity. Looking and asking provide insights and suggest hypotheses — and can shoot-down wrong ideas — in ways that go beyond introspection and reading. Let me offer a personal example that has changed my thinking.

An Example

When our pin factory groups visited a manufacturing plant — whether it was Gillette's facility for making razor blades or a Ford assembly plant or a steel mill — the most striking thing was that there were virtually no workers. Most of the workers in the Gillette plant (where billions of blades are made) are maintenance engineers. In the steel mill, the workers were mostly in control rooms guiding the process by computer. And in both cases, they were very few in number.

One of my favorite questions to managers on these visits was therefore: "What percent of your costs are represented by these production workers?" As economists, we tend to think of labor's share of costs as about two-thirds or three-fourths. But the typical answer would be, "Oh, about 10 percent."

Now further questioning would bring out that one reason for the low labor share is that a significant part of the cost is purchased material that embodies labor inputs in other companies. But that leads to the natural follow-up question: "Leaving out purchased inputs, what percent of costs are represented by these production workers?" The typical answer was about 20 percent, far less than the share of costs that we normally associate with labor.

Well that tells you a lot — and it raises a lot of questions. As we probed we learned that most of the labor costs in a manufacturing firm are for non-production workers like product designers, accountants, sales people, other managers, and so on. An important implication of this is that raising the productivity of the production workers does relatively little to raise overall productivity. The big scope for productivity gains is with the non-production workers. That may help to explain why the information technology revolution may be a major source of the recent productivity improvements.

The more general lesson that I have taken from the Pin Factory experience is that plant visits and discussions with management can be an interesting and useful source of hypotheses. I emphasize hypotheses — not definitive answers. The next step in an economic research project is to elaborate these ideas and test them with data.

Seeing a company and talking with managers is likely to lead to a more interesting and informed approach to the data. But testing is important in order to know that the specific example that you've seen or the answer that you've heard is not just a special case — or indeed a misleading answer. When questions go beyond a simple fact — like the percentage of payroll outlays in total costs — there is room for ambiguity and for outright deception (perhaps self-deception) on the part of the person answering.

So the researcher has to ask: Are they telling me the truth? Or are they trying to paint a favorable picture of the way the company works — a more rational picture of the innovation process or a more favorable story about the effectiveness of compensation incentives. That's why looking and asking must be seen as a source of insights and hypotheses and not of definitive results. Looking and asking is not an end in itself — but it can be a very useful beginning.

Cambridge MA
January 2000

ECONOMISTS AND FIELD RESEARCH: "YOU CAN OBSERVE A LOT JUST BY WATCHING"

Susan Helper*

Forthcoming, American Economic Review, May 2000

Modern economics began with Adam Smith's visit to a pin factory, which helped him to explain how the division of labor worked (1985, p. 6). However, not many economists today do fieldwork, which involves interviews with "economic actors" and visits to the places they live and work. To help economists get out more, the Sloan Foundation in 1994 funded the National Bureau of Economic Research to promote field research among economists via plant tours, conferences, and commissioned papers.

I start by discussing how fieldwork can improve economic research, drawing largely on interviews with participants in the NBER/Sloan project. I then describe techniques that can improve economists' field research. (For an expanded version of this paper, see www.nber.org/sloan).

I. Fieldwork complements other methods.

Economists today typically do research using econometrics and mathematical modeling. These techniques have many strengths, but share the weakness of distance from individual economic actors. In contrast, field research allows direct contact with them, yielding several advantages:

Researchers can ask people directly about their objectives and constraints. It is not always easy to figure out someone's incentives or strategies by looking only at outcomes. For example, Jim Rebitzer wondered why many professionals complain about long hours, yet few firms offer the option of short hours. In talking with lawyers, he learned that partners found it difficult to decide whom to promote in order to maximize their incomes. "One comment that stuck with me was a partner saying about an associate, 'She does really good work, but I wonder, does she like money enough?' That is, he wanted to know, will she work *really* hard?" Comments by this partner and others implied that they used work hours as a proxy for this propensity to work. This insight led Renee Landers, Rebitzer, and Lowell Taylor (1996) to build a model and collect survey data which suggested that these observability constraints led to incentives to work inefficiently long hours.

Fieldwork allows exploration of areas with little pre-existing data or theory. I started my dissertation research thinking I would look at automakers' make/buy decisions. But when I started interviewing and

* Economics Department, Case Western Reserve University, Cleveland OH, 44106. I thank Martin Feldstein and Adam Jaffe for involving me in the NBER/Sloan program, and David Levine for extensive discussions. I also received much useful advice from those cited herein, Zvi Griliches, Carol Heim, and Don Siegel. The quote is from Yogi Berra (nd).

reading trade journals, I realized that important changes — not reflected in the existing literature — were occurring on the "buy" side. U.S. automakers were moving from adversarial deals (in which they "would steal a dime from a starving grandmother," one supplier said) to "voice" relationships in which they worked with suppliers to improve performance. My qualitative study (1991) suggested that information exchange and commitment were important determinants of supplier performance, leading me to collect survey data on these factors. One finding was that voice relationships were associated with more cost reduction, but only if complementary policies were adopted (Helper, 1999).

Fieldwork suggested to Zucker, Darby, and Brewer (1998) that the number of gene-sequence discoveries was a good proxy for intellectual capital in biotechnology; their regression results were consistent with intellectual capital being the main determinant of the location and growth of biotechnology firms.

Fieldwork facilitates use of the right data. Ichniowski, Shaw, and Prennushi (1997) used interviews to determine that steel finishing lines had homogeneous technology and that there were enough such lines to allow econometric investigation of the impacts of innovative human resource policies unconfounded by technology or industry differences. They then visited 45 plants to collect production data (interviewing managers to insure compatible measures across plants) and observed what human resource practices were in place.

Fieldwork provides vivid images that promote intuition. Edward Lazear said of his work on the change from time-rates to piece-rates at Safelite Autoglass (1996), "It's one of my most-cited papers — I think it's because everyone can imagine those guys working harder to install windshields once they're on piece-rates, and it's an image they remember a lot more than the regression coefficients."

Because of fears about the reliability of field methods (discussed below) some economists get ideas from the field, but do not discuss their fieldwork in their published articles. But this tactic causes us to lose the vividness that is a principal benefit of fieldwork. Understanding the setting can help explain differences in findings between cases, by making clear the mechanism by which variables are linked. For example, while Lazear found that a move to piece-rates increased profits at the auto-glass installer, Richard Freeman and Morris Kleiner (1998) found that a change *away* from piece-rates increased profits at a shoe manufacturer. Understanding the production process at the two firms is key to making sense of these results: while both papers found that productivity was higher under piece-rates, time-rates at the shoe firm facilitated the introduction of a new production process that brought reduced inventories and faster new product introduction.

Many of these insights can be translated into the language of econometrics or theory. It is possible that economists using only these methods could have generated these insights, but in fact, they didn't. Fieldwork offers a new source of inspiration, one that is complementary to more conventional methods. "It's important to go out and discover the facts for yourself," said Ronald Coase, who developed his ideas about the "nature of the firm" (1937) during a year of visits to firms throughout the United States.

II. Many criticisms of fieldwork can be answered with improved methods.

Many economists remain skeptical of qualitative research, fearing that it is not objective, replicable, or generalizable. "We don't have standards for what good fieldwork is," said Eli Berman. "In econometrics, we know to look for things like identification and specification issues, but what are the analogues in field work? How is it different from journalism?"

Better techniques can alleviate these problems. "Economists think that while econometrics requires years of training, field research is easy," said Rebecca Henderson. "But we need to pay just as much attention to things like careful research design and sample selection as we do in quantitative research." Below are tips from economists and others experienced in field methods. (See Robert Yin (1989), Robert Thomas (1994) and www.nber.org/sloan for additional suggestions. David Lodge (1990) presents information in novel form.)

Objectivity. A way for economists to avoid confirmatory bias is to test hypotheses coming from several competing theories. (Economists often go into the field with hypotheses to test. In contrast, disciplines like anthropology emphasize understanding the world as their informants do.) One source of alternative theories is to let respondents talk, even when you think they are getting off the subject. You will probably learn something you never thought to ask about and besides, getting to tell their story is part of their payoff for talking with you (Michael Piore, 1979).

It is important to do more than look to confirm or discredit preconceived hypotheses. Claudia Goldin said about a 1995 visit to Joseph Pollak, an auto-parts manufacturer in Stoughton, MA: "I didn't have any particular expectations going into the plant, but I remember vividly looking down from a mezzanine from which you could see the whole shop floor. As I looked down, I realized I was observing — in one moment — the transition from 19th century technology to 20th century technology. Just scanning across the room, I could see the relative increase in the demand for skill. The 'continuous process' machinery required lots of skilled labor to set up the machines and mechanics to maintain them; there were few operators...The old-fashioned areas [making similar parts], however, were filled with semi-skilled workers and almost no skilled workers. The scene sparked my imagination and I wrote two papers (Goldin and Lawrence Katz, 1996, 1998). For years I had been reading the history of technology, but it wasn't until I went to Pollak that I made the connection that adoption of continuous-process technology was complementary to skill."

Another way to enhance objectivity is to "triangulate," asking questions so that answers can be checked against information from other interviewees, company documents, trade journals, and so on. Site visits also provide a way to get a perspective other than the one presented by interviewees. A dramatic example of this came at our 1995 visit to LTV Steel in Cleveland, when workers stopped us to explain their anger at management's participation in a non-union joint venture in Alabama, which they thought violated the spirit of union-management partnership. Managers hadn't mentioned this dispute to us in their presentation on determinants of plant performance, even though it had caused the union to pull out of (productivity-enhancing) employee-involvement programs.

