NBER WORKING PAPER SERIES

THE GUY AT THE CONTROLS:
LABOR QUALITY AND POWER PLANT EFFICIENCY

James B. Bushnell
Catherine Wolfram

Working Paper 13215
http://www.nber.org/papers/w13215

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2007

This research was generously supported by the Sloan Foundation-NBER International Productivity Project. We are grateful to Rob Letzler, Amol Phadke and Jenny Shanefelter for excellent research assistance. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

The Guy at the Controls: Labor Quality and Power Plant Efficiency
James B. Bushnell and Catherine Wolfram
NBER Working Paper No. 13215
July 2007
JEL No. J24,L51,L94,M54

ABSTRACT

This paper examines the impact of individual human operators on the fuel efficiency of power plants. Although electricity generation is a fuel and capital intensive enterprise, anecdotal evidence, interviews, and empirical analysis support the hypothesis that labor, particularly power plant operators, can have a non-trivial impact on the operating efficiency of the plant. We present evidence to demonstrate these effects and survey the policies and practices of electricity producing firms that either reduce or exacerbate fuel efficiency differences across individual plant operators.

James B. Bushnell
University of California Energy Institute
2547 Channing Way
Berkeley, CA 94720-5180
bushnell@haas.berkeley.edu

Catherine Wolfram
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720-1900
and NBER
wolfram@haas.berkeley.edu

1 Introduction

In this paper we explore the impact of labor policies on the operations of electric power plants. At first glance, it might seem that workers should have little scope to influence the performance of the electricity industry and that this should be particularly true of the generation sector of the industry, where costs are dominated by the capital required to build plants and the fuel required to operate them. Overall, labor costs constitute a small fraction of generation costs. Yet, in extensive interviews with plant managers and utility executives in the US and Europe, most expressed the belief that the individual skill and effort of key personnel could make a significant difference in the performance of generating plants.

We focus on the role of the plant operator, an individual whose decisions have direct impact on many facets of plant operation. We describe both anecdotal evidence drawn from our interviews and empirical analysis documenting that individual operators do influence the efficiency of plant operations. The existence and tolerance of such an 'operator effect' might seem counter-intuitive. The cost of fuel in power plant operations is orders of magnitude greater than the salary of any individual operator. The savings in fuel costs reaped by highly skilled operators far outweigh any pay premiums they earn.

Having documented the existence of an operator effect, we describe circumstances where companies have taken steps to foster the practices of efficient operators and discourage those of inefficient ones. Generally, however, these appear to be the exception more than the rule. Because labor makes up such a small fraction of industry costs, it is possible that managers have not made human resource polices a priority. Further, it seems likely that the history of regulation in the industry dampened the incentives for operational efficiencies both among managers and workers. This trend may begin to change with the adoption of various forms of regulatory restructuring throughout the industry.

This paper is related to an emerging empirical literature that uses high frequency data to measure productivity differences across workers (see, e.g., Hamilton, Nickerson, and Owan (2004), Bandiera, Barankay, and Rasul (2005), and Mas and Moretti (2007)). While the previous work has focused on measuring the impacts of the workers' environments on their productivity (e.g., teams, compensation scheme, and co-workers), we focus on the size of the differences in productivity across workers at the same firm. Worker heterogeneity is not ordinarily captured in descriptions of firm efficiency based on production functions, but may be an important component of technical efficiency differences across firms. We also place a straightforward economic value on the productivity differences across power plant operators, and show that it is quite large relative to the pay received by the workers.

We begin by giving a general description and historical overview of the electricity industry. We then describe the power production process and the key role of plant operators in that process. We present empirical evidence, drawn from shift and production data from several U.S. power plants, that operators can indeed have a non-trivial impact on plant efficiency. We then conclude with a discussion of labor policies in the industry and describe some isolated attempts to confront and take advantage of the differences in operator skill and effort levels.

2 The Electricity Industry

The electricity industry provides a foundation for much of the industrial and commercial activity in the developed world. In the US, total sales in 2004 were nearly $300 billion per year, making electricity industry revenues comparable to those in the automotive, petroleum products, and telecommunications industries. Yet the industry has typically been viewed as a sleepy one, where innovation, quality improvement, and efficiency efforts have not yielded the rewards garnered in other industries.

Historically, electricity was viewed as a natural monopoly industry. Typically, a single utility company generated, transmitted and distributed all electricity in its service territory. In much of the world, the monopoly was a state-owned utility. Within the U.S., private investor-owned companies supplied the majority of customers although federally- and municipally-owned companies played an important minority role. These companies operated under multiple layers of local, state and federal regulation.

A primary feature of regulation or government ownership, was that revenues were based on costs rather than market factors. Under a typical rate-of-return regulatory structure, electric utilities would be responsible for making investments and operating power systems such that the demand of its franchise customers was met. In return operating expenses would be recovered fully from rates, and capital expenditures would earn a guaranteed rate-of-return. Typically, only the most egregiously wasteful expenditures would be overturned by regulators. It has long been observed that this form of "cost-plus" pricing structure naturally weakens incentives for cutting costs and improving efficiency of operations.[1] The lack of direct competition also made the industry relatively amenable to unionization. The electricity industry has traditionally featured one of the highest union membership rates among U.S. industries. Although deregulation and restructuring has reduced that rate somewhat, as of 2001 the membership rate was around 30%, higher than telecommunications and trucking, and more than twice the level of the U.S. workforce overall.[2]

Industry Structure

The electricity industry is comprised of three main sectors, generation, transmission, and distribution. The generating sector encompasses the power plants where electricity is produced from other energy sources. The transmission system transports the electricity over high-voltage lines from the power plants to local distribution areas. The distribution system includes the local system of lower voltage lines, substations, and transformers which are used to deliver the electricity to end-use consumers. Administrative activities associated with billing retail customers are often included with distribution. Each sector is strongly differentiated from the others in operating characteristics. Transmission is capital intensive, with minimal labor and operating costs. While the natural monopoly arguments for distribution point to the large capital costs associated with replicating the distribution system, from an accounting perspective, most

[1] Joskow (1997) gives a detailed overview of the history and performance of the industry in the US, and of the forces pushing regulatory restructuring and reform.

[2] See Niederjohn (2003).

of the capital in the sector is extremely long-lived, so the main accounting costs are related to operating and maintaining the distribution system. In the US in 2006, about 40% of the over 400,000 employees in the industry worked in distribution and, aside from approximately 25,000 in transmission, the remaining worked in generation.[3]

Within the generation sector, fuel accounts for the bulk of the expenses. For fossil-fired steam generation units, fuel accounted for about 75% of power plant operating costs in 2003 and still over half of the expenses when capital costs are included.[4] By contrast, labor expenses are less than 10% of total generation costs. Although power plants can be extremely large, complex, and expensive facilities, the fundamental process is the conversion of fuel (usually fossil fuel) into electricity. Since fuel is the dominant input into this production process, even small improvements in the efficiency at which fuel is converted into electricity (usually through an intermediate conversion into steam), can result in significant cost savings.

