Preliminary
Not for Quotation
Citations incomplete: omissions are inadvertent
Suggestions for further citations welcome

Insider Econometrics:
A Roadmap to Estimating Empirical Models of Organizational Performance

Casey Ichniowski
(Columbia University and NBER)

and

Kathryn Shaw
(Stanford University and NBER)

(November 2006)

"Great advances have been made in theory and in econometric techniques, but these will be wasted unless they are applied to the right data."

-- Zvi Griliches (1994, 2) Presidential Address, American Economic Association

A fundamental objective of the emerging field of organizational economics is an improved understanding of the theory of the firm. As always, theory needs to be grounded in systematic empirical observations and hypothesis testing. Yet it is probably safe to say that Griliches' observation in his presidential address of over a decade ago – that theory and technique often outpace empirical analysis based on the "the right data" – still applies to the economic analysis of organizations.

In this paper, we discuss an approach we have termed "insider econometrics" that grapples with the task of conducting empirical research on organizations based on the right data. The approach is not new. It is one inspired by the body of work of many researchers like Griliches that seeks to combine convincing data with a reasoned use of contemporary econometric methods and, perhaps most importantly, an informed perspective on the issues being studied. "Insiders" in any area are defined as individuals who have "special knowledge or access to confidential information." The insider econometric approach we describe in this paper is one that emphasizes the importance of "special knowledge" and "access to confidential data" for gaining new insight into how firms operate. It relies on going inside organizations not only to gain access to confidential information to construct the "right" data sets but also to gain the special insights that insiders have about why their own organizations are the way they are.

While field research inside businesses has long been used to generate rich case study descriptions of organizational realities, insider econometric studies have an additional objective. The approach also tries to test hypotheses about theoretical regularities about organizations and their workers. The "right data" thus are also data that permit rigorous tests of hypotheses derived from organizational theory. Since data on real organizations are not generated by experiments, the hypothesis testing relies on multivariate econometric methods. In the discussion to follow, we describe this insider econometric approach in more detail and discuss the kinds of topics and level of analysis to which one might apply this approach most productively.

We consider this paper to be a roadmap because it is a guide about this insider methods rather than a review of empirical results. If Griliches is still at least partially correct, then a roadmap may be of some practical benefit. A goal of the paper is to provide a useful guide to the some of the important issues an empirical researcher faces in testing models about how organizations function, rather than a review of the results that confirm or refute theories from some select set of organizational models. While the paper does not include a literature review, we highlight several papers as examples of research that rely on insider data to conduct econometric tests of hypotheses relevant to broad field of organizational economics. As this discussion will imply, systematic empirical research utilizing data on organizations or on workers inside organizations is making important contributions to our understanding about how firms operate.

I. Insider Econometrics: Motivating the Method

This section introduces the basic methodological principles of the insider econometric (IE) approach by considering an illustrative example about real patterns in organizational design that raise the basic questions addressed by IE research. There have now been several large economy-wide studies of human resource management (HRM) practices covering U.S. firms or their workers in the last fifteen years.[1] Several surveys also exist for firms or workers in western European countries as well.[2] A basic empirical pattern emerges from all of these studies, one that is not the explicit focus of studies using these data. HRM practices vary considerably within narrowly defined industry and occupational groups.

An example from research using the most recent of these surveys illustrates this fact about organizational design. Data in the 2002 worker-level General Social Survey[3] was

---

[1] Other nation-wide cross-industry surveys containing data on HRM practices include the 1994 Employee Benefits Research Institute Survey, 1999 Worker Representation and Participation Survey, 2000 National Employment Survey, 2005 National Compensation Survey. Osterman (1994, 2000) reports on data about HRM practices from his survey research of establishments.

[2] Studies by Michie and Sheehan (1999), Leoni, Cristini, Labory and Gaj (2001), and Eriksson (2001) report on performance effects of HRM practices using organizational data from Great Britain, Italy, and Denmarke, respectively.

[3] The GSS survey is administered by the National Opinion Research Center and is now being completed for 2006.

recently analyzed by Kruse, Blasi and Park (2006) to investigate the incidence of profit-sharing, gain-sharing, or other incentive pay plans. Consistent with earlier surveys, they find profit sharing or other related pay-for-performance plans cover roughly one-third of U.S. private sector employees. They also report that 32% of production workers, 33% of administrative support workers, 34% of professional and technical workers, and 47% of sales workers are covered. Industry distributions of these types of compensation plans reveal that they cover 50% of manufacturing workers, 40% of transportation workers, and 29% of finance and insurance industry workers are covered. These surveys about businesses' work practices find the same pattern of large intra-occupation and intra-industry variation for practices in any area of personnel management, and additional research finds this basic pattern is present even within the narrowest of industry classifications and for very precisely defined work practices.

A. Motivating Questions

Research based on these cross-industry surveys of establishments certainly identify that various work practices are observed more commonly in some industries than others, so "industry effects" are important. At the same time, these surveys do not offer rich or comprehensive answers about the causes for the considerable within industry and occupation variation in these practices. IE research takes patterns like these as its starting points and ultimately focuses on two related questions about the adoption of organizational policies and their effects on performance outcomes:

> Why do competing organizations make different choices about key policy decisions that determine an organization's overall design and operations?
>
> Are there differences in performance that are attributable to the different policy decisions that competitors make?

B. Methodological Steps

From this starting point, IE research is characterized by several basic methodological steps that are useful in developing answers for these questions in different settings.

Identify a sample. Samples for analysis in IE research can be samples of organizations or workers, but within the sample, some design or policy feature of theoretical interest varies and performance outcomes that those design features can affect are measurable. Importantly, the sample is drawn narrowly. Observations are homogeneous along many observable dimensions. While this sample design decision will obviously have the beneficial effect of limiting omitted variable concerns as we discuss in more detail below, such samples pose interesting economic puzzles. “Outsiders” would not see obvious reasons for differences in the decisions.

Conduct preliminary field research. Preliminary field research can take many forms from interviews with practitioners to personal inspections of worksites, but the purpose of the field research is to explore the practitioners’ view points on why decision makers in their position have the different policies and practices that they do. By definition, outsiders would find the observed organizational differences puzzling. Here, the “special knowledge” that defines insiders is important. Their knowledge helps to focus subsequent steps of IE research in important ways. It can inform decisions about which theoretical models are most relevant and about the kinds of data one would like to collect. Relevant questions for insiders concern the relationship of the design feature to performance outcomes, other determinants of performance besides the organizational feature of interest, and determinants of the variation in organizational design.

Collect micro-level data. Specific data collection decisions depend on the research context and on access, and of course depend on the extent to which the analyst can tap into the second defining characteristic of insiders – their “access to confidential data.” However, decisions in this step are informed by the preceding research steps. Data collection tries to encompass information about the variation in performance outcomes, the design feature, other determinants of performance outcomes and the design feature. A number of approaches for collecting data will be illustrated in the examples considered in this paper. These examples of IE research will also help illustrate features of data sets that have proven useful in these studies for crafting more convincing answers to the two basic IE questions about the adoption and performance effects of organization-specific policies or decisions, such as an exhaustive “universe” of all observations that fit the criteria for defining the narrowly drawn sample, panel data before and after the observation-specific intervention, and others. But perhaps the most important point to mention in this introductory description is that the data collection efforts are designed to anticipate the basic issues of econometric estimation and interpretation concerning selectivity in the adoption of the design feature and potential effects of unmeasurable variables that exist in any study of an (organization- or worker-specific) treatment on a (performance) outcome in a non-experimental setting.

Conduct econometric analysis. The next section considers this step of testing specific theory-based hypotheses concerning the observed variation in organizational design and performance outcomes, and highlights the kinds of estimation issues that the preceding steps should try to anticipate.

II. Econometric Issues in Estimating Effects of Organizational Treatments

IE research that follows the guidelines presented in section I is ultimately an investigation of an organizational "treatment" on some performance outcome. Because the treatments are, by definition, ones that vary considerably within narrowly defined industries or occupations, the level of analysis is organizations or workers and below the level of the firm which in turn determines the kinds of treatments and performance measures that can be considered.

A. Performance Outcomes and Treatments

While economic performance is ultimately profits, the level of analysis of IE studies will likely preclude a full analysis of the effects of the given treatment on economic profits. Even an analysis of accounting profits cannot be conducted for analyses of individual workers, work groups, or parts of establishments where profits are not measured. Thus, at these very micro-levels of analysis, research will study more detailed measures of worksite performance, such as productivity, product quality, choice of product mix, pricing strategy, or unit sales. Studies of workers might alternatively focus on team output, worker output, wages, incentive payouts or other compensation measure.