This incident also illustrates the value of interviewing people in a variety of positions. High-level people will give you an overview of the firm's intended strategy; low-level people offer detailed examples of the incentives and constraints they actually face.

Letting respondents tell their story does not mean taking everything they say at face value — you should be as skeptical of their statements (and as appreciative of them) as of any other data. Phrase questions concretely, and ask for examples. Until I asked for examples of trustworthy behavior in customer-supplier relations, I didn't realize that the definition affected the measured relationship between trust and performance. For example, one manager said, "I have a very trusting relationship with the plant I supply. If they find defects, they'll call us up and we can fix it, without anything ever showing up in the records." Another manager gave a quite different description of a trustworthy customer: "They are incredibly strict on quality — they'll send back a whole lot if even one part is defective. But they're always there to teach us, to make us better." The first manager felt little incentive to improve from his "trustworthy" customer; the second one felt a strong incentive.

Replicability. Often interviewees won't talk freely unless they are promised confidentiality, making it difficult for other researchers to replicate by interviewing the same people. However, there are other ways to enhance replicability. Rebecca Henderson and Iain Cockburn (1994) coded interview transcripts to create a dummy variable measuring whether pharmaceutical firms had "pro-publication" policies for their scientists. (This variable was a significant determinant of firm's research productivity. It is hard to imagine how else they could have obtained this information, other than by asking the actors directly about their incentives.) The more clearly described the coding, the more other researchers will be able to replicate it elsewhere.

Care in writing up results can increase replicability and readers' confidence in qualitative findings. For example, some field-research papers present arguments about what drives their regression results, and then say "interviews confirmed these findings," without providing information about how interviewees were selected, or enough information about what they said for readers to judge for themselves. This is like asking readers to believe a summary of econometric results without tables of regression coefficients. (See David Levine (1993) on describing selection of interviewees, Stephen Barley (1986) on using quotes as evidence, and Kathleen Eisenhardt (1989) on using "textual tables" to present this information concisely.)

Some changes at journals could alleviate the space problems that cause economists to truncate description of their fieldwork. Authors could use journal web sites to post information about field findings that didn't fit in the published article, and more journals could accept stand-alone qualitative analyses (the *Journal of Financial Economics, Journal of Industrial Economics,* and *Industrial Relations* already do).

Generalizability. How can one be sure that one's conclusions go beyond the few firms one has visited? Careful selection of cases according to theoretical principles, and use of cross-case analysis can alleviate this problem (Eisenhardt, 1989). If one is building theory, in-depth understanding of a few cases may be appropriate. If one wants to use observations from plants as data, visiting a large number of them is quite convincing (Ichniowski *et al*, 1997). As more cases are done on a topic, we will better understand what is specific to each setting, as the piece-rate example shows. That is, the solution to the generalizability problem is to do more field research, not less!

No methodology is perfect. "Regressions also have serious problems of generalizability (they predict poorly out of sample), subjectivity (researchers may stop specification searches when their favorite *t* statistic rises over 2), and measurement error (critical concepts like 'income' and 'capital' are very poorly measured). This is why research is hard — and why we should believe only findings obtained with multiple methodologies," said Levine.

In summary, good field research should be like good journalism in containing accurate, vivid examples. It should go beyond journalism in explaining potential biases in selection of cases, describing the construction of concepts used, and having a theoretical starting point and/or outcome. In conclusion, field research can make us better economists — whatever our current technique — by increasing our understanding of the objectives, constraints, and incentives that economic actors face.

REFERENCES

Berra, Yogi. "Famous Yogi Berra Quotes." http://www.yogiberraclassic.org/quotes.htm, nd.

Barley, Stephen. "Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners." *Administrative Science Quarterly,* 1986, *31,* pp. 78-108.

Coase, Ronald. "The Nature of the Firm." *Economica,* 1937, *4,* pp. 386-405.

Eisenhardt, Kathy. "Building Theories from Case Study Research." *Academy of Management Review,* 1989, *14*(4), pp. 532-550.

Freeman, Richard and Morris Kleiner. "The Last American Shoe Manufacturers." NBER Working Paper No. 6750, October 1998.

Goldin, Claudia and Lawrence Katz. "Technology, Skill, and the Wage Structure." *American Economic Review,* May 1996, *96*(2), pp. 252-257.

_______. "The Origins of Technology-Skill Complementarity." *Quarterly Journal of Economics,* August 1998, *113,* pp. 693-732.

Susan Helper. "Strategy and Irreversibility in Supplier Relations." *Business History Review, 65*(4), Winter 1991, pp. 781-824.

_______. "Complementarity and Cost Reduction." NBER Working Paper No. 6033, May 1997, revised 1999.

Henderson, Rebecca, and Ian Cockburn. "Measuring Competence? Exploring Firm Effects in Pharmaceutical Research," *Strategic Management Journal,* 1994, *15,* pp. 63-84.

Ichniowski, Casey, Giovanna Prennushi, and Kathryn Shaw. "The Effects of Human Resource Management Practices on Productivity." *American Economic Review,* May 1997, *87*(3), pp. 291-313.

Landers, Renee, James Rebitzer, and Lowell Taylor. "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms." *American Economic Review,* June 1996, *86*(2), pp. 329-348.

Lazear, Edward. "Performance Pay and Productivity," NBER Working Paper No. 5672, July 1996.

Levine, David I. "Fairness, Markets, and Ability to Pay: Evidence from Compensation Executives." *American Economic Review,* December 1993, *93*(4), pp. 1241-1259.

Lodge, David. *Nice Work.* New York: Penguin, 1990.

Piore, Michael. "Qualitative Research Techniques in Economics." *Administrative Science Quarterly,* December 1979, *24,* pp. 560-569.

Smith, Adam. *An Inquiry into the Nature and Causes of the Wealth of Nations.* New York: Modern Library, 1985.

Thomas, Robert. *What Machines Can't Do.* Berkeley: University of California Press, 1994.

Yin, Robert. *Case Study Research.* Beverly Hills: Sage Publications, 1984.

Zucker, Lynne, Michael Darby, and Michael Brewer. "Intellectual Human Capital and the Birth of US Biotechnology Enterprises." *American Economic Review,* March 1998, *88*(1), pp. 290-306.

WHO BENEFITS MOST FROM EMPLOYEE INVOLVEMENT: FIRMS OR WORKERS?†

Richard B. Freeman* and Morris M. Kleiner**

Employee involvement (EI) programs are the leading edge form of personnel and labor relations in the United States. While many managers believe that these programs raise productivity and profits, the statistical evidence that EI improves the performance of firms is equivocal. The coefficients on measures of EI in production functions are usually positive but often insignificant or small (Commission on Labor-Management Relations, 1994, Chapter 2; Cappelli and Neumark, 1999) or contingent on other factors (Ichniowski, Shaw, Prennushi, 1997; Black and Lynch, 1997). A detailed case study of EI has further confirmed these small effects that were found in large datasets (Kleiner, Leonard, and Pilarski, 1999).

If EI programs do not greatly affect productivity, then why does business think so highly of them? In this study, we argue that the main beneficiaries of EI are workers and managers. We estimate the effects of EI on productivity using panel data on firms and the effects of EI on workers using a survey of employees; we find that EI barely affects firm productivity but substantially improves worker well-being. We offer two explanations for this result.

Firm-based Productivity

The data for our analysis of firm productivity comes from a 1993 mail survey by the Society for Human Resource Management (SHRM) conducted by Cheri Ostroff. The SHRM survey was sent to 3402 firms, of whom 373 responded, giving an 11 percent response rate; that is similar to other studies that attempt to measure human resource practices by mail survey (Kato and Morishima, 1996). We matched these data with firm-level data on production and financial outcomes from Compustat for the years 1983 to 1993 and obtained 273 usable observations. We then used a difference-in-difference design to compare the performance of firms year by year as their EI programs changed; contrasted firms by EI status in the final year of the survey; and examined changes in productivity and EI practices among firms over a decade.

The SHRM survey asked about eight EI practices: self-managed work teams; worker involvement in the design of EI programs; TQM;

† We thank Hwikwon Ham for his excellent research assistance.

* Professor of Economics, Harvard University and Program Director, Labor Studies, National Bureau of Economic Research, 1050 Mass. Ave., Cambridge, MA 02138

** Professor of Public Affairs and Industrial Relations, University of Minnesota, and Humphrey Institute, 260 Humphrey Center, 301 19th St. South, Minneapolis, MN 55455, and NBER.

committees on productivity; worker involvement in work processes; suggestion or complaint systems; information-sharing with employees; and opinion surveys. The survey also asked whether the firms' use of the practice was "very great," covering 80 percent or more of jobs/workers; "great," 60-79 percent; "moderate," 40-59 percent; "some," 20-39 percent; or "little," 1-19 percent of workers; and when the practice was implemented: 10 or more years earlier; 5-9 years earlier; 1-4 years earlier, or within a year.

Both the number and use of EI practices grew greatly between 1983 and 1993: in 1983 nearly half of the sample had no EI practices and only four companies had all of the practices, whereas in 1993 virtually all had some practice and 94 had all of them. In addition, firms extended the coverage of their programs over time. By 1993, practices introduced earliest, such as information sharing or suggestion/complaint systems, were used more intensively than practices introduced later, such as opinion surveys or giving workers a role in designing work processes. Because firms that use rarer practices generally have the most common ones, the practices fall into a reasonably well-ordered single dimensional "EI" scale, per Guttman scaling or the more general Rasch-type models (Bartholomew, 1996).