However, within the paradigm of cost-of-service regulation, efficiency of fuel conversion is usually taken to be an immutable, exogenous characteristic of operations rather than a parameter within management's control. In the United States, rates often contained *fuel adjustment clauses*, that would allow for automatic adjustment to electricity rates based upon the costs of fuel consumed by the utility. Thus fuel costs for many utilities were automatically passed on to customers. Although incentive mechanisms have been applied to certain activities, they have rarely extended to fuel consumption within the regulatory framework. One plant manager interviewed for this project indicated that, under regulation, management would not seriously consider an investment aimed solely at improving the efficiency of fuel conversion.

By contrast, environmental considerations can be powerful drivers of investment and operational decisions, both under regulation and competition. A common theme to our interviews was the high degree of focus on how plant operations could be modified to deal with emissions restrictions, or other environmental concerns such as water temperature. The design of the plant

[3]This information is from the Bureau of Labor Statistics, "Employment, Hours and Earnings from the Current Employment Statistics" survey. Information for the industry overall is is based on NAICS code 2211, while the generation, transmission and distribution sectors are five-digit subsets of this.

[4]This figure is also taken from EIA's Electric Power Annual.

and the actions of individual operators can have impacts on these environmental factors. In many cases the goals of fuel-efficiency and emissions mitigation are in conflict with each-other. For example, an oxygen rich fuel mix can reduce NOx emissions, but also reduce fuel-efficiency.

Regulatory Restructuring and Market Liberalization.

Over the last two decades, governments in many countries have privatized and restructured their electricity industries. Restructured electricity markets now operate in much of Europe, North and South America, New Zealand and Australia. These changes were primarily motivated by the perception that the previous regimes of either state ownership or cost-of-service regulation yielded inefficient operations and poor investment decisions. Restructuring of the electricity industry also reflected the natural progression of a deregulation movement that had already transformed infrastructure industries, including water, communications and transportation, in many countries.

Within the United States, electricity restructuring has proceeded unevenly, driven by state-level initiatives. Restructuring has reached an advanced level in much of the Northeast, California, Illinois and Texas. By contrast, the organizational and economic structure of the industry in most of the Northwest and Southeastern US remains unchanged from the 1980s.

Restructuring is primarily aimed at the generation sector. Within restructured markets, wholesale electricity is sold at market-based, rather than cost-based prices. Many power plants have been divested to non-utility owners, many of which have been unregulated affiliates of the former utility owners. During the period from 1998 through 2004, the industry has also experienced an enormous amount of investment in new generation facilities by non-utility operators.

There is some evidence that restructuring, and the ensuing changes in the incentives of generation firms, has had an effect on efficiency in the industry. Bushnell and Wolfram (2006) find that fuel efficiency rates at divested power plants improved roughly 1-2% relative to non-divested plants. Aggregate statistics suggest that employment in the industry has declined substantially, from over 550,000 in 1990 to 400,000 in 2005. Figure 1 plots employment relative to 1990 both for the whole industry, and, beginning in 1997 when employment is broken out by five-digit NAICS code, distinguishing between the generating sector and the transmission and distribution sector of the industry. At least post-1997, the major cuts in the industry were driven by employment reductions at power plants. While these trends are suggestive of a regulatory restructuring effect, there could have been other factors driving the reductions. The results in Fabrizio, Wolfram and Rose (2007) suggest that restructuring was at least partially responsible for the decline, as they demonstrate that regulated power plants operating in states that passed restructuring legislation reduced the number of employees and the level of nonfuel operating expenses by more than both power plants in states that did not pass restructuring legislation and municipally-owned power plants.

3 Plant Operators and Generator Efficiency

In this section, we will focus on the largest single cost in the electricity industry, the consumption of fuel in power plants. Despite the fact that billions of dollars are invested in the research, design, and construction of power plants, and the fact that labor is a relatively small component of power plant costs, there is a widespread belief in the industry that the *quality* of the workforce can have a non-trivial impact on performance. In particular, the decisions of one key employee, the plant operator, can affect the efficiency with which the plant converts fuel into electricity.

As described above, power plant operations are fundamentally the process of converting potential energy in fuel into electrical energy. In general, this process can be further separated into the handling and processing of fuel, the combustion of the fuel, and the generation of electricity from either the exhaust heat or steam produced by the combustion. Depending upon the fuel type, technology, and location of the plant, the processing and monitoring of emissions and other waste products can be another significant component of plant operations. The complexity of these individual processes depends upon the specific technology of the plant. The materials handling and processing is very involved at coal facilities and relatively straightforward at natural gas plants. The combustion process can either entail burning the fuel in boilers to heat water into steam, which in turn rotates a turbine, or the direct use of hot exhaust from combustion to rotate a turbine. The former technology is often described as steam combustion and the latter a combustion turbine (CT).

While power plants employ teams of widely varying sizes and roles, all fossil fired conventional power plants staff a *plant operator*, whose responsibilities are central to the performance of the plant. The plant operator is primarily responsible for the monitoring and control of the combustion process.

At more complex plants, such as coal facilities, an operator controls several aspects of the process that can influence both fuel-efficiency and emissions. These include the rate at which coal *mills* feed pulverized fuel to burners, or even the number of mills and burners in operation. The operator controls the mix of oxygen in the combustion process, and through *dampers* the mix of air and fuel in the mills. Some boilers also allow for adjustment of the angle or *tilt* of the burners within the boiler chamber.

In all cases, these settings are automated to some degree, but the operator has the ability and responsibility to adjust or override automatic settings in the context of monitoring the operational status of the generation unit. The degree to which these decisions have been automated and optimized varies greatly across facilities. As we discuss below, development of automated combustion optimization systems is an area of active commercial and research interest.

In many interviews plant managers and executives expressed a belief that individual operators can have a non-trivial impact on the combustion process. Each facility has idiosyncratic aspects that experienced and motivated operators learn to account for. The act of balancing all of these input parameters was described by one manager as "playing the piano," and one star operator was considered a virtuoso on the instrument.

Another important responsibility of plant operators that was often cited in interviews at coal plants is the operation of *soot blowers* within boilers. In the combustion process, pressurized water is run through pipes or *tubes* and heated by the boilers into steam. As a by-product of the combustion, various impurities and uncombusted material form into soot that settles onto the tubes. The soot forms an insulating layer on the tube that reduces the transfer of heat from the boiler to the water. To counteract this effect, boilers are equipped with soot blowers to jet steam at the tubes and knock off the soot.

While the operator needs to ensure that soot does not accumulate to a detrimental level on tubes, the manner in which the soot is removed can also impact boiler performance. Ideally, blowers would be operated in a sequence that is calibrated to current boiler operations. Alternatively, one unmotivated operator described in interviews, would "trigger all the blowers at once and go have a sandwich." Triggering all the blowers can cause excess soot to circulate throughout the boiler and also reduce the efficiency of combustion.