The "treatment" refers to a discrete organizational choice variable. It may be the use of piece rate pay or of another human resource management practices, such as the use of teams, the methods for selecting or training workers, the quality of workers, the design of jobs or career ladders, or the degree of information sharing. The treatment can refer to other organizational policies besides work practices: strategies about choice of product mix, the pricing policy, re-organization of worksites under new rules, or an organizational change from deregulation. While these discrete treatments would not encompass adjustments of the quantity of labor or capital, treatments could be the adoption of specific new computer-based technologies for production, a decision to use capital in a just-in-time operations approach, or the use of some new employee recruiting and selection mechanism that affects the quality of the workforce.

B. Modeling Organizational Treatment Effects

Given the goal of modeling the effects of a treatment (X) on performance (Y), one can represent two distinct "switching" regimes for the performance of treated (T) and not treated (N) observations:

(1) $Y_{it}^{T} = \alpha^{T} + X^{T}\beta^{T} + u^{T}$ for $T = 1, E(u^{T}) = 0$

(2) $Y_{it}^{N} = \alpha^{N} + X^{N}\beta^{N} + u^{N}$ for $T = 0, E(u^{N}) = 0$

and a choice of T equation

(3) $I = Z\gamma + v$ where T=1 if $I \geq 0$

Observed performance is

(4) $$\begin{aligned} Y &= Y^{T}T + Y^{N}(1-T) \\ &= \alpha^{T}T + \alpha^{N}(1-T) + X^{T}\beta^{T}T + (X^{N}\beta^{N})(1-T) + Tu^{T} + (1-T)u^{N} \end{aligned}$$

Using (4), we want to identify the effect of the treatment, T. However, the estimated treatment effect is likely to be subject to selection bias as implied by equation (3). In addressing selectivity in the adoption of T, we need to do two things: we need to define what specific kind of "treatment effect" we want to measure, and then we need to provide econometric methods for addressing the selection effects. In the next subsection, we discuss these alternative treatment effects. In the subsection that follows, we discuss selection on the dependent variable and then other issues of measurement error and omitted variable bias.[4]

---

[4] Note also that most researchers assume that $\beta^{T} = \beta^{N}$, so that the effects of the treatment operate only through the intercept, not through interactions with the X variables. Whether this is a reasonable

C. The Average Treatment Effect

The average treatment effect (ATE) is the effect that researchers often identify as their goal in the broad treatment literature: it is the effect from the experiment if the treatment were randomly assigned to firms (or individuals in the firm). In (4), it is:

(5) $$E\left(Y^T - Y^N \middle| X_{it}\right) = \left(\alpha^T - \alpha^N\right) + \left(\beta^T \overline{X}^T - \beta^N \overline{X}^{NIT}\right) + E\left(u^T - u^N\right)T + u^N$$

where $\overline{X}^{NIT}$ is the value of the X's for the organization (or worker) that was not treated but could have been $\left(\overline{X}^{NIT} = E\left(X^N \middle| T = 1\right)\right)$.

The problem is that the treatment is often not assigned randomly, either across organizations or within them. Rewrite the expected residual in (5) as

(6) $$E\left[\left(u^T - u^T\right)T + u^N\right] = E\left[\left(\left(u^T - u^N\right) \middle| T = 1\right) \cdot P_r\left(T = 1\right)\right]$$

Often, this value will be correlated with the $E\left(\alpha^T \middle| T = 1\right) - \left(\alpha^N \middle| T = 0\right)$ which is the treatment gain readily estimated in most data sets.

D. The Treatment of the Treated

While estimating the "average treatment effect" is often the goal in many studies, it is often not a relevant goal for researchers studying organizational performance. For example, consider the treatment of introducing piece rate pay as a replacement for hourly pay. To estimate the average treatment effect, one would take a sample of companies and randomly impose piece rate pay on a subsample and estimate the outcome. Clearly, such an experiment would be hard to imagine. What both researchers and the managers of organizations want to know is the answer to the question: under what production conditions is it optimal to use piece rate pay and given that it is optimally introduced, what are the expected performance gains from it?

Answering these questions requires an estimate of the treatment of the treated. Here, the relevant treatment effect is the expected gains conditional on treatment.

(7) $$E\left(Y^T - Y^N \middle| x, z, T = 1\right) = \left(\alpha^T - \alpha^N \middle| T = 1\right) + \left(\beta^T \overline{x}^T - \beta^N \overline{x}^{NIT}\right) \Big| T = 1 + E\left(\left(u^T - u^N\right) \middle| T = 1\right)$$

---

assumption depends on the particular circumstances of the treatment effect, which insiders are better able to address for each individual study.

which is the gain from choosing the treatment (T=1) given the values of X=x and Z=z. In the example concerning piece rate pay, we need to consider two points. First, among a sample of firms adopting piece rate pay, what was the mean performance gain? This "treatment of the treated" effect can be estimated with first differences (before and after treatment) or instrumental variables. Second, it is important consider that the estimated treatment affect is conditional on the decision to choose piece rates in (3), whether optimal or not. If we can assume that all firms choosing piece rates did so optimally, then we have a clear model of the treatment of the treated for optimal adoption. If firms are not adopting piece rates "optimally," then the interpretation of the estimated effect from (7) is unknown. This leads to estimation of THE alternative mean treatment effect.

E. Methods of Estimating Treatment Effects

As a way of summarizing the intuition about the differences in these kinds of treatment effects, consider Figure 1. There are two sets of firms: those that have the treatment, and those that don't. Figure 1 assumes we are tracking changes in performance over time for the T and N type observations. Both groups experience gradual performance increases over time for reasons such as learning-by-doing or increased human specific capital in production. The treated group experiences a big gain in performance following treatment, of T. We do not observe the performance gain for the non-treated group had it been treated. Figure 1 displays one potential example of this unobserved counterfactual. It shows a hypothetical example that assumes that gains from the treatment would be lower among observations where no treatment is observed.

Figure 1 clearly displays the selection biases. If we estimate the model in OLS, the estimated gain attributed to the treatment will be larger than the effect experienced by those who were treated. In particular, the treated group has a higher level of performance prior to the treatment than the non-treatment group and this difference would be erroneously attributed to treatment. There are a large number of approaches available for addressing the selection bias that arises when aiming to estimate the ATE. This approach is equivalent to conducting an experiment so that there is random assignment and would require that the researcher can ask the firm to implement an organizational change, such as incentive pay or teams, and then estimate the outcome when the change is randomly assigned to individuals in the organization. Experimental data should provide a control group and identical treated group to calculate the unbiased ATE. However, as one can see in Figure 1, even with experimental data for assignment of an organizational treatment would not avoid selection biases, as long as people continue to select into treated and untreated organizations.

Another approach is to take non-experimental data without random assignment of the treatments, and then use instrumental variables or selection correction methods that aim to remove the correlation between the residual and the treatment. These methods would posit equations for the treated and non-treated in Figure 1, and condition the performance outcome on selection into treatment. In one example discussed below, this approach was executed is to use semi-parametric estimation within fixed effects models for changes in performance that difference out the selection bias variable. To estimate the treatment of the treated effect in such examples where panel data contain observations before and after treatment, one can estimate the top performance line in Figure 1, dropping from the sample the non-treated group. Assuming that you care about the treatment of the treated group, this can be a reasonable approach (as long as no other variables are correlated with the treatment).

In essence, we are estimating the expected mean treatment effect for those individuals or firms who are at the margin of adopting or not adopting the treatment (such as incentive pay or some other organizational treatment). This parameter – the mean effect for those at the margin – has been labeled the Mean Treatment Effect (Bjorklund and Moffitt, 1987, Heckman and Vytlacil, 2001, and references therein) or the Local Average Treatment Effect (LATE, Imbens and Angrist, 1994, and other forms of it in Heckman, 1997, Angrist, et. al, 2000). Since this parameter pertains to those at the margin of adoption, we are considering the treatment of the treated. Given the objective of estimating this parameter, there are nevertheless a large number of alternative estimation methods that depend on several practical issues: the nature of the treatment being applied being conducted, the observed variables that are available for identifying the selection equation, what one is willing to assume about the value of the observed variables in predicting selection, and the assumptions about the distributions of the unobserved variables. A consideration of these important issues is available in the references above. But the overall point for this discussion of insider studies of the effects of organizational changes on performance outcomes is that there is tremendous heterogeneity in the nature of the parameter to be estimated and in the actual estimate of the parameter given alternative assumptions.