To measure firm EI activity over time, we constructed a rating scale of employee involvement based on the existence, intensity, and timing of the eight practices. We gave each practice a measure from 0 to 5, depending on its presence and extent, and then summed the eight measures. Firms that applied every practice to almost all workers in a particular year received the value 40. Firms with no practices received the value 0. If the practices fit a perfect (Guttman) scale, firms having the lowest scores would have the fewest/least intensive practices while firms with higher scores would have the practices/intensity of those with lower scores and then some. They fit this pattern well enough to make the scale meaningful (Freeman, Kleiner, and Ostroff, 2000).

Because the survey asked for a range of years when the program was implemented, we cannot identify the precise years when the firm has a program and when it does not. Our solution is to approximate the existence of a program in a given year by assuming that the program had a uniform probability of being introduced in one of the years in the reported range. For a program introduced 3-5 years ago, this leads to a one third chance the program existed 5 years ago, a two thirds chance it existed 4 years ago, and a singular chance that it existed 3 years ago. These figures estimate the probability that a firm had a program in year *t*. Because the survey asked for intensity of use only in 1993, we estimated the intensity of use in earlier years by exploiting the fact that intensity of use is highly correlated with the length of time a program has been in place. We regressed the intensity of use on the length of program life for all companies and programs taken together in 1993 (for the details of this estimation, see Freeman, Kleiner, and Ostroff, 2000) and used the estimated coefficients to predict the unobserved intensity in years prior to 1993.

Given the estimated probability that a firm had a particular EI program in a given year (p') and the estimated intensity of use (U') in that year, we calculated a scale of EI activity for each year and program as the product p'U' and then summed the values across programs for each firm in each year to obtain an EI variable that measures the number of programs and their intensity of use in a given year. Finally, we estimated the effect of EI using the production function:

(1) $\ln Q = a + b \ln K + c \ln L + d\, EI + YR + FIRM + u$

where Q = sales; L = employment; K = book value of assets, from Compustat. FIRM is a vector of firm dummy variables; YR is a vector of year dummy variables; EI is our scale; and u is a residual. With firm and year dummy variables, identification of an EI effect comes from the differential variation of EI over time across firms with productivity.

Table 1 gives the results of our analysis using two different estimating procedures: ordinary least squares and median regressions. The OLS results in line 1 show little or no impact for EI. By contrast, there is a small noticeable effect of the EI scale on ln sales in the line 2 median regression, indicating that results are sensitive to the mode of estimation. These regressions make maximum use of our data but suffer from the possible problem that EI programs require considerable time to bear fruition. The regressions in lines 3-4 deal with this by focusing on the final year of our survey, 1993, by which time many of the programs should have been in a mature state. Here too the estimated coefficients on EI are negligible. Finally, in lines 5 and 6 we regress changes in sales on changes in inputs and the change in EI over the 1983-1993 decade, with the data transformed into average annual changes so that we could include firms for whom we did not have data going back to 1983. Put differently, in this regression we took the longest period for which we had data and annualized the changes. Again, we find no EI effects.

Many analysts believe that a linear specification of the EI effect is incorrect. Rather, EI has a substantial non-linear effect on productivity, so that firms that introduce one involvement activity may gain nothing or even lose because "a single tree does not a forest make" whereas firms that introduce a full spectrum of complementary policies gain from EI (Ichniowski, Shaw, and Prennushi, 1997). We mined our data in search of non-linearities, but found little evidence that any such patterns are confounding the EI impact. In sum, our data show that EI has little or no effect on productivity, with a positive effect in only one median regression. Perhaps a larger data set might uncover something that our data fail to reveal, but recent work by Capelli and Neumark (1999) support our finding in a larger file.

Table 1: Regression Coefficients (Standard Errors) on Production Function Estimates of Effects of EI Scale, Employment, and Assets on LN Sales

	LN Emp	LN Assets	EI Scale	R^2/ Pseudo R^2
Annual Data (n=2127, with year and co. dummies)				
OLS	.48	.52	-.000	.93
	(.02)	(.02)	(.002)	
Median	.41	.57	.003	.82
	(.01)	(.01)	(.001)	
1993 Cross Section (n=237)				
OLS	.46	.55	.000	.92
	(.05)	(.05)	(.005)	
Median	.42	.56	-.003	.81
	(.03)	(.03)	(.003)	
1983-93 Change (n=229)				
OLS	.42	.56	-.001	.84
	(.05)	(.05)	(.002)	
Median	.45	.48	.000	.61
	(.03)	(.03)	(.001)	

SOURCE: COMPUSTAT, Standard and Poors, various years.
1983-93 regressions include all years for which we have data; 1993 regression covers 1993 cross section; 1983-93 change includes all companies for which we could get at least 2 years, with changes calculated as average annual changes over the longest period for which data exists. The 1993 and 1983-93 regressions include 7 industry dummy variables, which have no noticeable impact on the results.

Effects on Employees

To find the effects of EI programs on workers, we turn to the Workplace Representation and Participation Survey, WRPS, (Freeman and Rogers, 1999). This is a nationally representative survey of some 2400 workers in firms with over 25 employees that focuses on employee attitudes toward various labor practices. The WRPS is a cross section, so that we cannot follow workers from a firm with an EI program to another firm or conversely. But we can contrast EI participants to non-participants in companies with programs and to employees in companies without programs, and can determine how employees view their firms' EI program.

Table 2 presents cross-tabulations that summarize the responses of non-managerial employees by EI status to questions relating to employee decisionmaking at work and attitudes toward the firm, and presents the views of EI participants toward the program. More detailed analysis controlling for covariates gives comparable results, justifying our focus on the cross-tabs.

Panel A compares the percentage of workers who report "a lot" of involvement in the company decisions that affect their work life in six different areas. In each area, workers on EI committees report greater involvement on decisions than other employees. Averaging across areas, EI participants have a 14 percentage point edge over non-participants in firms with programs and a 17 point edge over employees in firms without programs. The similarity between the responses of non-EI participants in firms with EI programs and those in firms without any program indicates that the EI impact is not a simple "good company" effect. Note that the survey asked about workers' influence on the job prior to the module on EI, so responses are not affected by questions about EI. In some areas the differences are striking. A substantial proportion of EI workers have a lot of direct involvement in setting goals for their work group, deciding what training is needed, or how to work with new equipment or software.

Table 2: Involvement in Decisions and Attitudes of Non-Managerial Workers, by EI Status

	Firm Has EI Program		No EI
	Partici-pates	Does Not Parti-cipate	
Percentage Answering Have "A lot" of Direct Involvement in...			
Deciding how to do job	68	52	50
Setting goals for work group	44	29	24
Setting work schedules	39	30	24
Deciding what training is needed	43	23	23
Setting safety standards and practices	44	31	30
Deciding how to work with new equip/software	38	26	23
AVERAGE across Areas	46	32	29
Attitudes toward Work			
Percent very satisfied with influence on decisions that affect job/work life	34	19	19
Percent who look forward to work	74	63	61
Percent who feel a lot of loyalty to firm	63	42	39
Percent who trust firm to keep promises a lot	49	36	30
Percent who rate system for resolving workplaces problems as very effective	38	26	22
Perceived Impact of EI (asked only of participants)			
Percent who benefitted by greater influence on job	79		
Percent who benefitted from better wages or benefits	36		
Percent for whom elimination of EI would be bad/very bad	71		

SOURCE: Tabulated from Worker Representation and Participation Survey, Detailed Tabulations, See http://www.nber.org/data_index.html (Freeman and Rogers, 1994).

Panel B gives the responses from questions relating to satisfaction or attitudes toward work. A much larger proportion of EI participants report themselves as very satisfied with the "influence (they) have in company decisions that affect (their) job or work life" compared to other workers. In addition, proportionately more participants look forward to going to work (as opposed to not caring one way or the other or wishing they did not have to go); are more loyal to their employers; trust that their company will keep its promises to them and other workers; and view their firms' program for dealing with the workplace problems as very effective.

Panel C examines the experiences of EI participants toward their firms' program. Over three quarters say that they personally benefited by gaining greater influence on the job; over a third say they obtained better wages/benefits. Most important, the vast majority of EI participants said that getting rid of their firms program would have bad or very bad effects on them personally.

Interpretation

What explains the negligible effect of EI in our (and other) production function analyses compared to the strong effect that workers report EI makes on their working lives? One possible explanation is statistical. Firm-based studies of EI like ours may fail to find productivity effects because the "true" effects are too small for the production analysis to uncover with any degree of certainty given the sample size, unexplained variation in productivity or changes in productivity, and imperfect measures of EI. Consider, for example, the situation in which a correctly measured EI program raises productivity by 3 percent over, say, a decade. This is modest given the variability of productivity for a firm over time, but sufficiently large as to raise profits considerably if the gain accrued primarily to capital. In our sample the standard deviation of the residual of changes in ln sales regressed on changes in ln employment and changes in ln assets over the decade and industry dummies is around 0.46 across firms. Assuming that the true population variance in productivity is of comparable magnitude, the *critical effect size* for detecting the EI impact is about 0.065 (= .03/.46). Then, for a 5 percent two-tailed t test on the EI coefficient, with 95 percent power, we would need a sample of about 3,124 observations (interpolated from Kraemer and Thiemann, 1987, p. 109). But this calculation assumes that EI is measured correctly. If the signal-to-noise ratio in the EI variable was non-negligible, we would need a larger file to obtain a significant relation and would have to correct for the measurement error to uncover the true effect of EI. Alternative models for the critical effect size will yield somewhat different results, but the general point is clear: if productivity varies a lot, as it does, and if EI effects are modest, which seems likely, and if EI is measured with error, which is surely true, we need more studies with better measures of EI to get the right answer or enough studies to permit a careful meta-analysis.