Overall, most managers we spoke to believed that operators could have a non-trivial impact on the performance of plants. In the next section we present empirical evidence that this is in fact the case.

4 Measuring Efficiency Differences Across Operators

In this section we develop an empirical model to test whether individual shifts or operators impacted the fuel efficiency of their power plants. This task is facilitated by the continuous emissions monitoring system (CEMS) dataset collected by the U.S. Environmental Protection Agency (EPA). The CEMS program was developed to monitor power plant emissions systematically in order to implement environmental controls such as the cap-and-trade system for SOx. The CEMS data track many attributes of generation unit performance on an hourly basis, including the fuel burned and the power output of each facility. We can use these data to obtain an hourly measure of the fuel efficiency of each generation unit.[5] We combine the fuel efficiency data with shift information we obtained from several power companies.[6] Power plants typically comprise multiple boilers and turbines, and each boiler-turbine pair is usually referred to as a

[5]We used a compilation of the CEMS dataset obtained from Platts. The data are described in more detail in the Appendix.

[6]In all cases, the specific identity of the operators was masked in the data.

generating unit. Some multi-unit plants are organized around a single control room, so that the same plant operator controls multiple units (up to seven in our data). By contrast, some plants, typically plants with larger units, have separate control rooms for each unit. To mask their identity, we will refer to the five entities from which we received shift schedule information as "Plant A" through "Plant E," recognizing that in some cases, the operator controls less than the entire plant. The key characteristics of the plants are described in Table 1. Although by no means a comprehensive sample of U.S. generation technology, they do represent some of the standard technologies in use in the U.S. today.

4.1 Empirical Strategy

To test for efficiency differences across operators, we estimate versions of the following equation:

$$ln\left(HEAT\_RATE_{ijt}\right) = \alpha_i + \beta_1 ln\left(OUTPUT_{ijt}\right) + \beta_2 \Delta ln\left(OUTPUT_{ijt}\right) + \beta_3 X_{ijt} + \kappa_j + \varepsilon_{ijt} \quad (1)$$

where t indexes an hour, i indexes the operator and j a generating unit. We estimate this equation for each plant for which we have shift-schedule information.

The dependent variable, $HEAT\_RATE_{ijt}$, is a generation unit's heat-rate, measured as the ratio of the heat content of the fuel input (in Btus) per units of electricity output (measured in kWh). It is inversely proportional to a unit's fuel efficiency and is the industry standard measure of fuel use. We obtained information on the hourly heat rates from the EPA's Continuous Emissions Monitoring System (CEMS) database. As part of the Sulfur Dioxide (SO2) Emissions Permit program, all electric power plants larger than 25 MWs were required to install pollution monitoring devices in their smokestacks. They transmit the data from the monitoring devices to the EPA on a quarterly basis, and the EPA posts it on their website. For some types of units, the fuel input is calculated based on the carbon in the smokestack, while for others, it is measured directly.

The main variables of interest for this study are the $\alpha_i{}'s$, the operator-specific effects. These capture the mean difference in heat rates across operators, controlling for the other variables in the regression. To code them, we needed information on exactly which person was in the control room during a particular hour. We obtained this kind of detailed shift information from three US companies covering five fossil-fuel fired plants. Table 1 summarizes information on the five plants.

For Plant A, a large coal plant in the Southeast, company personnel transcribed entries from the operator logs for one unit at the plant for 2003. Though there are two approximately 1000 MW units at the plant, each unit has its own control room and its own operator at any given hour. Operators are asked to sign into the log when they begin their shift, although for 33% of the hours (24% of the hours when the plant is producing power), the operator did not sign the log. We estimate a single operator effect for all hours when the operator information is missing. The plant operates on a 3-shift schedule, with a morning shift (7AM to 3PM), afternoon shift (3PM-11PM) and a night shift (11PM-7AM). We have information on a total of 12 people, who logged anywhere from 120 to 780 hours over the course of the year. Operators who logged few hours did not necessarily have less industry experience since they could have been assigned mainly to the second "sister" unit at the plant.

For Plant B, a gas plant with two units in the West, company personnel sent us three years worth of spreadsheets with the planned shift schedules. The plant operator controls both units at the same time, so we estimated versions of equation (1) including observations for each unit. We also include a unit fixed effect to capture mean efficiency differences across the two units. These will impact our operator effect estimates to the extent the allocation of output across units varies systematically by operator. There was a fair amount of operator turnover over the three years we analyze, as the time period followed the divestiture of the plant from a regulated utility to a non-regulated merchant firm. Some of the more senior employees at Plant B left to take jobs with the utility parent in part to maintain their favorable treatment in the company benefits programs. Also, for some shifts, two people were scheduled as the operator. We estimate a separate operator effect for each team, giving us 16 total operator effects, though only 12 distinct individuals are represented in the data. Plant personnel work 12-hour shifts, either from 7AM-7PM or 7PM-7AM. Plant B installed combustion optimization software in August 2002 at unit 3 and in August 2003 at unit 4.

Plants C, D and E are all owned by the same firm (Firm X), but the information we have from this firm is the sparsest. Company personnel gave us two single page printouts with the schedules for the four different shifts over two years. The same shift schedules apply to the three Firm X plants that are located in the same state. This means that shift A is always working at the same time at all three plants, but the employees on shift A at Plant C are different from the employees on shift A at Plant D, and the composition of shift A at a particular plant no doubt varies over time. Unfortunately, we don't know anything about the turnover of the personnel working on the shift. Shifts work for twelve hours at a time, either from 7AM-7PM or from 7PM-7AM. The three plants are also quite different from one another. Plant C has two natural gas-fired units that were still in operation as of 2004 with a combined capacity of 760MW (five of the units at the plant were already retired). Plant D is a large plant with seven total natural gas- or oil-fired units ranging in size from 100 to 700 MWs, with the combined potential to generate over 2000MW of total capacity. Some of the units are quite old and run infrequently. Plant E is a natural gas-fired unit with one unit still in operation.

For all units, we control for the unit output level ($ln(OUTPUT)$), change in output over the previous hour ($\Delta ln(OUTPUT)$) and the ambient temperature.[7] The output variables are taken from the EPA CEMS database. We obtained hourly temperature (dry bulb temperature measured in Fahrenheit) by picking the closest weather station from the NOAA surface weather data base (see: http://www.ncdc.noaa.gov/servlets/ULCD). We also include dummy variables for the four hours directly after the unit is started and dummy variables for the type of shift (e.g. night shift versus day shift).