Two advantages of IE style research for assessing these treatment effects should be noted. First, in some cases, studies with confidential data on detailed worker or organization characteristics will provide better data for constructing the selection equations in the model. Second, even when confidential sources do not provide quantitative data for modeling the selection process, the insiders' perspectives often provide an understanding of the selection process and guide the analyst toward sensitivity analysis that tests whether returns to the treatment vary within various identifiable subsamples of observations. The essential point is that inherent heterogeneity in people and production processes in turn produces heterogeneity in how a given treatment impacts performance outcomes.

III. Empirical Applications

This section considers several examples of IE style research as illustrations of insider knowledge and data about a specific production process can be combined with a thoughtful application of econometric estimation to produce a rich understanding of the ways the various organizational policies and practices affect performance outcomes. We consider two main categories of studies – single firm studies in where worker selectivity must be carefully considered when estimating effects of the organizational policy, and multi-establishment studies where the central questions center on which worksites did and did not adopt the given policy.

A. <u>Single Firm Studies: Estimation of Treatment Effects with Worker Selection</u>

Many of the most compelling papers on organizational performance have been based on data sets from firms when a large firm adopts some new organizational practice. In some cases, it can even be argued that the firm's decision about which workers, work groups, or work sites to the treatment is random. As we review studies that illustrate IE research in such single firm settings, we will find that selection issues are still important even when the policy treatments have features of random assignment.

1. Insider Knowledge of Worker Selectivity after Different Treatments

Lazear (2000) tackles the fundamental economic question of the effects of monetary incentives on employee performance, specifically the effect of piece rate compensation (the treatment) on worker-specific output (the performance outcome). The sample he analyzes is drawn narrowly and pertains to all workers performing one specific job in the same firm, the Safelite Company, before and after the implementation of a piece rate compensation plan. Lazear describes the production process that generates the data for this firm-specific setting. A worker drives a truck to a customer's location and installs a windshield in a customer's car that has a broken windshield. Thus, the production function is worker-specific in this case. Lazear estimates performance equation (1) with number of windshields as the dependent variable, piece rate pay (versus hourly pay) as the treatment, and controls for the workers' tenure. Because each worker forms an individual production unit, he is able to estimate the production function for the firm. The average value of the worker productivity measure is 44% higher after the firm implements piece rate pay.

Lazear then re-estimates the effect of piece rate incentives controlling for fixed (worker) effects. From personnel records, Lazear identifies employees who worked for the firm in the period before the treatment but left before piece rates were implemented and employees who joined the firm after the piece rate plan began. In Figure 2, outcomes for the former group are represented by line BB while outcomes for the latter group are shown in line CC. As is clear from the figure, the fixed effects model that examines outcome changes only within the treatment group produces the gain of AA', while the OLS estimation across groups produces the gain of B'C. The fixed effects estimate of the effects of piece rates on worker productivity falls to 22%. The fall for the fixed effects is due to selection.

Specific features of this study allow us to make several key points about selection. First, from insider knowledge about the treatment, the firm randomly assigned workers by region to the piece rate policy. It would be tempting and natural to assume that the estimated treatment effect is the average treatment effect since workers (who are the production units) are randomly assigned to piece rate pay. Selection occurs through worker turnover which he can identify in the confidential personnel data. Thus, Lazear estimates a treatment of the treated effect. He shows that when production function (1) is estimated with worker fixed effects, the gains from piece rate pay among a fixed set of workers before and after the adoption of piece rate pay is 22%, while the overall change in output across all workers in the firm before and after the policy treatment is much larger, or 44%, a difference which Lazear attributes higher quality workers self selecting into the firm after incentive pay is adopted.

Selection Bias as an Advantage in Organizational Research

Lazear uses endogenous selection as a key advantage for learning more about how incentive pay affects worker performance in (certain) real organizations. We learn that when a firm introduces a new incentive compensation scheme, the firm gains for two reasons. Effort for a given worker rises in response to the incentive pay. But in addition, more productive workers self sort into the firm in response to the incentive pay scheme. In terms of Figure 2, workers of low quality BB leave and workers of high quality CC are hired after piece rates. Often in the treatment literature, analysts try to eliminate the effects of any selection issues when estimating the treatment effect since understanding the nature of the selection will not be the aim of the research. In organizational studies, we want to understand why the selection occurs, and the inside data and analyses often provide that understanding.

Still, Lazear's estimate of a 22% increased worker productivity due to piece rates is a mean treatment effect conditional on selection. The paper estimates mean treatment effects that attempt to measure the return to incentive pay. But in organizations, the effect of one treatment is likely to be conditional on other treatments. In the case of the windshield repair study, the estimated mean effect would change if the firm also introduced a new method of hiring workers that changed the quality of new workers beyond what occurs under self-sorting, or of lowering turnover that retained more of the lower quality workers than was actually the case. Again, insider knowledge about the context is important for understanding how to interpret the empirical patterns in the study. The simplicity of the production function excludes some potential complicating effects – for example, there is no teamwork. Thus, by picking the homogeneous production environment and understanding it well, he is able to focus on two key changes – pay practices and worker selection – and not worry that certain classes of omitted variables such as other practices are responsible for the estimated results.

The above discussion considers one of the two key selection issues that arise in all models attempting to estimate the effects of some policy treatment on organizational performance – the selection of production units receiving the treatment. In the case of Safelite, this is the selection of workers who do and do not work under the new piece rate treatment. The second selection issue concerns the treatment itself – why did Safelite adopt piece rate pay when it did? According to the study, Safelite adopted piece rate pay in part because they purchased new information technology that kept immediate records of individual output. One can consider adoption to be a result of an exogenous technology shock. One can thus argue that piece rates would have had similar effects in earlier time periods but the costs of the necessary performance measurement would have made piece rate pay unprofitable. While the study also reports that incentive compensation remained uncommon for managers and workers in other occupations at Safelite whose output is still difficult to measure, the important point to note here for IE research more generally is that a single treatment applied at one point in time across a broad group of employees does not permit a model of the selection of treatment itself.

2. Insider Knowledge and Data to Assess Causality Behind Treatment Effects

Persuasive studies not only measure differences in outcomes for a given set of subjects under different treatment conditions but they also follow their subjects closely in a way that provides richer insights. One can imagine a medical study comparing those treated by a new drug versus those receiving a placebo. If the drug has desirable effects, one could also imagine the benefits from collecting additional data that tracked the treated subjects that offered insights into how the new drug produced the effect that it did. IE studies can also address the analogous issue for organizations and provide insight into mechanisms through which some treatment had a measurable effect on performance.

Bandiera, Barankay, and Rasul (2005) also use firm-specific data to study the effects of piece rate incentive pay, but in the context of a fruit picking farm. The farm begins with a relative payment scheme in which each worker earns more per unit of fruit if he picks more than his peers, but as overall productivity goes up, the farm lowers the wage per unit of fruit picked. The farm switches to a flat piece rate with a fixed fee per piece picked. Productivity of workers under the plain piece rate plan is over 50% higher compared to the 'relative' payment scheme. The paper is interesting and informative not just because they follow the same workers in the same work environment under two different compensation "treatments," but Bandiera et al. go on to identify why piece rate pay is preferred. By using data on who is friends with whom among the pickers, they show that under relative pay workers hold back on their effort when they are concerned about hurting their friends: if a picker works too hard it lowers the pay of their friends. The key to the study of organizations is not just that piece rate pay is more effective, but that using inside knowledge and data allowed the analysts to document why workers behave differently under different treatments, and thus establish why one treatment leads to a different outcome than another.

3. Insider Knowledge about the Selectivity Behind Which Workers are Treated

In an example like Lazear's study of Safelite, we noted the inherent limitation in the researcher's ability to examine certain aspects of selectivity in the organizational treatment. For Safelite, in regions where the piece rate system was implemented, why was it rolled out when it was? In other studies, however, when there is endogenity of the organizational treatment itself, analysts can develop insider data document important economic aspects about that selectivity.

One example that illustrates this aspect of IE research in a single-firm context is Hamilton, Nickerson, and Owan's study (2003) of teamwork in apparel production. The use of teams in many forms is becoming increasingly common in many firms. Theoretical models of teams posit that there are gains to teamwork when individuals on the team can work in a complementary fashion – when the multiplicative interaction between team members raises output more than the sum of the individual parts. At the same time, productivity can decline in moving from individual work to team production when free rider issues are important. Since some employees in a given occupation or industry work in teams while others do not, the question of the effects of teams on performance is one that IE style research could usefully address. Hamilton, Nickerson, and Owan address this increasingly important question by analyzing data on worker and work group productivity in apparel making, and document an increase in productivity after workers shift to teamwork. This finding suggests that the gains from worker skill complementarities or peer monitoring seem to offset the potential losses of output from a free-rider effect in these teams.