The second explanation is that in fact EI is an innovation whose economic gain accrues largely to workers (and managers) rather than the firm and shareholders. Why might this be? One possibility is that an increasingly educated and knowledgeable work force wants more independent decisionmaking at their job (even if their decisions will largely mimic those of a more authoritarian management). Since EI has no adverse effect/slight positive effect on the bottom line, firms will offer it to please their workers. Some may even be able to get a compensating differential for the practice, as do firms that offer PhD researchers greater opportunity for independent work (Stern,1999). Modern information technology may also be complementary with employee decisionmaking, so that the growth of EI reflects technological change embodied in workers rather than managerial innovation. Whatever the explanation, since EI has at worst a non-negative effect on productivity and a positive effect on the lives of workers, it is a net benefit to the U.S. labor market.

REFERENCES

Bartholomew, David. *The Statistical Approach to Social Measurement,* Academic Press, London, 1996.

Black, Sandra E. and Lisa M. Lynch, "How to Compete: The Impact of Workplace Practices and Information Technology on Productivity," NBER Working Paper No. 6120, August 1997.

Cappelli, Peter, and David Neumark, "Do 'High Performance' Work Practices Improve Establishment-Level Outcomes?", NBER Working Paper No. 7374, October 1999.

Commission on the Future of Worker-Management Relations. *Fact Finding Report.* Washington, DC: U.S. Department of Labor and U.S. Department of Commerce, May 1994.

COMPUSTAT, Standard and Poors, various years.

Freeman, Richard, Morris M. Kleiner, and Cheri Ostroff. "The Anatomy and Effects of Employee Involvement," working paper, Harvard University and the University of Minnesota, 2000.

Freeman, Richard and Joel Rogers. *What Workers Want,* Cornell University Press and the Russell Sage Foundation, Ithaca, NY 1999.

Freeman, Richard and Joel Rogers. Worker Representation and Participation Survey (WRPS). http://www.nber.org/data_index.html 1994.

Ichniowski, Casey, Giovanna Prennushi, and Kathryn Shaw. "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," *American Economic Review,* June 1997, pp. 291-313.

Kato, Takao, and Motohiro Morishima. "The Productivity Effects of Human Resource Management Practices: Evidence From New Japanese Panel Data," working paper, December 1996.

Kleiner, Morris, Jonathan Leonard, and Adam Pilarski. "Do Industrial Relations Affect Plant Performance? The Case of Commercial Aircraft Manufacturing," NBER Working Paper No. 7414, November 1999.

Kraemer, Helena and Sue Thiemann. *How Many Subjects?* Newbury Park, CA: Sage Publications, 1987.

Stern, Scott. "Do Scientists Pay to Be Scientists?", NBER Working Paper No. 7410, October 1999.

IS COST CUTTING EVIDENCE OF X-INEFFICIENCY?[†]

Severin Borenstein* and Joseph Farrell**

X-inefficiency is surely among the most important topics in microeconomics. Yet, economists have found it difficult to study. If a given level of X-inefficiency were inevitable and changeless, it would be of little interest (indeed, would not really deserve to be called X-inefficiency at all). So our attention should focus instead on actual and potential changes in X-inefficiency: that is, on causes of changes, internal to a firm, that shift the firm's cost function. We explore the use of firms' "cost-cutting" announcements to study the causes of *changes* in X-inefficiency.

Cost cutting announcements by large corporations are made frequently and are reported in the business press. One might be tempted to interpret these announcements as indicating efforts to reduce X-inefficiency, and indeed we think that a substantial number of them are just that: our discussions with managers confirm that finding and trimming "fat" within the corporation is an important, yet somewhat intermittent, activity.

Yet, it is also clear that not all cost-cutting announcements concern X-inefficiency. Some, perhaps many, may instead be re-optimizing (input and/or output) quantity responses to changes in exogenous factors, such as the prices of inputs or outputs. The announcements seldom make it clear whether the activity is re-optimization or fat trimming.[1]

We aim to explore the existence and nature of fat trimming within a firm, and how one might distinguish this from normal re-optimization. In particular, we seek evidence bearing on a central hypothesis in the informal theory of X-inefficiency (Harvey Leibenstein, 1966), with very broad support in news reports, in popular belief, and in our interviews with managers. This *fat hypothesis* is that a firm is most apt to cut costs to reduce X-inefficiency when it is under financial pressure. This hypothesis, if correct, has implications both for firm strategy and for competition policy.

[†] Erin Mansur, David O'Neill, and Evan Rose provided outstanding research assistance. We also thank many managers in the gold mining and oil industries for discussing cost cutting with us.

* Professor of Business, Haas School of Business, University of California, Berkeley, CA 94720-1900 and NBER.

** Professor of Economics, Department of Economics, University of California, Berkeley, CA 94720-3880.

[1] There are other data problems, including the following: The division(s) in which the cost cutting is to occur are usually not reported precisely. Magnitudes are not systematically reported (and usually are projections). Multiple reports of a single cost-cutting effort, sometimes months apart, are common.

While Olivier Blanchard *et al* (1994) studied the effects of idiosyncratic cash shocks, it seems desirable to find a more systematic source of wealth shocks. One such source is exogenous changes in the prices of competitively supplied inputs or outputs. In our 1999 paper, however, we explain that not only must the price changes be out of the firm's control, but they must also leave the firm's production possibilities unchanged. For instance, increases in the price of oil caused by political instability raise the expected profits of U.S. oil companies on the oil they will be able to sell in any case, but might simultaneously indicate a reduced opportunity to explore for oil in the future. Likewise, if the technology for gold mining improved, then the price of gold would fall, but the availability of the new technology to firms under study would offset the price change and could result on net in a positive wealth shock (and increase optimal production quantities). One might assess the source of price changes by examining quantity changes (potentially differentiating between supply and demand shocks), news reports, or the technology involved.

Drawing inferences from cost-cutting announcements presents another problem, because a price shock that (say) lowers the firm's overall profitability is likely also to lower its marginal profitability. Thus, the *null hypothesis* that firms simply are re-optimizing likely predicts that the firm will reduce output, and presumably also reduce at least some inputs; this might be announced as "cost cutting." So, it would seem that cost-cutting announcements in response to adverse price changes fail to distinguish between the null and fat hypotheses.[2]

Knowing the elasticity of the firm supply curve might allow one to distinguish the hypotheses. If the supply curve is highly elastic then an output price change will induce relatively large changes in optimal quantities, with relatively small effects on firm wealth. In contrast, if the supply curve is highly inelastic, then a price change will have significant wealth effects, but will not induce much re-optimization. If supply elasticities vary across firms in an industry, one might test whether exogenous output price reductions cause greater cost cutting among elastic- or inelastic-supply firms. In the remainder of the paper, however, we turn to strategies that instead use information on more than one market.

I. Cost Cutting in Multi-Divisional Firms

The internal capital markets literature in corporate finance suggests another approach to diagnosing fat trimming. That literature has shown that divisions within the same firm in various industries cross-

[2] For example, consider a company that owns oil wells. When the price of oil falls, the null hypothesis predicts that company might reduce output because the marginal cost of producing from some wells now exceeds the price. But the price decline also lowers the wealth and cash flow of the firm, so the fat hypothesis also predicts cost cutting.

subsidize one another in financing investment. Owen Lamont (1997) shows that after the 1986 oil price crash many oil producing companies cut back investment in divisions *unrelated* to oil (or divisions whose marginal profitability would likely rise with oil price reductions).

Some authors ascribe such changes to principal-agent problems within the firm that cause managers to make negative-NPV investments with free cash flow, a form of fat in the firm (Michael Jensen, 1986). Others suggest that evidence of internal capital markets could reflect the firm's optimal response to inefficiencies in external capital markets. If the cross-subsidy reflects non-optimizing behavior within the firm, it need not be limited to investment decisions. Negative wealth shocks to one division in a firm could trigger fat trimming in other divisions.

Anecdotal evidence of this sort of corporate-wide belt-tightening is abundant. Our plan is to see if this effect can be documented empirically. We identify a number of potentially testable empirical hypotheses.

First, under the null hypothesis that fat trimming does not take place (or is unrelated to firm wealth) one would expect that a wealth shock to one division should not induce cost cutting behavior in another unrelated division. The fat hypothesis would predict that a negative wealth shock in one division will trigger fat trimming in other, even unrelated, divisions of the same firm.

Second, and closely related, under the null hypothesis, a division's cost cutting would be unrelated to the size and degree of diversification of the corporation within which it is situated. Under the fat hypothesis, wealth shocks in a division would have implications for the entire firm and the cost cutting behavior of a division would be less responsive to its own financial performance if it were located within a large conglomerate.

Two cautions are in order. First, if capital markets are imperfect, such cross effects could arise even under the null hypothesis, as in the corporate finance literature. Second, there is a problem with treating corporate structure as a natural experiment: one should ask why a firm chose the structure it did. For instance, many oil refiners are integrated into oil extraction; we should ask whether the reason might be synergies that would affect cost cutting of the re-optimizing type. Thus, if it were highly advantageous to refine "one's own" oil (which does not appear to be the case in fact), the quantity decision problem for an integrated refiner/extractor could be different from that of a similarly sized stand-alone refiner.[3]

Economies of scope, in general, could make it difficult to distinguish re-optimizing from fat-trimming types of cost cutting. A change that optimally lowers production in one division could, in the presence

[3] We say "could be different," because an integrated refiner that is buying its *marginal* crude oil externally faces a marginal decision that is likely to be identical to a firm that has no infra-marginal crude supplies at all.

of scope economies, raise marginal production costs in another division and, thus, optimally lower production in the latter division as well. Note, however, that this explanation requires scope economies *on the margin*. If the scope economies are only common fixed costs that are independent of scale of the divisions, the change in output of one division will not cause re-optimizing in the other division.