One issue we confront in estimating equation (1) is the possibility that the choice of output

[7]Personnel at one of the plants we visited in the UK showed us calculations they do to benchmark the plant versus a target efficiency value and the main adjustments they make are for unit load, starts and ambient temperature.

level is correlated with the unit's efficiency. This would be the case if, for instance, the plant operator scaled back output when malfunctioning equipment reduced the unit's efficiency. This is equivalent to the endogeneity problem faced in estimating production functions (see, for example, Griliches and Mairesse (1998), Olley and Pakes (1996), Levinsohn and Petrin (2003)). To account for the possibility that both $ln(OUTPUT)$ and $\Delta ln(OUTPUT)$ are endogenous, we instrument for them using electricity demand within each plant's state ($ln(STATE\ DEMAND)$ and $\Delta ln(STATE\ DEMAND)$). Since electricity is not storable, plants are dispatched to meet hourly demand. Depending on congestion on the transmission grid, a plant may serve anywhere from a very local geographic area to a multi-state area. We take the state level as a reasonable representation of the average geographic area a plant could serve.

While it might be interesting to examine whether there are differences in the extent to which individual operators adjust output in response to efficiency shocks, we leave that for future work. Based on our discussions with plant personnel, we perceive that individual operators have some but by no means complete discretion to respond to efficiency shocks. Some of the output adjustments are purely mechanical, for instance, when a malfunctioning pulverizer reduces the amount of fuel that can be fed into a plant boiler. Also, many decisions about output are made by personnel outside the plant, since deciding by exactly how much production should be scaled back when efficiency drops requires coordination across plants in the same geographic area.

4.2 Empirical Results

The $\alpha_i's$ from an instrumental variables estimation of equation (1) for Plant A are summarized in Figure 2. The red squares are at the mean effect for the operator and the blue lines are drawn over the 95% confidence interval. Operator 27 collects all of the missing log entries. Four of the eleven operators (five including operator 27) had statistically significantly lower average heat rates than operator 4, the operator with the highest average heat rate. The estimates suggest that the best operator achieved an average heat rate that was more than 3 percent lower than the average heat rate achieved by the worst operator. To gain perspective on the magnitudes of the estimated effects, consider that if every operator were able to achieve the same average heat rate as the best operator, the unit would save approximately $3.5 million in fuel costs each year.[8] These savings are no doubt considerably larger than the annual payroll costs for operators.

The coefficient estimates on the control variables associated with the specification of equation (1) depicted in Figure 2 are reported in column (1) of Table 2. The second to last row in Table 2 also reports the F-statistic on the joint test that all of the operator effects are zero.[9] For Plant A, the F-statistic is 2.23, suggesting that we can reject the hypothesis that all operators are the same at the one percent level.

Figure 3 summarizes the operator effects estimated for personnel at Plant B, and column (2) of Table 2 reports the coefficient estimates and F-statistic for the specification used to generate the effects summarized in Figure 3. As with Plant A, eight of the fifteen operators are significantly different from the worst operator and the F-statistics suggests that we can reject that all operators are the same at better than the .1% significance level. The operator effects may be more significant at Plant B than they were at Plant A because we have three times as long a time period for Plant B, so the estimates are tighter. The range of operator effects is smaller for Plant B than it is for Plant A, with the most efficient operator only 1.9% better than the least efficient operator. We spoke with engineers from both coal and gas plants who suggested that operator decisions are likely to have more impact on efficiency at coal plants.

Unlike for Plants A and B, the operator effects at Plants C, D and E (recall that they are all owned by Firm X) were estimated to be small and statistically indistinguishable from zero. The largest difference between the best and worst shifts was .0020 (s.e. .0019) at Plant C. This point estimate is an order of magnitude smaller than the similar measures at Plants A and B. Overall, the results suggest there are no discernible differences between the four shifts at any of Firm X's plants. It is instructive to consider why we might find differences across operators at Plants A and B, but not at Plants C, D and E. For one, the shift information that we received from Firm

[8]This calculation assumes the plant operates at a 90% capacity factor, with fuel costs of $25/MWh and that the best operator worked for 10% of the time.

[9]The F-test for Plant A excludes operator 27, the operator effect used to collect all hours when the operator log was left blank.

X is much less precise than the information for Plants A or B, so the estimates could be biased to zero because of classical measurement error. For instance, since we only have information on four shifts, the shifts were scheduled to work almost 2,200 hours per year. No doubt operators, especially those with considerable seniority, are working much less than this per year, suggesting that each shift contains more than a single operator. Also, as we noted in comparing Plant A to Plant B, operators have less room to affect efficiency at gas plants. Finally, plant personnel at Plant C described an in-house computer program that they used to instruct operators about the optimal setting for plants, suggesting that operators at the Firm X plants are less likely to make different decisions about plant operations.

Note that there is reason to believe that all of the operators effects we measured are biased to zero. For one, we only have information on the operator and not the plant staff supporting him (all of the operators we have on record were male). It's possible that we could see larger differences if we could control for the supporting staff as well. Second, even for Plant A, where we have operator log information, there may be measurement error in our independent variable.

The coefficient estimates on the control variables summarized in Table 2 are for the most part as expected. For all plants except Plant A, the coefficient on $ln(OUTPUT)$ are negative and statistically significant, suggesting that plants are more fuel efficient at higher output levels. Also, as would be expected if operators are reducing output in response to negative efficiency shocks, instrumenting for $ln(OUTPUT)$ causes the coefficient to fall towards zero. For example, an OLS estimate of equation (1) using data on Plant B yields a coefficient on $ln(OUTPUT)$ of -.121 (s.e. = .002).[10] Similarly, the coefficient on $\Delta ln(OUTPUT)$ is positive and statistically significant at all plants, suggesting that increases in output degrade efficiency and reductions improve efficiency.[11] Also, the F-statistics on the first stages are large, suggesting that our

[10] The significance of the operator effects are not sensitive to the estimates of $ln(OUTPUT)$. In addition to the specifications we report, we also estimated other specifications that allowed $OUTPUT$ to take different nonlinear forms. The estimates of the F-statistics on the operator effects were qualitatively very similar, i.e. suggesting that operators at Plants A and B differed from one another but those at Firm X's plants did not.

[11] We also estimate specifications that allowed the effect of an output change to differ for positive and negative changes. Both effects were positive, suggesting that a reduction in output does lead to a lower heat rate (more efficient).

instruments work quite well.

The coefficient estimates on $TEMPERATURE$ are all positive and statistically significant, consistent with what engineers told us to expect. Only two of the five $DAYSHIFT$ variables are significantly different from zero, and one is positive and small and the other is negative and quite small (suggesting at most a .5% difference across shifts). Except at Plant A, the $START_{t-X}$ for $X \subset 2, 3, 4$ dummies are positive, suggesting that fuel efficiency is compromised after starts. There were only 13 starts at Plant A, so these variables are imprecisely estimated. Also, since starts are associated with rapid changes in output, the heat rate variable can be very noisy.

5 Labor Policies and Operator Performance

We have described the critical role that plant operators play in the operation of power plants, and presented anecdotal and empirical evidence that operators can have a significant impact on the efficiency of plant operations. Given this evidence, two important questions arise. Why is such a variation in performance tolerated by firms, and what can firms do to take advantage of the skills and experience of the strong performers?