Their analysis goes further and uses rich data on the workers to provide a deeper understanding of the team selection process. In particular, their study provides an example of the endogenous selection of production units that undertake the teamwork "treatment" – the decision by workers to work as a team is endogenous. Since the gains to teamwork are hypothesized to be a result of skill heterogeneity, selection of team members is critical. The composition of the team should affect the estimated returns to teamwork. However, there could also be adverse selection in this case, where the best sewers should avoid teams and stick with individual piece rate pay. How do these competing forces play themselves out?

Using insiders' knowledge about the organization and detailed confidential data on the history of workers' performance before and after teams, Hamilton et al are able to test the key selection hypotheses by looking at the quality of the individual performers on the teams and the pre-team performance of workers who form the teams. Their analysis documents different ways that worker selection on the teams is important to the estimated differences in productivity before and after team formation. First, they document that better sewers are more likely to form teams earlier. Thus their results follow the general pattern laid out in basic Figure 1 illustration (which when adapted to the specifics of this study, can be represented by Figure 3). Those who choose teams first have higher gains than those who choose teams later because the early choosers are more productive in terms of both their level of output and in the changes in output after moving to teamwork. They also show that other selection issues about where the organizational practice is applied are also important when considering performance differences before and after team performance. Team productivity grows more when the team is comprised by individuals who exhibited greater heterogeneity in individual performance.

In all of the subsections above, data from a single firm are used to understand the size of the productivity gains of a new organizational treatment in a way that accounts carefully for issues related to the selection of workers.[5] Estimated productivity differences are affected by workers' individual quality (Lazear, 2000), the workers' quality prior to forming teams (Hamilton et al.), and detailed characteristics related to social relations among workers (Bandiera et al.). There is no one 'treatment effect.' The magnitude of the estimated effect of the treatment varies as a function of the selection of workers. Organizational research does not just try to difference out the effects of such selectivity. Instead, organizational research must consider differences in the estimated size of organizational treatment effects under different econometric models to gain insights about the nature of the selectivity and about how these treatments affect the performance of firms.

B. Effects of Treatments Applied in Different Organizations

The above studies focus on the estimated effects of a change in an organizational practice within one firm. These within-firm models seek to understand the impact of the organizational change by comparing workers who are and are not subject to the organizational change in ways that carefully consider the endogeneity of which workers are covered by the given change. The estimation in these studies ultimately rely on a difference-in-difference model, for workers who do and don't adopt.

Of course, by looking only within one firm, the models cannot make the availability of the practice itself endogenous. In Hamilton et al (2003), insider data help identify which workers for teams once the firm decided to adopt their teamwork policy. The single firm study however cannot identify why this particular apparel making firm adopted the overall firm-wide policy about teamwork while another apparel maker did not,[6] or how much some other firms might gain. In terms of the econometric estimation, using within-firm data in which there is only one firm-wide decision to adopt the overall

---

[5] In other studies in which the person is the production unit, worker selection affects the impact of the treatment. For example, Landers, Rebitzer, and Taylor (1996) model the hours of lawyers, showing that large firms using tournaments produce longer work hours. Here selection is also important. Given shared compensation under the partnership structure of law firms, adverse self selection of low quality lawyers into high hours firms is a concern. These firms must screen workers by using their record of billable hours in the promotions decision.

[6] Cross-establishment research on HRM practices including teamwork among apparel makers does offer information on this kind of selectivity. See Dunlop and Weil (1996) and Berg, Appelbaum, Bailey, and Kalleberg (1996)

firm policy, the analyst cannot estimate difference-in-differences with the treatment decision as the unit of analysis rather than the worker.

For example, to develop even richer estimates of the effects of piece rate pay on performance than are developed in the studies by Lazear (2000) or Bandiera et al. (2005), we would like to have observations on the introduction in piece rate pay in organizations that are at the margin in introducing it, find a matched sample of firms that are also at the margin but did not introduce it (for exogenously determined reasons), and then estimate difference-in-difference models to compare the effects of piece rate across the two groups. No study using existing firm-level data can achieve this. However, this objective is an important benchmark to keep in mind when considering existing studies that do examine the effects of different organizational treatments across organizations. This study considers estimation of the effects of organizational treatments on performance outcomes using insider data from many organizations.

1.Insider Data on Practices Across Organizations and Treatment of the Treated Estimates

Consider again the effects of teams on productivity analyzed in Hamilton et al. (2003), but now the unit of observation is the plant, not the individual. In terms of the basic comparisons illustrated in Figure 1, there are two sets of plants: those that have adopted teams, and those that have not. As before, the performance of both groups will change over time, with Figure 1 assuming gradual performance increases for both groups. The team-treated group experiences a big gain in performance following treatment, of P. We do not observe the performance gain for the non-treated group, but the figure assumes that the plants that did not adopt teams would experience smaller gains from adopting teams since they did not adopt them.

Figure 1 suggests the possibility of differences-in-differences estimation with fixed effects models. Clearly full sample OLS estimates without controls for plant fixed effects are biased. If we estimate such an OLS model, the gains to teams will appear bigger, because the treated group has a higher level of performance that would be erroneously attributed to treatment. However, difference-in-differences assumes that the treatment sample is the same as the untreated conditional on observed X's. To be specific, the difference-in-differences estimation is appropriate for Figure 1 if we assume that the estimated effect of time can be estimated parametrically so that the non-treated group has a lower growth of productivity over time ensure than the treated group. That is, as in equation (4) above, the βX varies across the treated and untreated subsamples, and the parametric estimation of βX captures these differences. If instead, the slope of the non-treated is lower than the slope of the treated (as in Figure 1), and this is not measured well with the observed X's, then the gains to teams with be overestimated because the higher learning curve of the group that selects teams will be attributed to the team effect, not to differences in selection. If there are sufficient data, a reasonable choice for estimating the treatment of the treated without this bias is to limit the sample to only those who select teams.[7]

Using longitudinal data on the sample of U.S. steel minimills, Boning, Ichniowski, and Shaw (forthcoming) estimate the effects of teams on measures of productivity in this industry according to this procedure that focuses on productivity changes within the group that adopted teams. Comparing estimates of the effects of teams in models that do and do not restrict the sample to only those mills that introduced teams, they find that estimated treatment effects are very similar. Apparently, in this case, there is little difference in the growth rates of productivity for the non-adopters and the adopters prior to adoption.

Insider knowledge here is used to rule out certain explanations for differences in adoption of practices like teams across plants. The "special knowledge" that insiders offer can also be used to identify factors that do help explain why some firms were 'treated' and some were not. In the case of minimills, site visits and interviews revealed that teamwork and problem solving was more important in mills that made more complex products. Without site visits, one might simply have assumed that homogeneous production facilities like these, all within the exact same narrow standard industrial classification, is also homogenous. For four different classes of minimill products observed in these plants, from the least complex product type to the most, the frequency of teams was: 23% of all plants in product class one (plants making the simplest minimill

[7] One of the first "organizational treatments" to receive considerable attention in empirical econometric studies was union status, and this literature offers examples of treatment of the treated estimation within samples of plants that became unionized. Clark's analysis (1980) of cement plants is an early exemplar of an IE approach that generates compelling treatment of the treated estimates of the positive effects that unions can have on productivity. This research also reflects the use of interview evidence with cement plant managers to uncover mechanisms that explain the estimated productivity effects.

steel products); 33% of plants making product class 2; 67% of plants making product class 3; and 100% of the plants making the most intricate type of steel products. These patterns then lead to additional theoretical analyses to explain why problem solving teams are more commonly adopted when organizations make more complex products and empirical analyses that examine changes in the productivity metrics within plants making different kinds of products. As the adoption patterns suggest, the introduction of teams in minimills making more (less) complex products leads to larger (smaller) changes in mill productivity.

The example here offers an interesting point about the value of IE research for addressing the two basic questions about the adoption and subsequent performance effects of organizational practices. Broad-based surveys of establishments in many industries may be able to identify that an establishment in one industry like minimills is more or less likely to have teamwork policies than an establishment in some other industry. Such research may also be able to identify certain establishment-level characteristics correlated with team policies. However, research following an IE approach was able to identify a very powerful predictor of the adoption of teamwork policies based on field work that also revealed that problem solving activity by teams was particularly useful for mills making more intricate steel products.