II. Vertical Integration as a Special Case

In a firm that is vertically integrated (for instance, a firm that both extracts and refines crude oil), a shock to the intermediate-good price (crude oil) may raise the profits and optimal scale of one division even as it lowers them for the other. Intuition might, at first, suggest that a fully vertically integrated firm (one that refines all of, and only, its own oil) is insulated from the wealth effects of a change in the price of the intermediate product. But the effect is actually more complex. There are, for instance, significant Ricardian rents in the oil extraction business that change one-for-one with the price of oil. The refining business, on the other hand, is considered very competitive. The effect of an oil price change on the profits from oil refining is likely to be comparatively short-lived and small. Thus, for a firm to be wealth-neutral with respect to supply-shock driven crude price changes it would have to maintain much larger operations in refining than in extraction.

A price change exogenous to the firm may come from demand or supply shocks. A change in the intermediate good price caused by a supply shift will indeed affect the upstream and downstream divisions of the firm in opposite ways, so vertical integration will provide some wealth insulation. If, for instance, the price of crude oil declines due to an oil field discovery that does not involve the observed firm, this will cause a negative wealth shock to the firm's oil extraction division. However, it likely also will raise the expected profits, and optimal production quantity, from the firm's oil refining business. Thus, if we observed cost cutting in the firm's refining business after an exogenous decline in the price of crude oil, this would be consistent with the fat trimming hypothesis and hard to square with the null hypothesis of optimizing behavior (absent significant scope economies).

In contrast, oil price changes attributable to demand shocks are apt to produce positively correlated wealth, and optimal output, effects in the upstream and downstream divisions. If a weak world economy pushes down oil prices, its effect on the returns to operating an oil refinery also is likely to be negative. Both divisions of the firm could then plausibly engage in cost cutting of the reoptimizing type, reducing output and laying off workers. Thus, demand-driven oil price shocks are not especially helpful for distinguishing between the null and fat hypotheses.

III. Cost Cutting in the U.S. Oil Industry

We have collected 122 cost-cutting announcements by major U.S. oil companies since 1984. Figure 1 shows the time series of monthly oil prices and cost cutting announcements. Because of the lag between price changes and the responses of firms (and to keep the scales of the two series in the same range), the cost-cutting series for each month is the sum of cost cutting announcements in that month and the five following months.

Figure 1 indicates that cost cutting announcements tend to follow oil price declines. As discussed above, however, this is just the beginning of an attempt to statistically discern fat trimming. Less quantitatively, the texts of some cost cutting announcements support the fat trimming hypothesis. For example, a March 1986 *Wall Street Journal* report on Amerada Hess, an integrated oil company, says that the company has announced it will respond to the plunge in oil prices by, among other actions, reducing output from its Virgin Island refinery since it has been losing money in its refining business.

Figure 1. Oil Prices and Cost Cutting

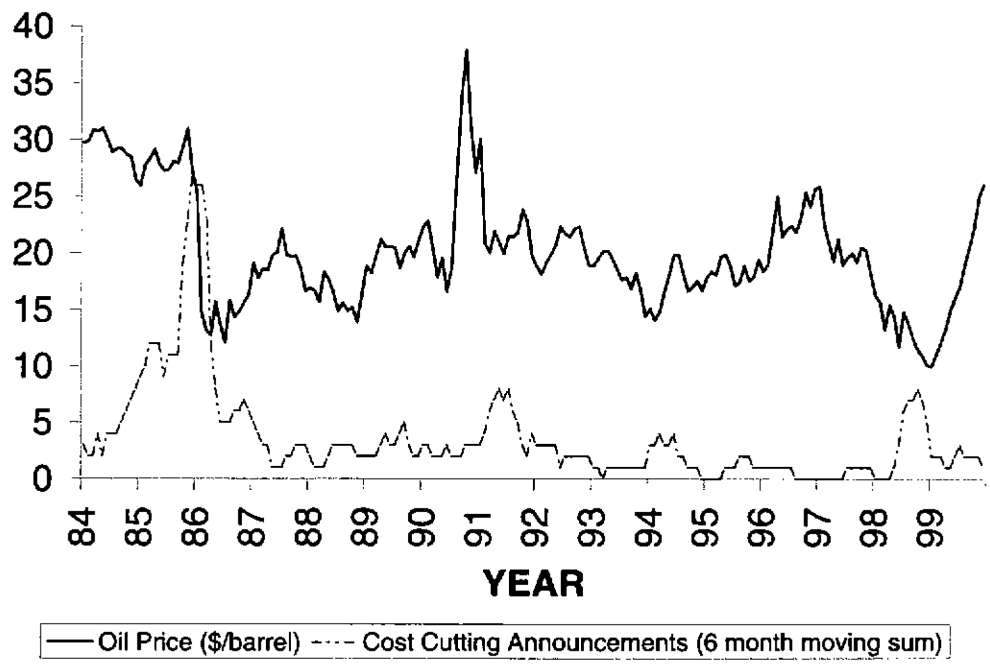

IV. Conclusion

We think that X-inefficiency merits much more empirical analysis. Our interviews with managers in two industries in which companies are subject to large wealth and profit fluctuations — gold mining and oil production/refining — strongly support the hypothesis that fat trimming occurs in response to wealth and profit declines. We hope to use a panel of observations on major U.S. oil companies to distinguish fat trimming, which conflicts with standard microeconomic analysis, from re-optimizing behavior that is central to standard microeconomics.

REFERENCES

Blanchard, Olivier, Florencio Lopez-de-Silanes, and Andrei Shleifer. "What do Firms do with Cash Windfalls?", *Journal of Financial Economics*, December 1994, *36*(3), pp. 337-360.

Borenstein, Severin and Joseph Farrell. "Do Stock Price Movements Reveal Profit Dissipation: An Investigation of the Gold Mining Industry." NBER Working Paper No. 7075, April 1999.

Jensen, Michael. "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers." *American Economic Review*, May 1986 (*Papers and Proceedings*), *76*(2), pp. 323-329.

Lamont, Owen. "Cash Flow and Investment: Evidence from Internal Capital Markets." *Journal of Finance*, March 1997, *52*(1), pp. 83-109.

Leibenstein, Harvey. "Allocative Efficiency versus X-efficiency," *American Economic Review*, June 1966, *56*(3), pp. 392-415.

Williams, John D. "Dividend News: Amerada Hess Omits Payout, Cites Price Drop." *Wall Street Journal*, March 6, 1986.

KNOWLEDGE SPILLOVERS AND PATENT CITATIONS: EVIDENCE FROM A SURVEY OF INVENTORS[†]

Adam B. Jaffe*, Manuel Trajtenberg**, and Michael S. Fogarty***

The "non-rival nature" of knowledge as a productive asset creates the possibility of "knowledge spillovers," by which investments in knowledge creation by one party produce external benefits, thus facilitating innovation by other parties. At least since Zvi Griliches's seminal paper on measuring the contributions of R and D to economic growth (1979), economists have been attempting to quantify the extent and impact of knowledge spillovers. One line of research of this type has used patent citations to identify a "paper trail" that may be associated with knowledge flows between firms.[1]

However, very little of this research has attempted to determine the modes or mechanisms of communication that actually permit knowledge to flow. Further, most of the work has simply assumed that citations or other proxies are sufficiently correlated with knowledge flows to allow statistical analysis of the proxies to be informative regarding the underlying phenomenon of interest.

This paper reports on a preliminary attempt to improve on this situation. The idea for this survey came from R and D managers whom we were interviewing to test whether the picture of knowledge flows produced by patent citations was consistent with the managers' impressions. One of these managers commented that he couldn't let us talk to the scientists who worked for him, but that he couldn't stop us from contacting them via the postal address that appears on their issued patents.

The survey results suggest that communication between inventors is reasonably important, and that patent citations do provide an indication of communication, albeit one that also carries a fair amount of noise.

[†] We thank Richard Anderson for his help in editing and administering the survey, and Melinda Miller and Margeret Lister Fernando for research assistance.

* Associate Professor of Economics, Brandeis University, MS 021, Waltham MA 02454-9110, and NBER.

** Associate Professor of Economics, Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, ISRAEL, and NBER.

*** Center for Regional Economic Issues, Wickenden 311, Case Western Reserve University, Cleveland OH 44106.

[1] Patent citations or references appear on the front page of a granted patent. They serve the legal function of identifying "prior art" upon which the current invention builds. For more detail, see Jaffe, Trajtenberg, and Rebecca Henderson, 1993.

I. Survey Design[2]

We surveyed two groups of inventors: one we asked about citations made in their patents to previous patents (the "citing inventor" survey); and one we asked about citations *received* by their patents from subsequent patents (the "cited inventor" survey). For the *citing* inventor survey, we asked a series of questions about the extent, timing, and nature of communication that they had had with the inventors of three previous patents. Two of these previous patents had appeared among the citations made by the surveyed inventor's patent. The third previous patent was a "placebo" patent that was not cited by the surveyed inventor, but which was granted in the same patent class and year as one of the actually cited patents. In the survey questionnaire this placebo was not identified or distinguished in any way; all three of the earlier patents were referred to as "cited patents."