5.1 Human Resource Policies

Aggregate statistics and our interviews with power plant managers both suggest that labor policies in electricity generation have been undergoing a dramatic transformation over the last 10-20 years. This transformation has coincided with the rise of non-utility power producers, the privatization of publicly owned utilities outside of the U.S., and the advent of regulatory restructuring. It is reasonable to conclude that the competitive pressures created by these developments provided the impetus for these changes. However, it is worth noting that these changes have not been limited to regions where power plants have been divested or deregulated. Also, many interviewees cited the adoption of automated monitoring technology beginning in the late 1980s as a factor in the declining employment rates.

In general, the historic labor picture at power plants was heavily unionized with inflexible work rules and promotion policies. There were several layers of job categories and restrictions on utilizing employees in roles outside of their categories. Staffing levels were also, by today's standards, quite high. Promotion was largely based upon tenure at a plant or with the company. Certainly a minimum level of competence was required for promotion, particularly to the operator level. However, among those employees able to exceed a certain minimum threshold of performance, there was little effort to differentiate among the quality of employees when determining promotions.

Since the mid 1980's employment levels have steadily declined. Plant F, a coal-plant in England visited for this project, is representative of this trend. There were 285 employees at the plant when we visited, down from a peak of over 700 before the plant was privatized in the early 1990's. This trend is shared among most liberalized electricity markets, but not restricted to those facing full competitive pressures. Plant G, a coal plant in Alabama also visited for this project reported 320 employees in 2004, down from a peak of over 450 despite the fact that its regulatory status has remained unchanged. Among the positions eliminated was a full-time groundskeeper, cited to us as an example of previous excesses given the paucity of grass around the plant.

As mentioned above, aggregate statistics suggest a pronounced reduction in power plant employment throughout the U.S. These reductions are most pronounced in areas actively pursuing some form of deregulation (see Fabrizio, Wolfram, and Rose, 2007). The largest reductions overall appear to be a plants divested from regulated utilities to non-utility operators (see Bushnell and Wolfram, 2006).

The reduction in employment has coincided, at least in restructured states, with a declining influence of unions and increasingly flexible work rules. In two separate interviews, managers described how previously, a shift was staffed with a number of specialists, including mill workers, electricians and boilermakers. Union work rules prohibited job sharing. In the late 1990s, management had been able to renegotiate union contracts, in some cases when the plants were divested to new owners, to allow workers to be classified generally as power plant operators. As a result, workers at the restructured plants we visited were valued for their broad skill sets, and staffing levels fell.

According to managers at some plants, wage levels have in many cases risen as the number of employees has been reduced and responsibilities expanded.[12] Promotion policies have also become less rigid. One operator at Plant F in England rose to his position in just over two years, much faster than would have been possible under the plant's previous tenure-based promotion scheme. The merchant owner of Plant B replaced a large fraction of the employees it inherited from the regulated utility when it purchased the plant, drawing its new employees largely from ex-Navy technicians and engineers. By contrast, Firm X, also a merchant company operating plants it had purchased from regulated utilities, has retained most of the employees at the plants it purchased.

Despite these broad trends that indicate increasing productivity at power plants in liberalized electricity markets, in most cases we found little focus on the quality of specific employees, beyond standard promotional policies. In particular, in most cases there were no specific initiatives designed to address the operator effects on fuel efficiency that have been described above, despite a widespread consensus that such effects are meaningful. That said, there were some efforts at linking bonuses to corporate or plant performance, and one specific effort to link employee pay to the efficiency of the plant. We describe these programs below.

5.2 Performance Pay

All plants we visited paid bonuses to their employees loosely based upon some measure of performance. In some cases, as with Plant G in Alabama, these bonuses were largely linked to corporate financial performance and therefore were more a version of "profit-sharing" than incentive pay. Bonuses at many plants also reflected conventional HR policies, such as a linkage to favorable performance reviews by supervisors, the completion of assigned tasks on time, and

[12] Shanefelter (2006) uses BLS data to describe a picture consistent with these claims.

limited absenteeism. In several cases, such as Plant F, bonuses were linked to aggregate measures of plant's performance, such as the achievement of certain fuel efficiency and availability targets. For the most part, however, such bonuses did not attempt to distinguish between the performance of specific employees within a given plant.

One notable exception to these policies was a performance pay initiative attempted at Plant F in England in the mid 1990's. Plant F is a large coal-fired plant that had been built by the government-owned Central Electricity Generation Board (CEGB) and privatized in the early 1990s. The plant has since changed hands multiple times. Since before privatization, substantial efforts were made to monitor and document the plant's performance along a large number of efficiency measures. These efforts evolved into an automated system able to monitor, quantify, and report the "cost of [efficiency] losses" at the plant. The cost of losses calculation was highly sophisticated and attempted to control for all relevant exogenous impacts on plant operations, such as fuel quality, ambient temperature, and the output level of the plant. It generated detailed reports breaking down efficiency losses to specific processes within the plant.[13] Initially (and currently) these data were aggregated into monthly performance reports and utilized by managers as a general tool for helping to focus efficiency efforts. These measures would be reviewed at monthly meetings of all section heads, including representatives from operations, commercial performance, and maintenance.

In 1995, managers attempted to utilize the cost of losses system in a more direct fashion by linking it to performance bonuses for specific shifts. Recognizing the disparity in performance and losses between shifts, manager's believed that the incentives provided by such a linkage would help to focus under-performing operators and shifts and help to improve their efficiency at least to levels attained by higher performing shifts. In doing so, managers implicitly expressed a belief that these performance disparities were largely effort-based, rather than a result of differences in the inherent acumen or talent of the operators. The pay differentials created by the bonuses were still quite modest, amounting to about 1 percent of annual pay.

[13]The cost of losses report would decomposes performance measures to report the losses due to several factors including turbine losses, boiler efficiency, fuel feed trains, and exhaust pressure.

Even with this modest incentive, however, managers did notice marked changes in performance between shifts. Unfortunately they were not the kinds of effects that they intended to induce. The incentive scheme was based upon the *relative* performance in the cost of losses of each shift. Operators quickly discovered that a degradation in the performance of *other* shifts could be as rewarding as an increase in their own efficiency. It appears that there are more and easier options for sabotaging other shifts than for improving own performance. Managers found that operators would sometimes avoid blowing soot throughout their shift, forcing excessive blowing upon the next shift, or triggering all the blowers simultaneously at the very end of their shift, leaving the next shift to deal with the resulting residue. In such an environment there was growing acrimony between shifts and operators. Eventually, managers at Plant F dropped the incentive scheme, and shifted toward a system of rewarding the pooled performance of all shifts. Although the direct influence of individual effort and performance on such pooled incentives is diluted, managers claimed that efficiency improved roughly half a percent under this new scheme.