2. Insider Data on Practices Across Organizations with Exogenously Determined Treatments

While policies and practices that vary across plants in these intra-industry samples are not randomly assigned by plant managers, some IE research that considers the effects organization-specific practices on performance concludes that estimated changes in productivity observed among policy adopters might well be a reasonable estimate of the change in productivity that non-adopters would experience. Ichniowski, Shaw, and Prennushi (1997) study the effects of different systems of HRM polices on productivity in production lines in integrated steel mills. This analysis also includes estimates from econometric models with controls for plant-specific fixed effects and finds significant productivity benefits associated with adopting new HRM systems.

Their conclusion for differential adoption of these productive HRM systems across similar plants is that certain plants have higher transition costs of adopting these practices. In particular, the most productive HRM systems were more likely to be observed in "greenfield" lines starting up production and in older "reconstituted" lines that started up operations by new owners after being closed temporarily than they were in older lines that operated continually since inception (Ichniowski and Shaw, 1999). While the desired counterfactual on the productivity effects among non-adopters can never be observed, the authors conclude in this case that the most compelling conclusion for this setting is that adoption is governed not by the expected productivity gains from the new HRM practices, but by differences in transition costs across sites. It was less costly to implement new work systems in an old plant that was closed and being opened by new owners than it would be to implement them in an old, but continuously operating, worksite.

Perhaps the most important point to make for this analysis of IE methods is that fine-grained distinctions can be what differentiates adopters of organizational policies from very comparable non-adopters. Even survey questions about age of plants would not capture the distinctions in this industrial setting. Furthermore, the IE research here directs theoretical analysts to consider transition costs related to social relations and attitudes in the workplace rather than the age of capital equipment as the more important determinant of HRM adoption. Finally, the example here illustrates the importance of remembering that performance metrics in micro-level studies cannot measure overall profitability to the firm. In this case, the study argues that HRM adoption would raise steel-making productivity among non-adopters, but transition costs would make this decision an unprofitable one.

Mas (2005) provides an illustration of IE style research from a very different setting than manufacturing in which exogenous determination of the treatment may be important in interpreting observed empirical patterns. Mas looks inside public sector organizations and examines how well New Jersey police perform in solving crimes. The outcomes in this study are police pay and the treatment is the use of interest arbitration to determine the police union contracts. Mas finds that when there is a large loss for the union workers in arbitration – measured by larger differences between the union demand and the arbitrator award – workers are subsequently less productive. In interpreting these patterns, Mas argues that, while arbitration may be endogenous to various firm-specific factors, the wage outcome of arbitration can be considered exogenously determined and this is what drives his organizational change results. Ultimately, he concludes that it is not the level of the wage that affects performance it is the loss relative to expectations. Like other IE research, the study pertains to the very specific group of unionized police and those subject to state law permitting contract arbitration. But the patterns observed for this setting suggest a specific conclusion that workers are motivated to perform in part by the difference between wage outcomes and their wage expectations that may apply to other settings.

C. Omitted Variable Bias and IE Estimation

In any econometric study of the effect of some treatment condition on a performance outcome based on non-experimental data, one can always imagine reasons why estimated effects of the treatment could be attributable to some factor correlated with both the treatment and the performance outcome that is omitted from the analysis. The selection model presented in section II and represented in Figure 1 encapsulates this concern. In particular, one can still argue that various estimates in studies reviewed above (e.g., fixed effects estimates of the effect of teams on productivity in before/after comparisons in a sample of team adopters) are still overestimates of the treatment-of-the-treated effect. One need only argue that some other productivity enhancing factor was adopted around the same time as teams in these plants. In terms of the presentation of section II, the estimated treatment effect may still be conditional on the omitted selection variable. Certain advantages of IE research in considering omitted variable issues are clear. To the extent that the IE researcher has conducted extensive interviews or field research, he or she brings a richer understanding of the determinants of the organizational treatment and of the performance outcome. Furthermore, effects of omitted variables would have to ones that would be plausible for the narrowly drawn samples that characterize IE studies.

1. Measuring Unmeasurables

IE research also offers examples of a direct approach for considering whether estimated effects of organizational treatments might be due to omitted (time-varying) variables that is typically not available without an IE approach. In the studies considered in section III.2 that analyze data from multiple organizations, one could reasonably be concerned that more productive organizational practices simply capture effects really attributable to better managers. Better managers adopt different practices than do other managers, such as teamwork policies, but superior performance in the organizations is due to the managers and not the practices. In section II presentation, management quality enters the adoption equation (3) as a Z variable, and also is an element of the X vector in the performance equations.

Insider access to organizational information data can allow researchers ways of addressing such potential sources of bias in estimates. For example, in the case of HR practices adopted in integrated steel mills, Ichniowski, Shaw, and Prennushi (1997) create explicit measures for the time periods when different managers were in charge of the steel making operations. The combination of narrowly drawn samples that is sometimes coupled with unique opportunities to identify and incorporate measures of factors that would remain unmeasured without the special knowledge that an IE approach can often allow IE research reasonably convincing ways for addressing omitted variable concerns.

2. Was There More than One Organizational Treament?

Milgrom and Roberts (1990) offer a simple description of operations of contemporary manufacturing operations as an illustration of a potentially important point that would affect IE research seeking to estimate the effect of an organizational treatment on performance outcomes – many organizational practices may adopted as systems of practices because of complementarities among the practices. For example, if a plant decides that introducing teams is optimal, it may also institute additional training, more careful selection of workers for attitudes toward teamwork, information sharing among employees and other practices that support effective problem solving by teams. If these organizational practices are complements, then all should enter the production function multiplicatively. That is, if you are looking at the effects of teamwork on productivity, but omit a complementary work practice, the return to teamwork as a lone policy treatment will be biased upward.

The possibility of complementarities among organizational practices counsels the analyst to investigate the nature of the organizational treatment carefully. This task is perhaps more manageable within narrowly drawn IE-style samples since specific production settings like those considered in this study, like windshield replacement, apparel making teams, fruit picking, minimill steel lines, and others, may have a more limited set of relevant organizational policies. Existing empirical IE research suggests that complementarity among organizational practices is an important concern with regard to complementarities among HRM practices (Ichniowski, Shaw and Prennushi, 1997); among new HRM practices and just-in-time production practices (MacDuffie, 1995); and among new information technologies and new HRM practices (Breshnahan, Brynjolfson, and Hitt, 2002).

IV. Obtaining Insider Data

Several observations can also be made based on this review about alternative methods researchers have used to access insider data.

A. <u>Insider Data from Within Firms</u>

First, the seemingly most common kind of insider data used in IE studies is data from a single firm. Within the firm, IE research focuses on distinct production processes and within any firm there can be many different output generating "production functions." The production function could be individual workers, work groups within an establishment, different establishments within a division, and so on. With increasingly sophisticated information collecting and storage technologies now common place in many businesses and firms, and the greater likelihood of cultivating insider access through one or two key officers of a firm, we see this as a particularly promising area for growth of new IE research.

Second, the number of industry-specific studies about performance differences across different plants or observations that all share a very specific production process has also grown in recent years. This approach relies on data across firms. While this approach obviously requires significant time and effort, it can result in exactly the kind of rich understanding of an entire industry context that IE research strives for. Methods for managing this approach to data collection and its costs are therefore worth considering.

One development here is the commitment to network building in well-established industry councils supported by the Alfred P. Sloan Foundation. The number of industry-specific study councils supported by the Sloan Foundation has now grown to twenty-six, some of which have been in operation for many years. From a practical standpoint, this approach to research can benefit from the support of a small board of industry advisers and the existing councils could be beneficial in facilitating data collection.

B. Insider Data Available Through Surveys

Data collection for industry-specific research across multiple organizations and firms that encompasses many production units does not have be gathered from site visits to each production unit. Several alternatives are possible for managing the costs of IE data collection. In our own work, we conducted five plant tours in one industry with the aim of developing an industry-specific survey to study the effects of new information technologies and new HRM practices on productivity and worker skill requirements (Bartel, Ichniowski, and Shaw, 2006). The field visits and interviews did not generate the confidential IE data, but it did provide the "special insider knowledge" that we used to develop a survey in the language specific to the technologies and operations of that industry. Contacts for the survey sample were obtained through public data sources on business establishments in the U.S., and the survey was conducted by a telephone by a survey research firm. This approach again involves data collected personally by the analyst, but in a way that economizes on the costs of data collection.