For the *cited* inventor survey, we picked one of the actually cited patents about which the *citing* inventor was questioned, and contacted the primary inventor. This *cited* inventor was asked about communication with the *citing* inventor, and was also asked to form a judgement, based on his/her reading of the citing patent, about the likelihood that the citing inventor had used knowledge represented by the *cited* inventor's patent. Finally, both the *citing* and *cited* inventors were asked a series of questions about the economic and technological significance of their inventions.

We received 166 partial or complete responses to the *citing* survey, based on patents granted to these inventors in 1993. We received 214 partial or complete responses to the *cited* inventor survey, based on patents granted to these inventors between 1985 and 1993.

II. Results Regarding Extent, Timing, and Nature of Communication

We discuss here the responses to four questions asked of the citing inventors about their cited and "placebo" patents. One of the questions sought responses on a Likert scale to a question regarding the overall degree of familiarity of the citing inventor with the cited invention. For the patents that were in fact cited, 28 percent of the responses indicated a 4 or 5 on the Likert scale, indicating a significantly high familiarity; just under half of the respondents rated their familiarity at the low end of the scale. In contrast, over 80 percent of the respondents rated their familiarity with the "placebo" patent at the lowest possible level.

Another question related to *when* they learned about the "cited" invention. For the "true" citations, about 38 percent of respondents indi-

[2] The survey instruments were "pretested" on about 20 experienced inventors and the survey questions refined in response to their questions and comments. Space limitations preclude a detailed description of the survey and survey samples. For these details, see Jaffe, Trajtenberg, and Fogarty (2000).

cated that they had learned about the cited invention either before or during the development of their own invention. About one-third indicated that they had learned about it after essentially completing their invention, including cases where they learned about it during the preparation of their own patent application. A little less than one-third indicated that, despite the presence of the patent citation, they had not learned about the cited invention before receiving our survey. This is not surprising, because citations to inventions unknown to the inventor can be generated by the inventor's patent attorney or the patent office examiner.

A third question relates to the mode of knowledge spillover. Even for the "true" citations, only about 18 percent indicated that they had had either direct communication or had been exposed to some kind of presentation or demonstration of the cited invention. Another 18 percent indicated that they learned through "word of mouth" or actually had read the patent document itself. Consistent with the answers regarding timing, almost 40 percent indicated that it was the process of their own patent application that had caused them to learn of the previous invention.

The last question in this section, perhaps ambitiously, tried to get at the issue of the nature of assistance that the citing inventor may have received from the cited invention. Respondents were given a set of choices that we thought possible, and also invited to "write in" their own responses. Again, about 60 percent of the respondents indicated some specific way in which they had benefited from the "cited" invention; the single most common response was that the cited invention represented a concept that could be improved upon.

The "other" responses provided by the inventors also provide some insight into the nature of possible interactions. Examples include:

> "The technology from patent 1 was incorporated in the product which used my invention."
> "new market for our new technology!"
> "The other patents gave credibility to our idea — they showed our ideas were 'feasible' to the people not intimately involved in our idea."

Other explanations confirmed that many citations derive from the patent process and probably are not related to any spillover:

> "did not learn of patents before filing — therefore these patents were not a factor in our work"
> "a patent cited by the patent examiner with no direct ties to my patent"

Assuming that these responses can be taken at face value, they suggest that a significant, but not overwhelming, fraction of the "links" indicated by a patent citation correspond to some kind of spillover. Across the different aspects captured by each of these questions, typically one-quarter of the responses correspond to a fairly clear spillover; perhaps one-half of the answers indicate no spillover, and the remaining quarter indicate some possibility of a spillover. It is also clear that addition of citations by the inventor's patent lawyer or the patent examiner is the primary reason for citations to patents unknown to the inventor.

In order to test formally whether the respondents' answers relating to citations are different from those for the placebos, we constructed a composite spillover index for each "cited" patent, using the answers to all 4 questions. This index was constructed by consolidating the possible answers to each question to produce a score of 0, 1, or 2, and then adding these scores across the 4 questions. We then undertook an ordered probit analysis of this score, using as regressors variables that would seem likely to foster communication between the cited and citing inventors, variables that might foster the inventor's *remembering* that communication occured, other controls, and a dummy variable for whether the score pertains to a "true" citation as opposed to a placebo. Results of these analyses are presented in Jaffe, Trajtenberg, and Fogarty (2000).

Overall, the results confirm that citations can be interpreted as providing a (noisy) signal of spillovers. The difference in spillover score between the citations and placebos is quantitatively and statistically significant. The other variables generally have plausible and often significant effects. Overall, the spillover score is higher if the "cited" patent is more recent. Interestingly, analyses that separate the true citations from the "placebos" show that this combined effect mixes a significantly positive effect for the citations with a significantly negative effect for the placebos. For the citations, this is consistent with more recent patents being more useful, and older citations being more likely to be non-spillovers included by the lawyer or examiner. It could also reflect the possibility that the inventor's memory of actual communication is better with respect to more recent technology. Conversely, for the placebos, the spillover index is *lower* the more recent the "cited" invention. Since these represent patents that were not, in fact, cited, there should not have been communication; thus the negative coefficient for the placebos is consistent with the inventors' giving more accurate answers with respect to more recent patents, and more often "mistakenly" indicating communication with respect to older patents.

We included the (log of) total citations received by the "cited" patent to control for the overall "importance" of that patent. It exhibits a positive effect, meaning that more important patents are perceived to have generated greater spillovers, whether because the spillovers are truly greater or simply more likely to be remembered by the respondent. Similarly, cited patents whose inventors reside in the same state as the citing inventor are perceived to have generated greater spillovers. These two variables are *not* significant when looking only at the placebos, further confirming that citations are meaningful, in the sense that the perceived extent of spillovers is correlated with things that *ought* to be correlated with spillovers for the citations, but is uncorrelated with these things for the placebos.

III. Citations and Perceived Importance

In addition to the use of individual citation "links" as possible evidence of knowledge flow, a number of authors have used the total number of citations received by a patent as an indicator of the relative significance of patents. Both our citing and cited inventor surveys asked the inventors to rate the "technological significance" and the "economic importance" of the inventions, and also asked whether the patent had been licensed and whether it had been commercialized. Jaffe, Trajtenberg, and Fogarty (2000) presents results of regressions of the log of total citations received on these indicators of importance. For this purpose, the citing and cited responses were combined into one dataset. In order to control for variations in citation practice by field and changes in propensity to cite and extent of truncation over time, all regressions include technology field and grant year dummy variables. In addition, based on the findings of Jean Lanjouw and Mark Schankerman (1999), we also included the log of the number of claims made by each patent, to allow for the possibility that patents that consist of more claims are more highly cited.

The results provide some evidence that citations are correlated with significance or importance as perceived by the inventors themselves. Each of the indicators is positively correlated with log citations, with the coefficients achieving t-statistics that vary from just below to just above 2, depending on the question. There is no particular indication as to whether citations are more associated with technological versus economic significance. The claims variable is strongly significant, though its elasticity of about .25 suggests strong diminishing returns to increasing the number of claims, as distinct from the constant returns relationship suggested by Lanjouw and Schankerman. If the claims variable is excluded from the regression, the effect of the perceived importance variable increases, suggesting that importance, as perceived by the inventor, reflects both the "size" of the patents as indicated by the number of claims, and the importance or significance of each of the claims.

IV. Conclusion

Many of the important concepts in the economics of technological change are fundamentally unobservable. We routinely rely, therefore, on proxies or indicators for the concepts of interest. Often, our only test of the validity of these measures is the extent to which the proxies are correlated in the way that our theory says that their underlying concepts should be. In this paper, we provide an additional kind of evidence about the unobservable process of knowledge transfer, and the relationship of patent citations to that process.

The results suggest a "half-full cup" with respect to the validity of patent citations as indicators of knowledge spillovers. Taking the responses at face value, the likelihood of knowledge spillover conditional on the observation of a patent citation is significantly greater (in both the statistical and quantitative senses) than the unconditional likelihood. Nonetheless, a large fraction of citations, perhaps something like one half, do not correspond to any apparent spillover. We believe that these results are consistent with the notion of citations as a noisy signal of the presence of spillovers. This implies that aggregate citation flows can be used as proxies for knowledge spillover intensity, for example between categories of organizations or between countries. Further work is needed, however, to refine our understanding of the mechanisms by which these flows move and the relationship of those mechanisms to the citation process.

REFERENCES

Griliches, Zvi. "Issues in Assessing the Contribution of R&D to Productivity Growth." *Bell Journal of Economics*, Spring 1979, *10*(1), pp. 92-116.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics*, August 1993, *108*(3), pp. 577-598.

Jaffe, Adam B, Manuel Trajtenberg, and Michael S. Fogarty. "Report on a Survey of Inventors," forthcoming as an NBER Working Paper, 2000.

Lanjouw, Jean O. and Mark Schankerman. "The Quality of Ideas: Measuring Innovation with Multiple Indicators." NBER Working Paper No. 7345, September 1999.

REFLECTIONS ON PIN FACTORY VISITS

Robert J. Gordon
Northwestern University and NBER

These remarks are complementary to Sue Helper's paper, which focuses mainly on the methodology and benefits of plant visits. In contrast, my comments are on the production process itself. What does one learn from these plant visits, starting from the context of a standard microeconomic production function and the standard macroeconomic literature on productivity change? My reactions are based on seven group visits to six establishments (one of them twice).[1]

At this session Martin Feldstein expressed his surprise that there were so few people on the factory floor. If he had said this and I had not participated in the plant visits, I would have reacted, "why is this surprising, since output in U. S. durable manufacturing increased by 600 percent from 1949 to 1999 while hours of labor input rose only 41 percent?"[2] However, I did participate in some of the plant visits, and this generates another reaction, how *extremely heterogeneous* are the plants in the presence of people on the plant floor. True, at LTV Steel the people visible on the plant floor closely approximated zero, and the only visible employees were those staring at computer monitors in raised computer-control "pulpits" spaced every 300 or so yards apart. And in the Ford body shop the welding was done by tireless but noisy robots, while likewise at Toyota the process of boring cylinders in auto engines was carried out in rows of rectangular tank-like machines which functioned quietly and apparently without human intervention.