5.3 Combustion Optimization Software

The experiment with performance pay at Plant F can be viewed as an attempt to elevate the efficiency of under-performing operators at least up to the level observed in the better operators by applying incentives intended to increase focus and effort. A more recent trend at power plants may also result in more balanced performance among operators by reducing the impact of their individual performance. This trend is the adoption of automated combustion optimization software and systems. In general these systems use learning algorithms to attempt to customize operating protocols to the specific idiosyncrasies of a specific plant. The more ambitious of these systems take much of the influence over burner angles, fuel flow, oxygen content, etc. out of the hands of the operator. In theory such systems should reduce the disparities between operator performance. Indeed, the vendor of one such system, NeuCo, claims that its systems can help to ‘make the worst operator at least as good as the best.’ The adoption of these systems is still in its early stages, and we were not able to attain sufficient data to adequately evaluate such claims.

However, two factors that were raised during our interviews indicate that, at least in the near future, the impact of such systems on fuel efficiency may be small. First, these systems are being utilized primarily for the purpose of reducing emissions, rather than improving fuel efficiency. Second, in many cases operators have been hostile to yielding control over operations to these systems. In one plant we visited, an installed control system had been converted to an 'advisory mode' that provided recommendations, but direct control was left to the human operator.

That said, managers at the Firm X plants firmly believed that the optimization systems they had installed would significantly reduce if not eliminate any operator effects. Our empirical analysis supports their view. By contrast, Plant B installed a NeuCo system in the middle of our sample period. The system had been installed to help the plant address NOx emissions, rather than fuel efficiency. When we included a dummy variable equal to one after the adoption of the optimization software, we did not detect a statistically significant impact on either the overall fuel efficiency of the units or on the relative operator effects at the plant, although we observed only nine operators who worked before and after the installation.

6 Conclusions

Labor policies in the electricity industry have been significantly impacted by its historical status as either a publicly owned or regulated utility business. At the same time, evaluating and improving labor practices may have been given low priority due to the fact that labor costs constitute a small portion of industry costs. We present evidence that, despite the fact that overall labor costs are small, the *quality* of certain workers can have a significant impact on the operations of power plants. Power plant operators, in particular, can influence the fuel-efficiency of the plants under their control in a myriad of individually small, but in aggregate consequential, ways. There is good reason to believe that this effect is more prominent in the more complex coal facilities than in gas-fired power plants.

In our examination of performance data from U.S. power plants we find that the individual operators could influence fuel efficiency by more than 3%. While this figure may sound modest, it translates into a difference worth millions of dollars in annual fuel costs at larger facilities. Despite what appears to be a widespread belief in an 'operator effect' amongst plant managers, there have been relatively few attempts to address the impacts of these effects. We have documented one failed attempt at performance-based incentive pay, and described how the advent of automated combustion optimization systems may reduce or eliminate operator effects. Even the roll-out of such automated systems has been relatively slow, and more focused on environmental considerations than on efficiency concerns. It is worth noting that market incentives have only recently been introduced in the industry. The process of regulatory restructuring is less than a decade old in most of the world, and this is a relatively short-time in a historically slow moving industry. It remains to be seen whether firms facing more exposure to market incentives will prove to be more adept at taking advantage of operator effects, or whether such effects are an immutable characteristic of the power generation business.

More generally, our results provide a clean measure of the extent of worker heterogeneity within the same job description at a particular plant. It is possible that other industries would show less heterogeneity, perhaps because labor practices have received little attention in the electricity industry relative to other industries where labor is a larger fraction of overall employment. It is also possible that the true heterogeneity across workers would be larger in other industries, and the fact that managers have clean measures of worker output in electricity helps keep it in check. For example, Mas and Moretti (2007) report a 21% difference between supermarket cashiers in the top and bottom deciles. At any rate, worker heterogeneity is not ordinarily captured in descriptions of firm efficiency based on production functions, but may be an important component of technical efficiency differences across firms.

References

[1] Bandiera, Oriana, Iwan Barankay and Imran Rasul (2005). "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics*, 120(3), 917-962.

[2] Bushnell, James and Catherine Wolfram (2005). ""Ownership Change, Incentives and Plant Efficiency: The Divestiture of U.S. Electric Generation Plants." CSEM Working Paper WP-140, University of California Energy Institute. March. Available at www.ucei.org.

[3] Christensen, Laurits R. and William H. Greene (1976). "Economies of Scale in U.S. Electric Power Generation," *Journal of Political Economy*, 84 (4), 655-676.

[4] Fabrizio, Kira, Nancy Rose and Catherine Wolfram (2007). "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency," *American Economic Review*, forthcoming.

[5] Griliches, Zvi and Jacques Mairesse (1998). "Production Functions: The Search for Identification," in Steinar Str¿m ed., *Econometrics and Economic Theory in the 20th Century.* Cambridge University Press: Cambridge, UK.

[6] Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan (2003). "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation," *Journal of Policital Economy.* 111(3), 465-497.

[7] Joskow, Paul (1997). "Restructuring, competition, and regulatory reform in the U.S. electricity sector." *Journal of Economic Perspectives.* 11(3), 119-38.

[8] Kleit, Andrew and Dek Terrell (2001). "Measuring Potential Efficiency Gains From Deregulation of Electricity Generation: A Bayesian Approach," *The Review of Economics and Statistics*, 83(3), 523-530.

[9] Knittel, Christopher R. (2002). "Alternative Regulatory Methods and Firm Efficiency: Stochastic Frontier Evidence from the US Electricity Industry," *The Review of Economics and Statistics*, 84 (3), 530-540.

[10] Levinsohn, James and Amil Petrin (2003). "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 70 (2): 317-41.

[11] Mas, Alexandre and Enrico Moretti (2007). "Peers at Work." NBER Working Paper No. w12508. Available at www.nber.org.

[12] Niederjohn, M. Scott (2003), "Regulatory Reform and Labor Outcomes in the U.S. Electricity Sector," *Monthly Labor Review*, 126(5), 10-19.

[13] Nerlove, Marc (1963). "Returns to Scale in Electricity Supply," in Christ et al. eds. *Measurement in Economics*. Stanford University Press: Stanford, CA.

[14] Olley, Steven and Ariel Pakes (1996). "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64 (6), 1263-1297.

Data Appendix

Our primary data sources are BaseCase and PowerDat, two databases produced by Platts (see www.Platts.com). Platt's compiles data on power plant operations and characteristics from numerous public sources, performs limited data cleaning and data analysis and creates cross references so that the data sets can be linked by numerous characteristics (e.g. power plant unit, state, grid control area, etc.). We relied on information from Platts for the following broad categories.

Unit Operating Profile

BaseCase contains hourly power-plant unit-level information derived from the Continuous Emissions Monitoring System (CEMS) database collected by the Environmental Protection Agency. The EPA assembles this detailed, high quality data to support various emissions trading programs. The CEMS data are collected for all fossil-fueled power plant units that operate more than a certain number of hours a year. The dataset contains hourly reports on heat input, gross electricity output and pollutant output. We calculate the heat rate by dividing heat input (measured in mmBtus) by gross electricity output (measured in MWh). By construction of the heat rate variable, our sample is limited to hours in which the unit was producing positive gross electricity output.