Other studies that fit the section II criteria for IE style research have made effective use of existing data on workers or establishments within certain industries. The Census Centers in the U.S. provide access to surveys with detailed information on specific industries. Perhaps the best known exemplars here are studies of the effects of the adoption of new information technologies on trucker productivity (Hubbard, 2003; Baker and Hubbard, 2003; 2004) and occupation-specific research on lawyers (Garicano and Hubbard, 2003).

C. Insider Analysis in Regulated Industries

A number of studies also make use of publicly available data on specific industries, often because regulated industries must file information with government agencies. Examples of insider work that has been done on public data include the following. [*References and short summaries to be added*.]

- Electricity generation – Bushnell and Wolfram (2006) on productivity as a function of the variance of individual performance.
- Education
- Health care
- Retail hotels in Texas – Kalinins (2006) on social networks and productivity using public hotel data; also reference Ingram and Roberts (2000) on hotel performance and social networks
- Police – Mas (2006) matches crime data to arbitration data.

D. Using Existing General Survey Data to Create Insider Studies

Finally, agencies of the U.S. government, as well as government agencies of many European countries, have begun to make available what have been labeled employer-employee matched data sets, often on a limited basis to researchers inside the government. These data sets contain administrative data from firms and on every worker in the firm, including data on wages, demographic features of the employees, and at times their occupations and promotions. In U.S data, workers are tracked over time between firms. These data sets are very broad in scope and may not contain the kind of rich institutional detail on many organizational characteristics an IE researcher would desire. However, in some cases, these data have been linked to specific production data of firms or to surveys of organizational practices of firms. Because these data sets are so large (often covering all individuals in the country for large sectors), subsets of the data can be used for industry analysis. For example, in the U.S., data have been collected on the performance and internal structure of firms for trucking, semiconductors, and software industries. (Abowd, Haltiwanger, Lane, 2006; Andersson, Freedman, Haltiwanger, Lane, and Shaw, 2006)

V. Conclusion

We introduced our summary of Insider Econometric studies of the effects of organization-specific policy treatments on performance outcomes with the Griliches' observation about the need for empirical research using the "right data" to catch up with advances in theory and technique. The discussion here does not offer a boiler plate check list of the characteristics that several informative IE style studies all share in common. But the discussion does offer a road map with a destination – reasoned and rigorous tests of hypotheses about the effects of organizational policies and practices on performance outcomes. Since true random assignment of policies to organizations and of workers to organizations does not occur in real organizations, and since important economic insights about organizations are gained by evaluating the precise ways in which selectivity in these treatments does take place, it is important for the analyst to anticipate these concerns in developing and collecting the quantitative data used in IE research.

IE research analyzing narrowly drawn samples by its very nature cannot produce findings that can be broadly generalized across many different industrial settings with their own distinctive production processes. But, several important advantages of IE style research for assessing treatment effects should be noted. First, studies with confidential data on detailed worker or organization characteristics will provide better data for constructing the selection equations in the model. Second, even when confidential sources do not provide quantitative data for modeling the selection process, the insiders' perspectives often provide an understanding of the selection process and guide the analyst toward sensitivity analysis that tests whether returns to the treatment vary within various identifiable subsamples of observations. Third, models of the performance outcomes can be more thoroughly understood and specified. The results of the research can prove to be persuasive and offer insights about the processes that link organizational practices to performance.

References

Andersson, Fredrik, Brown, Clair, Campbell, Benjamin, Chiang, Hyowook, and Park, Yooki "The Effect of HRM Practices and R&D Investment on Worker Productivity," *Working paper*, 2006.

Andrews, Donald W. K. and Marcia Schagfans (1996) "Semiparametric Estimation of the Intercept of a Sample Selection Model," *Review of Economic Studies,* vol. 65, pp 497-517.

Appelbaum, Eileen, Thomas Bailey, Peter Berg, and Arne Kalleberg (2000) *Manufacturing Advantage: Why High-Performance Work Systems Pay Off.* Ithaca, NY: Cornell University Press.

Appelbaum, Eileen, and Rosemary Batt (1994) *The New American Workplace: Transforming Work Systems in the United States*, Ithaca, NY: Cornell University Press.

Arthur, Jeffrey (1992) "The Link Between Business Strategy and Industrial Relations Systems in American Steel Minimills," *Industrial and Labor Relations Review*, vol. 45, pp. 488-506.

Asch, Beth J. (1989) "Do Incentives Matter—The Case of Navy Recruiters." *Industrial and Labor Relations Review*, vol. 43 (Supplement 1), pp. s89-s106.

Ashenfelter, Orley (1978) "Estimating the Effect of Training Programs on Earnings," *Review of Economics and Statistics*, vol. 67, pp. 648-660.

Ashenfelter, Orley and Hosken, Daniel "The Effect of Mergers on Consumer Prices: Evidence from Five Selected Case Studies," *Working paper*, 2006.

Athey, Susan and Scott Stern (1998) "An Empirical Framework for Testing Theories about Complementarity in Organizational Design," *NBER working paper* #6600.

Autor, David, Frank Levy, and Richard Murname (March 2001) "The Skill Content of Recent Technological Change: An Empirical Investigation."

Autor, David, Frank Levy, and Richard Murname (March 2001) "Upstairs Downstairs: Computers and Skills in Two Floors of a Large Bank."

Autor, David H. (2001) "Why do Temporary Help Firms Provide Free General Skills Training?" *Quarterly Journal of Economics*, vol. 116, pp. 1409-1448.

Baker, George P. and Hubbard, Thomas N (2003) "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information," *American Economic Review*, vol. 93, pp. 551-572.

Bandiera, Oriana, Barankay, Iwan, and Rasul, Imran "Incentives and Favoritism in the Workplace: Evidence from Personnel Data," *Working paper*, 2005.

Bandiera, Oriana, Barankay, Iwan, and Rasul, Imran (2005) "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics*, vol. 120, pp. 917-692.

Bandiera, Oriana, Barankay, Iwan, and Rasul, Imran "Social Connections and Incentives in the Workplace: Evidence from Personnel Data," *Working Paper*, 2006.

Bandiera, Oriana, Barankay, Iwan, and Rasul, Imran (2006) "The Evolution of Cooperative Norms: Evidence from a Natural Field Experiment," *Advances in Economic Analysis & Policy*, Vol. 6: No. 2.

Bartel, Ann (2003) "Human Resource Management and Organizational Performance: Evidence from Retail Banking" *Industrial and Labor Relations Review*, vol. 57, no. 2.

Batt, Rosemary (July 1999) "Work Organization, Technology, and Performance in Customer Services and Sales," *Industrial and Labor Relations Review*, vol. 52, pp. 539-64.

Berg, Peter, Eileen Appelbaum, Thomas Bailey, and Arne L. Kalleberg (2000) "The Performance Effects of Modular Production" in *The American Workplace: Skills, Compensation and Employee Involvement*, eds. Casey Ichniowski, David Levine, Craig Olson, and George Strauss (Cambridge University Press).

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, (2004) "How Much Should We Trust Difference-in-Differences Estimates?" *Quarterly Journal of Economics*, *119*(1), pp. 249-75

Bhattacherjee, Debashish (2005) "The Effects of Group Incentives in an Indian Firm: Evidence from Payroll Data," *LABOUR*, vol. 19, pp. 147-173.

Black, Sandra and Lisa Lynch (2004) "What's Driving the New Economy?: The Benefits of Workplace Innovation", *The Economic Journal,* vol. 114, (February).

Black Sandra, Lisa Lynch and Anya Krivelyova (2004) "How Workers Fare When Employers Innovate", *Industrial Relations,* vol. 43(1), (January)

Black, Sandra and Lisa Lynch (2001) "How to Compete: The Impact of Workplace Practices and Information Technology on Productivity", *Review of Economics and Statistics,* (August).

Boning, Brent, Casey Ichniowski, and Kathryn Shaw (forthcoming) "Opportunity Counts: Teams and the Effectiveness of Production Incentives," *Journal of Labor Economics*.

Breshnahan, Timothy, Erik Brynjolfsson and Loren Hitt (2001) "Information Technology, Work Organization and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics*.

Cappelli, Peter, and David Neumark (2001) **"**Do High-Performance Work Practices Improve Establishment-Level Outcomes?" *Industrial and Labor Relations Review*, Vol. 54, No. 4 (July) pp. 737-775

Chevalier, Judith A. and Morton, Fiona Scott "State Casket Sales Restrictions: A pointless Undertaking?" W*orking paper*, 2006.