But at both Ford and Toyota much of the production process took place on assembly lines that would have made Henry Ford (the elder) feel at home, and they were not unlike auto factories I recall visiting more than three decades earlier.[3] At the Pollak auto parts plant earlier in the 1990s people seated at work stations were omnipresent; in the same rooms only a few years later, a return visit revealed mainly enclosed automated lines for small auto parts, most closely resembling a complex model train setup or a large model car racetrack.

---

[1] The plant visits were to Pollak (an auto parts manufacturer in Boston, MA) twice; LTV Steel in Cleveland; the Ford plant at Lakewood, OH which makes Ford Econoline truck bodies and does final assembly on the twin Mercury Villager and Nissan Quest minivans; Toyota auto assembly in Georgetown, KY; Johnson Control Industries, a manufacturer of auto seats and seat frames adjacent to the Toyota plant; and Chiron, a biotechnology company in Emeryville, CA.

[2] These percentages are from the BLS quarterly database on output and productivity and refer to the change between 1949:Q1 and 1999:Q4.

[3] My previous auto factory visits were at BMC in Oxford, England in 1964 and at GM in Framingham, MA in 1965.

This heterogenous choice between people and automation naturally leads to two questions: 1) which processes tend to use human workers and which don't; and 2) how is the decision made to substitute capital for labor?

The best way to generalize about the first question is to quote our Ford tour guide, "the last thing we'll automate is the marriage of chassis and body. Humans are simply better than robots at finding, manipulating, and fastening the wires and tubes that need to be hooked together at the marriage stage." At both Ford and Toyota, most of the fastening was done by humans on traditional assembly lines. Other processes not involving fastening, like welding and boring engine cylinders, were done by machine. Making finer distinctions, joining large parts involving "heavy lifting" tended to be done by machines, while assembly operations involving relatively small parts still involved humans.

As for the second question about the decision process in substituting capital for labor, we found different approaches. The clearest presentation was at Ford, where each separate operation is an ongoing candidate for capital-labor substitution in a continuing, incremental process. My memory is most vivid of a single man moving large vehicle parts with the aid of an overhead swing. Our guide said "do you see that man? Last year there were two people, now there's one, and next year there will be none, for this process is planned for total automation." Engineers explained to us that every operation is a candidate for capital-labor substitution and that each competed with all others based on a "hurdle rate of return," then (1996) roughly 20 percent.

But not every substitution decision is incremental. Sometimes entirely new processes are substituted for old, most notably at the Pollak plant where hand assembly visible everywhere on our first visit was replaced by the "race-track" automated assembly line of our second visit. From our small sample few generalizations are possible. At Ford the automated processes involved the assembly of large parts, for example body panels, leaving to humans the fastening of nuts and bolts, while Pollak had even succeeded in automating the assembly of switches and small gears involving numerous minute pieces of machined metal.

Perhaps the most surprising and interesting reaction is how different was the experience of visiting two auto assembly plants, Ford in November 1996, and Toyota only 17 months later. Ford was much as I expected: lots of machines, lots of people, and engineers talking about hurdle rates of return in deciding where and when to replace people by machines. Toyota was totally different. Most of what we saw involved assembly lines with workers attaching different types of items to the moving body or chassis, and there was much less (at least that we were shown) involving robots welding together auto bodies.

In fact, at Toyota we never heard the words "profit" or "rate of return," and we rarely even heard the word "productivity." Instead, the day's most frequently used word was "ergonomics." Both management and workers seemed obsessed with taking muscle strain out of routine assembly work and finding ways of replacing strain with comfort. An interesting aspect of ergonomics at Toyota's Georgetown (KY) plant is the evolving redesign of work processes to suit the strength and posture of women, who make up a substantial share of the Toyota U. S. workforce in contrast to the largely all-male Toyota workforce in Japan.

At one position, a worker moved a power fastener vertically at above-head level, a motion sure to create severe upper arm strain, except that the worker had designed a large spring-loaded arm brace that took all the weight off the upper arm. A plaque next to the work stall identified the particular worker who received an award for this invention, which was made in the plant's own machine shop. Similarly, another worker operating at a low level, only two feet off the floor, had designed a complicated seat that moved both from left to right and from forward to back, allowing the seated worker to position herself effortlessly next to the wheel well of the car moving down the assembly line.

At Toyota, ergonomics was a means to an end, in fact two ends: happy employees and high-quality cars. We were indoctrinated into the Toyota philosophy, that high quality products (as rated, for instance, in the *Consumer Reports* reliability surveys) made cars easy to sell, and thus consumer demand would "pull" out of the plant every car that could be made (implicitly at whatever price was necessary to achieve the required profit margin). Engineers reported that they had often driven down a highway and noticed the poor "fit and finish" of trunk lids on competing cars of U. S.-owned brands, clearly visible from their moving cars. They compared their "pull" system with the U. S. "push" system, in which there was a large emphasis on marketing, advertising, and dealer price incentives to "push" out onto the market poor-quality automobiles which could not be sold without this marketing investment.

The Toyota executives whom we met at lunch, none of whom were Japanese, were quite conscious that they were moving more slowly than Ford at automating individual plant processes, and they had an interesting reason to explain their slow pace. The Toyota philosophy of *Kaizan* means continuous, incremental improvements, and there is one problem with excessive reliance on automation: "machines don't make suggestions."

The economics literature on investment has struggled for decades to cope with adjustment costs, so we were a bit surprised to find that adjustment costs were on the front burner at both auto plants, especially at Toyota. The greatest possible disruption at an auto assembly plant occurs with the introduction of a new model. Both at Ford and Toyota, we were told of the elaborate plans made to establish pilot assembly lines up to two years in advance, so that every unexpected wrinkle in the manufacture of a new model could be ironed out long in advance. In fact, both plants planned for unexpected events, and we were told that at Toyota it was optimal to run the plant at 92 rather than 100 percent of capacity, as it would take too much extra labor to deal with the dislocation that would inevitably happen by an attempt to keep every line running at full speed all the time.

How did productivity show up at the auto plants? We did not walk away with any useful measures, except to be told that at Toyota the number of employees required to produce the 1998 Camry was about the same as the 1988 Camry, yet the newer model was superior in almost every dimension, both in size and the complexity of features. Also, productivity measurement is bedeviled by the ebb and flow of outsourcing; in the case of the Georgetown plant, much more of engine assembly and machining was done locally in 1998, whereas in 1988 most was done in Japan. Finally, with the same number of employees the plant was able to produce 45,000 spare parts per month, a new activity not performed previously.

Another set of questions was raised in our visits to two suppliers of auto parts, Pollak (which made small switches for doors and instrument panels) and Johnson Controls (which made auto seats and seat frames). What explains the lack of complete vertical integration in the auto industry? Several answers emerged, including one of the most obvious, the role of unions in raising wages in the plants owned by the auto companies themselves as contrasted with the plants of independently-owned suppliers. Another relatively obvious answer was the ability of independent suppliers to diversify their customer base across several automakers, thus protecting themselves against unforeseen shocks in the form of shifts in demand for particular final products or strikes at a particular customer.

Perhaps more interesting was the feeling we absorbed at both auto parts suppliers, that they had developed an admirable depth of expertise about specialized processes that might be of only marginal interest to executives of major auto companies. The CEO of Pollak who showed us around knew much more about the transition from electric to electronic auto switches than, we would presume, the CEO of GM or Ford, or even a second-level Vice President at GM or Ford. Nevertheless, at Toyota there was more vertical integration than we had expected, having previously visited Pollak (a major supplier to Ford). Toyota made their own engines, bumpers, and instrument panels, although they bought completed seats from the Johnson Controls factory that we visited across the road. Also, the Toyota staff was quite proud of the role of their own machine shop in crafting some of the unique braces and other devices which solved ergonomic problems in particular assembly tasks.

A more general reaction relates to the extensive attention given by economists to research, development, and patents. Clearly, much of the effort directed at productivity improvement that we witnessed was not being achieved within any kind of formal research and development activity, but could be classified under the general rubric of "incremental tinkering." *Kaizan* and R&D may be complementary, but in one sense they are diametrically opposed models of productivity improvement. It is the essence of *kaizan* that it originates in suggestions by line employees and their supervisors, that is, it takes place on the shop floor. In contrast, R&D connotes laboratories that are physically removed from the production process. Doubtless, the relative role of *kaizan* and formal R&D differs greatly across companies within a single industry and particularly across industries.

In fact, our only immersion in "pure" R&D occurred at our visit to Chiron, a biotech company, which seemed little different than visiting the science building in a well-equipped university, with rows of offices of young people staring at PC screens. As I remarked in our visit to this building, with its dramatic atrium and bright earth-tone colors, "this is just like visiting an economics department, but with much better architecture."

COMMENTS ON "JAFFE, TRAJTENBERG, AND FOGARTY" AND "BORENSTEIN AND FARRELL"

Samuel Kortum
Boston University and NBER

General Thoughts

I am a big fan of plant visits. As I see it, the payoff is from what Susan Helper calls "the vividness of field interactions." As economists, we get used to thinking in terms of theoretical abstractions, such as production functions. It is a refreshing change to instead think about something actually being produced. Theory then helps us to focus on what is common among the different production processes we see.