System-level Demand Characteristics

Data on system level demand are taken from the PowerDat database, also compiled by Platts. These data report the monthly minimum, maximum, mean,and standard deviation of load by utility, as well as the average daily maximum over a month. Platts compiles this information from survey data collected by the EIA and reported in its form 714.

Plant and Unit Characteristics

Unit characteristics are taken from the "Base Generating Units" and "Estimated Fossil-Fired Operations" data sets within BaseCase.

We merged data from Platts to several additional sources.

Shift Schedules

We obtained shift schedules from three companies covering operations at five power plants. For Plant A company personnel transcribed entries from the operator logs for one unit at the plant for 2003. Though there are two approximately 1000 MW units at Plant A, each unit has its own control room and its own operator at any given hour. Plant operators are asked to sign into the log when they begin their shift. For Plant B, a gas plant with two units, company personnel sent us three years worth of spreadsheets with the planned shift schedules. The plant operator controls both units at the same time. The information we have from Firm X is the sparsest. Company personnel gave us two single page printouts with the schedules for the four different shifts over two years. The same shift schedules apply to all three of Plant X's plants in the same Western state. This means that shift A is always working at the same time at all three plants, but the employees on shift A at plant 1 are different from the employees on shift A at plant 2, and the composition of shift A at a particular plant no doubt varies over time.

Ambient Temperature-Hourly

We obtained hourly temperature data by weather station from the Unedited Local Climatological Data Hourly Observations data set put out by the National Oceanographic and Atmospheric Administration. Further documentation is available at:

http://www.ncdc.noaa.gov/oa/documentlibrary/ulcd/lcdudocumentation.txt

We calculated the Euclidean distance between each weather station-power plant combination, using the latitude and longitude for each power plant and for each weather station. Then, for each month, we found the weather station closest to each power plant that had more than 300 valid temperature observations. For hours when the temperature was missing, we interpolated an average temperature from adjoining hours.

Table 1: Characteristics of Units Analyzed

	Plant A	*Plant B*	*Plant C*	*Plant D*	*Plant E*
Units under Operator's Control	1	2	2	7	1
Unit(s) Characteristics					
Size (MW)	950	700	700	2000	250
Primary Fuel	Coal	Gas	Gas	Gas & Oil	Gas
Year Installed	1975	1965	1965	1955-1970	1965
Operating Statistics					
Average Capacity Factor (%)	90	56	43	43	45
Starts/year	14	26	31	42	6
Efficiency (MMBtu/MWh)					
Average	8.9	10.2	10.5	11.4	10.4
Std. Dev.	.5	1.0	3.7	3.8	1.3
Positive Output (MW)					
Average	826	181	144	184	92
Std. Dev.	110	82	93	163	60
$Output_t/Output_{t-1}$					
Average	1.02	1.05	1.08	1.06	1.02
Std. Dev.	.85	.78	1.00	.73	.27
Combustion Optimization In Use ?	No	In later periods		In-house version	
Shift Schedule Information					
Source	Operator logs	Bi-weekly schedule		Annual schedule	
Period covered	2003	2001-2003		2002-2003	
Shift length	8 hour	12 hour		12 hour	
Total operators	12 individuals	11 individuals		4 teams	
N	7,578	33,490	18,003	28,790	15,339

Note: Unit size rounded to 50MW increments, and unit installation years rounded to half-decade.

Table 2: Efficiency Regressions

Dependent variable: ln(Efficiency)

	Plant A	*Plant B*	*Plant C*	*Plant D*	*Plant E*
ln(Output)	-0.040	-0.119***	-0.118***	-0.191***	-0.071***
	(0.040)	(0.002)	(0.003)	(0.003)	(0.003)
Δln(Output)	0.084***	0.051***	0.039***	0.024***	0.022***
	(0.019)	(0.003)	(0.007)	(0.006)	(0.007)
$Start_{t-2}$	-0.141	0.327***	0.278***	0.222***	0.261***
	(0.132)	(0.021)	(0.079)	(0.038)	(0.055)
$Start_{t-3}$	-0.155**	0.143***	0.096**	0.077***	0.091***
	(0.074)	(0.013)	(0.043)	(0.015)	(0.024)
$Start_{t-4}$	-0.061	0.060***	0.083*	0.035***	0.045***
	(0.051)	(0.009)	(0.043)	(0.007)	(0.015)
Day Shift	-0.008	<0.001	-0.004***	0.002*	-0.001
	(0.006)	(0.001)	(0.001)	(0.001)	(0.001)
Evening Shift	-0.001				
	(0.007)				
Temperature	0.0005***	0.0003***	0.0005***	0.0003***	0.0007***
	(0.0002)	(0.00009)	(0.00006)	(0.00008)	(0.0001)
Number of distinct operators	11	16	4	4	4
F-statistic on operator effects	2.23	3.90	.39	.11	.42
(p-value)	(.01)	(<.0001)	(.76)	(.95)	(.74)
N	7,578	33,490	18,003	28,790	15,339

All specifications estimated using instrumental variables with ln(State Load) and Δln(State Load) used as instruments for ln(Output) and Δln(Output). Unit fixed effects are included where operators control multi-unit plants and year-effects are included where data span multiple years. Standard errors (in parentheses) are robust to serial correlation within a day.
* significant at 10% level; ** significant at 5% level; *** significant at 1% level