Clark, Kim B. (1980) "The Impact of Unionization on Productivity: A Case Study," *Industrial and Labor Relations Review*, vol. 33, pp. 451-469.

Clark, Kim B. (1980) "Unionization and Productivity: Micro-Econometric Evidence," *Quarterly Journal of Economics*, vol. 95, pp. 613-639.

Clark, Kim B. (1984) “Unionization and Firm Performance: The Impact on Profits, Growth, and Productivity,” *The American Economic Review*, Vol. 74, pp. 893-919.

Drago, Robert and Garvey, Gerald T. (1998) “Incentives for Helping on the Job: Theory and Evidence”, *Journal of Labor Economics*, vol. 16, pp. 1-24.

Dranove, David, Kessler, Daniel, McClellan, Mark, and Satterthwaite, Mark (2003) “Is More Information Better? The Effects of “Report Cards” on Health Care Providers,” *Journal of Political Economy*, vol. 111, pp. 555-588.

Dunlop, John T. and David Weil (2000) “Diffusion and Performance of Modular Production in the U.S. Apparel Industry,” in *The American Workplace: Skills, Compensation and Employee Involvement*, eds. Casey Ichniowski, David Levine, Craig Olson, and George Strauss (Cambridge University Press).

Encinosa , William E. III; Gaynor, Martin and Rebitzer, James B. (1997) “The Sociology of Groups and the Economics of Incentives: Theory and Evidence on Compensation Systems,” Agency for Health Care Policy and Research.

Fernie, Sue and Metcalf, David (1999) “It’s Not What You Pay It’s the Way That You Pay It and That’s What Gets Results: Jockeys’ Pay and Performance.” *Review of Labour Economics and Industrial Relations*, vol. 13, pp. 385-412.

Freeman, Richard B. and Kleiner, Morris M. (2005) “The Last American Shoe Manufacturers: Decreasing Productivity and Increasing Profits in the Shift from Piece Rates to Continuous Flow Production,” *Industrial Relations*, vol. 44, pp. 307-330.

Galizzi, Monica, Miesmaa, Petra, and Slatin, Craig “Occupational Injuries, Workers’ Reporting and Firms Policies in the Health Care Industry: The Challenges and Rewards of Combining Qualitative and Quantitative Research Methodologies,” *Working paper*, 2006.

Garicano, Luis and Heaton, Paul “Computing Crime: Information Technology, Police Effectiveness, and the Organization of Policing,” *Working paper*, 2006.

Garicano, Luis and Hubbard, Thomas N. “Learning about the Nature of Production from Equilibrium Assignment Patterns,” *Working paper*, 2005.

Garicano, Luis and Hubbard, Thomas N. “Managerial Leverage is Limited by the Extent of the Market: Hierarchies, Specialization, and the Utilization of Lawyers’ Human Capital,” *Working paper*, 2005.

Gaynor, Martin, Rebitzer, James B., Taylor, Lowell J. (2001) “Incentives in HMOs” *NBER Working Paper* #8522.

Griliches, Zvi. (1994) “Productivity, R&D, and the Data Constraint” *American Economic Review,* vol. 84, no. 1, pp. 1-23.

Guertzgen, Nicole “Rent-Sharing and Collective Bargaining Coverage – Evidence from Linked Employer-Employee Data,” *Working paper*, 2005.

Hamilton, Barton, Nickerson, Jack and Owan,, Hideo (2003) "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation," *Journal of Political Economy*, vol. 111, pp. 465-497.

Heckman, James J. (1990) "Varieties of Selection Bias," *American Economic Review*, vol. 80, pp. 313-8.

Heckman, James J., Hidehiko Ichimura, Jeffrey Smith, and Petra Todd (1998) "Characterizing Selection Bias Using Experimental Data," *Econometrica*, vol. 66, pp. 1017-1098.

Heckman, James J., Robert Lalonde, and J. Smith (1999) "The Economics and Econometrics of Active Labor Market Programs," in *Handbook of Labor Economics*, vol. 3, eds. Orley Ashenfelter and David Card (Elsevier Science: Amsterdam).

Holmstrom, Bengt (1982) "Moral Hazard in Teams," *Bell Journal of Economics*, vol. 13, pp. 324-340.

Holmstrom, Bengt and Paul Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics and Organization*, vol. 7, pp. 24-52.

Holmstrom, Bengt and Paul Milgrom (1993) "The Firm as an Incentive System," *American Economic Review*, vol. 84, pp. 972-991.

Hubbard, Thomas N. (2003) "Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking," *American Economic Review*, vol. 93, pp. 1328-1353.

Hunter, Larry and John Lafkas (1998) "Information Technology, Work Practices, and Wages," *Proceedings of the 50th Annual Meeting of the Industrial Relations Research Association.*

Huselid, Mark and Brian Becker (1996) "High Performance Work Systems and Firm Performance: Cross-Sectional Versus Panel Estimates," *Industrial Relations*, pp. 635-672.

Ichniowski, Casey, David I. Levine, Craig Olson and George Strauss (Fall 1996) "What Works at Work: Overview and Assessment" *Industrial and Labor Relations Review*, vol. 35, pp.299-334.

Ichniowski, Casey, Kathryn Shaw and Giovanna Prennushi (1997) "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," *American Economic Review*, vol. 87, pp. 291-313.

Ingram, Paul and Peter W. Roberts (2000) "Friendships among Competitors in the Sydney Hotel Industry." 106: 387-423, *American Journal of Sociology*.

Jones, Derek C. and Kato, Takao "The Impact of Teams on Output, Quality and Downtime: An Empirical Analysis Using Individual Panel Data," *Working paper*, 2005.

Jones, Derek C., Kalmi, Panu, and Kauhanen, Antti "Teams, Performance-Related Pay, Profit Sharing and Productive Efficiency: Evidence from a Food-Processing Plant," *Working paper*, 2006.

Kalnins, Arturs and Chung, Wilbur (2006) "Social Capital, Geography, and Survival: Gujarati Immigrant Entrepreneurs in the U.S. Lodging Industry," *Management Science*, vol. 52, pp. 233-247.

Knoeber, Charles R. and Walter N. Thurman (1994) "Testing the Theory of Tournaments: An Empirical Analysis of Broiler Production," *Journal of Labor Economics*, vol. 12, pp. 155-79.

Kreps, David (1997) "Intrinsic Motivation and Extrinsic Incentives," *American Economic Review*, vol. 87, pp. 359-365.

Krueger, Alan B. (2003) "Strikes, Scabs and Tread Separations: Labor Strife and the Production of Defective Bridgeston/Fireston Tires," *NBER Working paper* #9524.

Kruse, Douglas (1993) *Profit Sharing: Does it Make a Difference?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Kruse, Douglas, Joseph Blasi, and Rhokeum Park (2006) "Shared Capitalism in the U.S. Economy: Prevalence, Characteristics, and Employee Views of Financial Participation in Enterprises" working paper, presented at NBER-Russell Sage conference on shared capitalism (October).

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

Lazear, Edward P. (2000) "Performance Pay and Productivity," *American Economic Review*, vol. 5, pp. 1346-1361.

Lazear, Edward (1995) *Personnel Economics*, Cambridge, MA: MIT Press.

Lazear, Edward (1999a) "Personnel Economics: Past Lessons and Future Directions," *Journal of Labor Economics*, vol. 17, pp. 199-236.

Lazear, Edward (1999b) "Globalization and the Market for Teammates," *Economic Journal*, vol 109, pp. 15-40.

Levenson, Alec "Knowledge Work and Teams: Analyzing Labor Productivity when Tasks are Interdependent," *Working paper*, 2006.

Lo, Desmond, Ghosh, Mrinal, and Lafontaine, Francine "The Role of Risk, Incentives, and Selection in Salesforce Compensation Contracts," *Working paper*, 2006.

MacDuffie, John Paul (1995) "Human Resource Bundles and Manufacturing Performance: Organizational Logic and Flexible Production Systems in the World Auto Industry," *Industrial and Labor Relations Review*, vol. 48, pp. 197-221.

MacLeod, Bentley (1988) "Equity, Efficiency and Incentives in Cooperative Teams," *Advanced Economic Analysis of Participatory Labor-Managed Firms*, vol. 3, pp. 5-23.

MacLeod, Bentley and Daniel Parent (1999) "Job Characteristics and the Form of Compensation," *Research in Labor Economics*, vol. 18: 177-242.