From this perspective, it seems possible to avoid dangerous biases arising from visiting only a small number of plants. The insights one gets often can be verified later (or not) using conventional data. Two examples may help to illustrate this point: Martin Feldstein reports being struck by the small share of production labor in the total costs of the plants we visited. I was surprised to see so many non-production workers engaged in the fixed-cost activity of redesigning products at a plant that produced small electrical components for automobiles. These two observations may have implications for pricing or for the costs of product variety. But, the basic fact — the small cost share of production labor — is easily checked using the Longitudinal Research Database at the Census Bureau.

On a number of visits I noticed that when a plant reported exporting, the exports were a small share of production, sent to just a few foreign locations. We did not hear about a plant supplying a product to the world. This observation suggests that plants are not reaping important scale economies by tapping into foreign markets. The basic fact — that most plants export no more than a modest fraction of their production — turns out to be well documented by Andrew Bernard and J. Bradford Jensen, again using Census data.

As these examples demonstrate, we may often be learning things from plant visits that are buried (perhaps not very deeply) somewhere in standard economic data. The plant visits still have value, however, because they get us to stop and take note. We are bombarded with facts; the problem is knowing which ones to keep in memory. I now turn to the specific papers that have been written for this session.

Jaffe, Trajtenberg, and Fogarty

Patent citations appear to be an incredibly rich source of data for tracking the elusive concept of knowledge spillovers. But, with a few exceptions, citations have been studied in a way that involves no outside validation. Although patent citations sound good on paper, how do we know that they are measuring what we think they are?

Jaffe, Trajtenberg, and Fogarty take a very direct approach to answering this question: ask the people involved. (In this sense, the paper reminds me of the important work of Edwin Mansfield.) To give crispness to what otherwise might have been mushy Likert-scale results, the authors apply a clever idea: Ask inventors not only what they learned from the patents they cited but also what they learned from a placebo patent (in the same technological area) that they did not cite (the inventors were not told about the placebo). The bottom line is a new stylized fact: roughly 50 percent of patent citations indicate some knowledge flow (with little knowledge flow reported for the placebo). I'd say this number is big enough to be good news for those using patent citations to study knowledge spillovers.

Borenstein and Farrell

It is a provocative hypothesis that many corporations could substantially cut costs if they really needed to. Why wait until it is absolutely necessary if your objective is always to make money? To answer such a question, it seems natural to ask the participants.

I applaud Borenstein and Farrell for investigating this hypothesis of X-inefficiency. Such an investigation needs to be carried out by neoclassical skeptics. A fuzzy type will see X-inefficiency everywhere and as a consequence won't succeed in changing the mind of any economist. I think the authors are up to the task. Their skepticism yet openmindedness comes through in the paper.

Unfortunately, the paper does not live up to its promise. Here is what passes for field interaction: "Our interviews ... supported the hypothesis that fat trimming occurs in response to wealth and profit declines." Vivid? I was hoping for vivid accounts of the types of costs that are cut, how the cost cutting initiatives are carried out, and how onerous it is for everyone involved. I would like to hear answers from executives and managers of how much was saved and why they did not cut costs earlier.

Somewhat related work by James Schmitz, on iron-ore extraction, delivers more. He documents big gains in labor productivity within individual Minnesota mines faced with the threat of closure. From visits, interviews, and bits of data he makes the case that these gains resulted from changes in work rules and job classifications that the union agreed to only when the future looked grim.

APPENDIX

NBER WORKING PAPERS IN THE NBER/SLOAN PROJECT ON INDUSTRIAL TECHNOLOGY AND PRODUCTIVITY, 1996 TO 1999*

The Meaning of Patent Citations: Report on the NBER/Case Western Reserve Survey of Patentees
Adam B. Jaffe, Manuel Trajtenberg, and Michael S. Fogarty
NBER Working Paper 7631

Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)
Wesley M. Cohen, Richard R. Nelson, and John P. Walsh
NBER Working Paper 7552

Heart of Darkness: Modeling Public-Private Funding Interactions Inside the R&D Black Box
Paul A. David and Bronwyn H. Hall
NBER Working Paper 7538

Geographic Localization of International Technology Diffusion
Wolfgang Keller
NBER Working Paper 7509

Do Patents Matter? Empirical Evidence after GATT
Jean O. Lanjouw and Iain Cockburn
NBER Working Paper 7495

Endogenous R & D Spillovers and Industrial Research Productivity
James D. Adams
NBER Working Paper 7484

150 Years of Patent Protection
Josh Lerner
NBER Working Paper 7478

Human Resource Management and Performance in the Service Sector: The Case of Bank Branches
Ann P. Bartel
NBER Working Paper 7467

Wage and Productivity Dispersion in U.S. Manufacturing: the Role of Computer Investment
Timothy Dunne, Lucia Foster, John Haltiwanger
NBER Working Paper 7465

*These papers are available at the NBER website: **www.nber.org/sloan**

Do Equity Financing Cycles Matter?
Evidence from Biotechnology Alliances
Josh Lerner and Alexander Tsai
NBER Working Paper 7464

Output-based Pay: Incentives or Sorting?
Edward P. Lazear
NBER Working Paper 7419

Do Industrial Relations Affect Plant Performance?
The Case of Commercial Aircraft Manufacturing
Morris M. Kleiner, Jonathan S. Leonard, and Adam M. Pilarski
NBER Working Paper 7414

Do Stock Price Movements Reveal Profit Dissipation?
An Investigation of the Gold Mining Industry
Joseph Farrell and Severin Borenstein
NBER Working Paper 7075

The Adoption of Offset Presses in the Daily Newspaper Industry in the United States
David Genesove
NBER Working Paper 7076

Organizational Change in French Manufacturing: What Do We Learn from Firm Representatives and from Their Employees?
Nathalie Greenan and Jacques Mairesse
NBER Working Paper 7285

Product Differentiation and the Use of Information Technology:
New Evidence from the Trucking Industry
Atreya Chakraborty and Mark Kazarosian
NBER Working Paper 7222

The Last American Shoe Manufacturers: Changing the
Method of Pay to Survive Foreign Competition
Richard Freeman and Morris Kleiner
NBER Working Paper 6750

Displaced Capital
Valerie Ramey and Matthew Shapiro
NBER Working Paper 6775

The Pay to Performance Incentives of Executive Stock Options
Brian Hall
NBER Working Paper 6674

Effects of Air Quality Regulations on Decisions of Firms in Polluting Industries
Randy Becker and J. Vernon Henderson
NBER Working Paper 6160

Environmental Regulation and Productivity: Evidence From Oil Refineries
Eli Berman and Linda T. Bui
NBER Working Paper 6676

Japanese Research Consortia: A Microeconomic Analysis of Industrial Policy
Lee Branstetter and Mariko Sakakibara
NBER Working Paper 6066

Evidence from Patents and Patent Citations on the Impact of NASA and other Federal Labs on Commercial Innovation
Bruce Banks, Michael Fogarty, and Adam B. Jaffe
NBER Working Paper 6044

Performance Pay and Productivity
Edward Lazear
NBER Working Paper 5672

Computer Use and Productivity Growth in Federal Government Agencies, 1987 to 1992
William Lehr and Frank Lichtenberg
NBER Working Paper 5616

Public-Private Interaction and the Productivity of Pharmaceutical Research
Iain Cockburn and Rebecca Henderson
NBER Working Paper 6018

The Control of Strategic Alliances: An Empirical Analysis of Biotechnology Collaborations
Joshua Lerner and Robert P. Merges
NBER Working Paper 6014

Are CEO's Really Paid Like Bureaucrats?
Brian J. Hall and Jeffrey B. Liebman
NBER Working Paper 6213

Environmental Change and Hedonic Cost Functions for Automobiles
Steven Berry, Samuel Kortum, and Ariel Pakes
NBER Working Paper 5746

Environmental Regulation, Investment Timing, and Technology Choice
Wayne Gray and Ronald J. Shadbegian
NBER Working Paper 6036

Rational Atrophy: The U.S. Steel Industry
Aaron Tornell
NBER Working Paper 6084

The Myth of the Patient Japanese:
Corporate Myopia and Financial Distress in Japan and the U.S.
Brian J. Hall and David E. Weinstein
NBER Working Paper 5818

Skill-Biased Technical Change
Eli Berman, John Bound, and Stephen Machin
NBER Working Paper 6166

Proofs and Prototypes for Sale: The Tale of University Licensing
Richard Jensen and Marie Thursby
NBER Working Paper 6698

The Institutional Context and Manufacturing Performance:
The Case of the U.S. Defense Industrial Network
Cynthia Cook and Maryellen Kelley
NBER Working Paper 6460

The Introduction of Pharmaceutical Product Patents in India:
"Heartless Exploitation of the Poor and Suffering"?
Jean O. Lanjouw
NBER Working Paper 6366

Capturing Technological Opportunity via Japan's Star Scientists
Michael Darby and Lynne Zucker
NBER Working Paper 6360

Consolidation in the Medical Care Market Place:
A Case Study From Massachusetts
Jason Barro and David Cutler
NBER Working Paper 5957

A Clinical Exploration of Value Creation and Destruction in Acquisitions:
Organization Design, Incentives and Internal Capital Markets
Steven Kaplan, Mark Mitchell, and Karen Wruck
NBER Working Paper 5999

Complementarity and Cost Reduction:
Evidence from the Auto Supply Industry
Susan Helper
NBER Working Paper 6033

NBER

National Bureau of Economic Research, Inc.

1050 Massachusetts Avenue, Cambridge, Massachusetts 02138
617-868-3900

www.nber.org