FIGURE 1: Electricity Industry Employment Relative to 1990

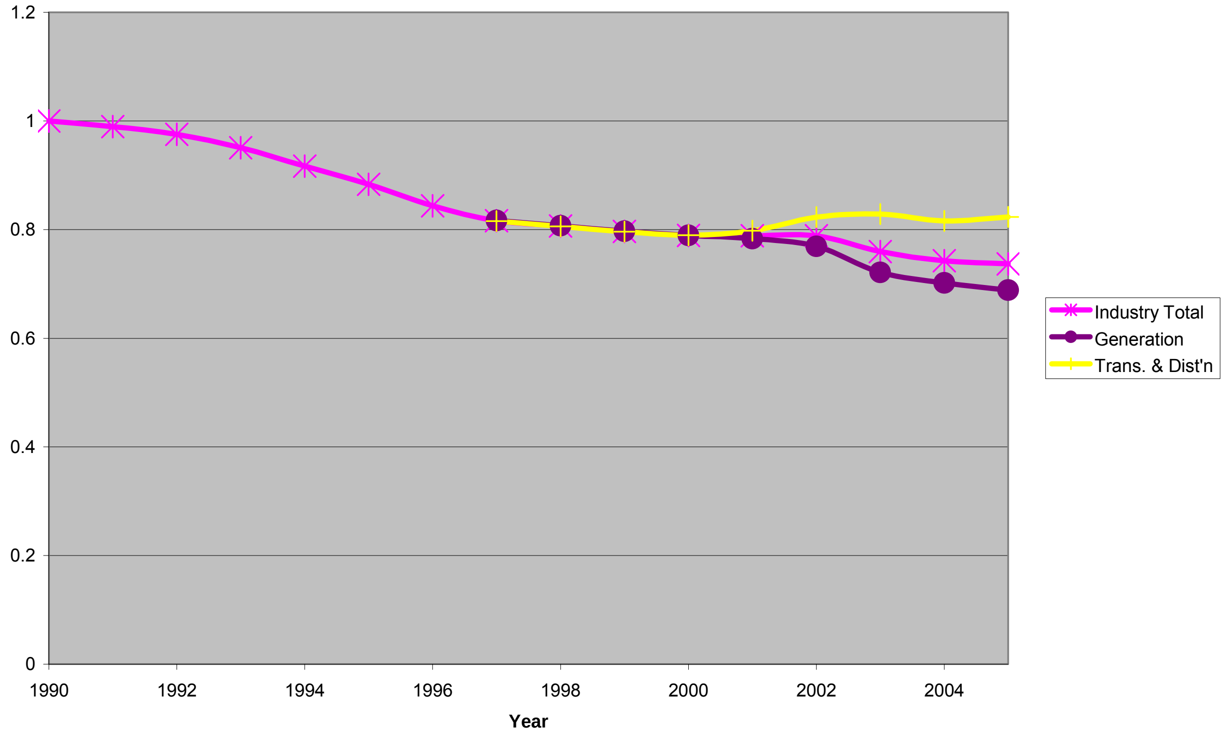

Source: Bureau of Labor Statistics, "Employment, Hours and Earnings from the Current Employment Statistics" survey.

FIGURE 2: Relative Heat Rates by Operator - Plant A

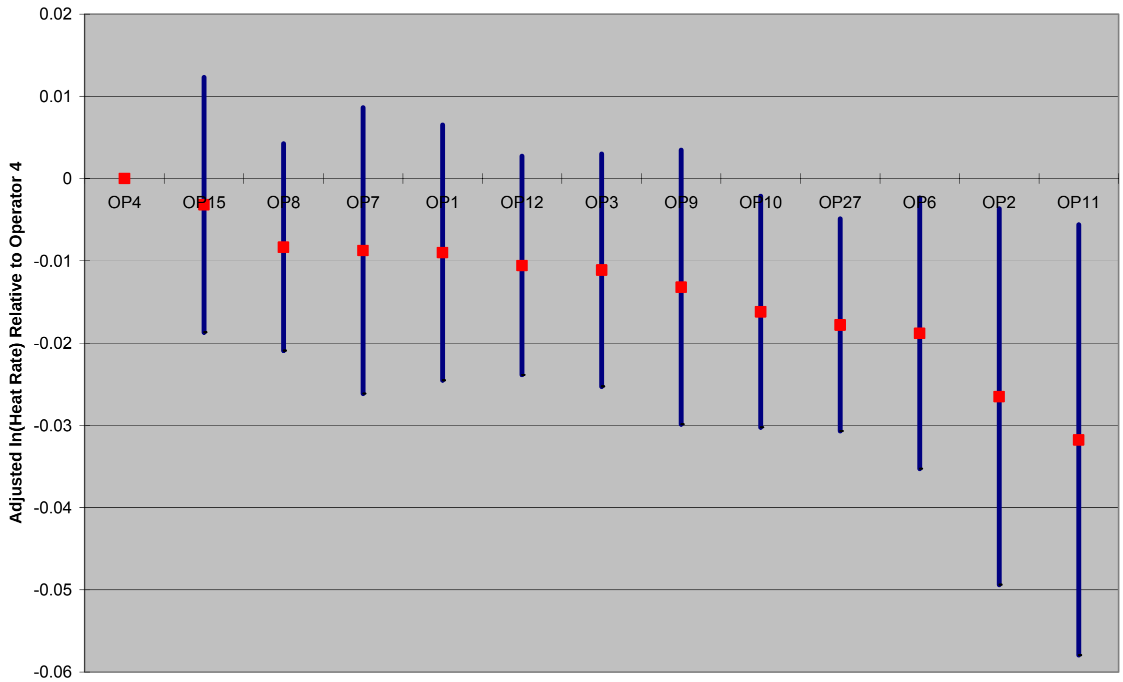

Note: The red squares are drawn at the estimated α_i from equation (1) for each operator, while the blue lines are drawn over the 95% confidence interval. Low values of α_i indicate that an operator achieved a lower average heat rate, i.e., was more efficient, relative to the least efficient operator (Operator 4).

FIGURE 3: Relative Heat Rates by Operator - Plant B

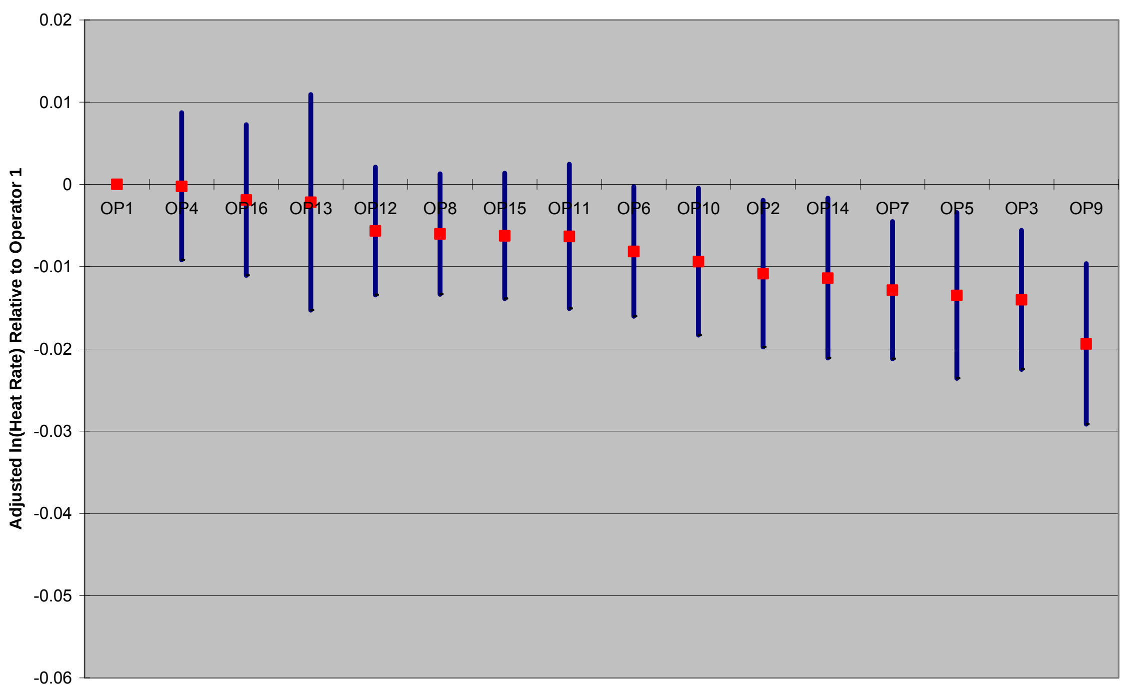

Note: See Figure 2.