Mas, Alexandre "Pay, Reference Points, and Police Performance," W*orking paper*, 2005.

Mas, Alexandre and Moretti, Enrico "Peers at Work," W*orking paper*, 2006.

Mas, Alexandre "Labor Unrest and the Quality of Production: Evidence from the Construction Equipment Resale Market," *Working paper*, 2006.

Milgrom, Paul and John Roberts (1995) "Complementarities and Fit: Strategy, Structure and Organizational Change in Manufacturing," *Journal of Accounting and Economics*, vol. 19, pp. 179-208.

Nagin, Daniel S., Rebitzer, James B., Sanders, Seth, and Taylor, Lowell J. (2002) "Monitoring, Motivation, and Management: The Determinants of Opportunistic Behavior in a Field Experiment," *American Economic Review*, vol. 92, pp. 850-873.

David Neumark, Roy J. Bank; Kyle D. Van Nort, "Sex Discrimination in Restaurant Hiring: An Audit Study," *The Quarterly Journal of Economics*, Vol. 111, No. 3. (Aug., 1996), pp. 915-941.

Osterman, Paul (1994) "How Common is Workplace Transformation and Who Adopts It?" *Industrial and Labor Relations Review*, vol. 47, pp. 173-187.

Osterman, Paul (January 2000) "Work Reorganization in an Era of Restructuring: Trends in Diffusion and Effects on Employee Welfare" *Industrial and Labor Relations Review*, vol. 53, pp. 179-196.

Oyer, Paul, "Fiscal Year Ends and Non-Linear Incentive Contracts: The Effect on Business Seasonality," *The Quarterly Journal of Economics*, CXIII, February 1998, 149-185.

Formatted: Underline

Deleted: Abstract

Formatted: Default Paragraph Font

Prendergast, Canice (July 1995) "A Theory of Responsibility in Organizations," *Journal of Labor Economics*, vol 13, no. 3, pp 387-400.

Prendergast, Canice (1999) "The Provision of Incentives in Firms," *Journal of Economic Literature*, vol. 37, pp. 7-63.

Rosen, Sherwin (October 1982) "Authority, Control, and the Distribution of Earnings," *Bell Journal of Economics*, vol. 13;pp. 311-323.

Van Biesebroeck, Johannes (2006), "Complementarities in Automobile Production," *NBER Working paper* #12131.

Figure 1

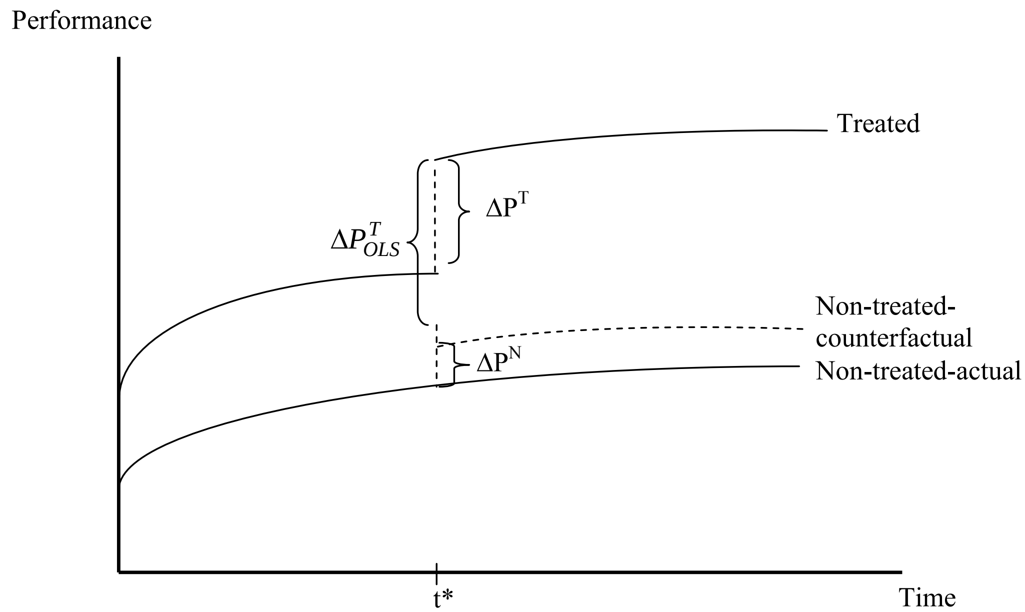

Where t* is the time that the treatment, such as an organizational change, occurs.

Figure 2: Using Selection as an Advantage in Estimating Treatment Effects
(Lazear, 2000)

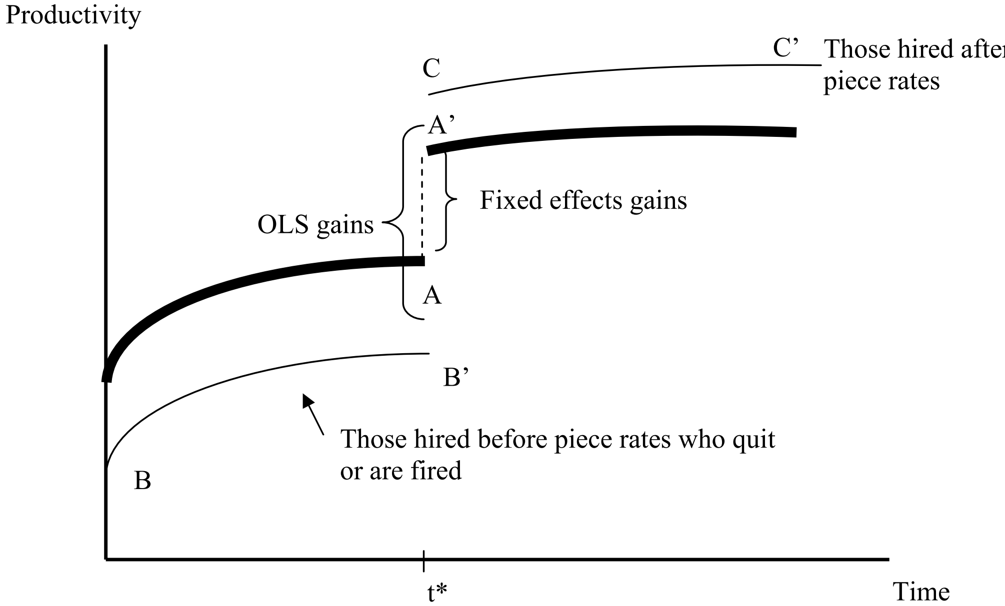

Figure 3: Heterogeneity in the Selection of Production Units
(Hamilton, Nickerson, and Owan, 2003)

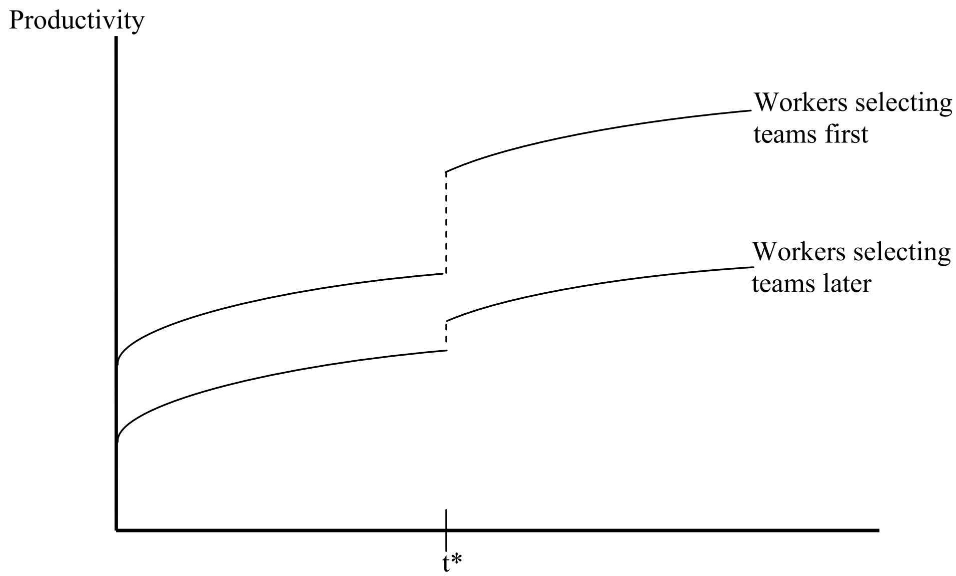

Figure 4: Using Organizational Data to Assess Causality
(Bandiera, Barankay, and Rasul, 2005)