Who Loses Under Power Plant Cap-and-Trade Programs?

Estimating the Impact of the NO_x Budget Trading Program on Manufacturing Employment

Mark Curtis
Georgia State University
ecurtis3@gsu.edu

August 20, 2013

Abstract

This paper tests how a major cap-and-trade program, known as the NO_x Budget Trading Program (NBP), affected labor markets in those regions where it was implemented. The cap-and-trade program dramatically decreased levels of NO_x emissions and added substantial costs to energy producers. Using a triple-differences approach that takes advantage of the geographic and time variation of the program as well as variation in industry energy-intensity levels, I examine how employment dynamics changed in manufacturing industries whose production process requires high levels of energy. After accounting for a variety of flexible state, county and industry trends, I find that industries in the top quartile of the energy intensity index lost 4.4% of employment relative to industries in the bottom quartile of the energy intensity index. Young workers experienced the largest employment declines and earnings of newly hired workers fell after the regulation began. Employment declines are shown to have occurred primarily through decreased hiring rates rather than increased separation rates.

I thank Spencer Banzhaf, Paul Ferraro, Ann Ferris, Barry Hirsch, Thomas Holmes, Josh Linn, Kyle Mangum, Erin Mansur, Karen Palmer, Joseph Shapiro and Ian Schmutte for helpful comments and conversations. I also thank seminar participants at the 2012 NBER Summer Institute, Camp Resources XIX, the 2012 AERE Summer Conference, Georgia State University and the Federal Reserve Bank of Atlanta for comments on an earlier version of this paper. Computing resources were provided by Cornell's Social Science Gateway which is supported through NSF grant #1042181.

1 Introduction

The Environmental Protection Agency's regulation of the energy sector has become a highly contentious topic in the public sphere. Proponents of the federal regulations emphasize the health benefits that accrue to society when fewer pollutants are emitted by power plants, while critics claim that the current regulations harm the economy by imposing significant costs on industry. Job loss in particular is cited as a primary means by which regulation inflicts damage on the economy and has gained special attention in recent years as the unemployment level has risen sharply.[1]

The federal law that lies at the heart of this controversy is the Clean Air Act Amendments (CAAA), which vests authority in the EPA to regulate the emissions of polluting industries. Broadly speaking, there are two types of regulations that the CAAA imposes on pollution-emitting establishments. The first, known as the National Ambient Air Quality Standards (NAAQS) began in the 1970's as a result of the 1970 CAAA and required polluting plants in counties with poor air quality to adopt "lowest available emission rate" technology. The labor market impact of the NAAQS and its subsequent expansions has been studied extensively by economists over the past decade (Greenstone 2002; Kahn and Mansur 2010; Walker 2011; 2012). These studies examine changes in manufacturing employment in counties that fail to meet NAAQS attainment standards and are thus subject to tighter regulations.

Since the passage of the 1990 CAAA, however, a second and wider reaching policy has taken form with the intent of regulating interstate air pollution. As counties began to realize that their own air quality was affected not only by local polluters but also by polluters from other counties and states, there was a push for the regulation of all establishments whose air pollutants crossed state boundaries. In 1990, the Acid Rain Program established a national cap-and-trade program for Sulfur Dioxide (SO_2) and in 2003 and 2004 the Nitrogen Oxide (NO_x) Budget Trading Program (NBP) was established in nineteen states east of the Mississippi. Despite the size and far-reaching impact of these two cap-and-trade programs, there has been no empirical research that has sought to evaluate their labor market implications.

From an efficiency standpoint, cap-and-trade programs produce desired outcomes by allocating permits which grant the owner the right to emit a given quantity of a pollutant. The total quantity of permits is limited by the environmental authority and

[1] See *Wall Street Journal* July 26, 2011 op-ed "The Latest Job Killer from the EPA". Also "Getting Ready for a Wave of Coal-Plant Shutdown" was the most read post on Ezra Klein's *Washington Post* "Wonkblog" in 2011.

firms are allowed to trade permits in an open and competitive market. Allowing firms to participate in a market for permits guarantees that firms for whom pollution abatement costs are cheapest will be the first to reduce their emissions. While they provide an efficient and market-based solution, tradable permit systems may still have potentially severe redistributional implications. Firms which previously paid nothing to pollute now face a new input cost. This may lead to an increase in marginal cost and may result in a re-optimization of their input bundle.

The Acid Rain Program was the first such large scale cap-and-trade program in the United States but, given that it applied to all power plants in the country, empirically estimating its impact on employment has been difficult due to the lack of a valid geographic counterfactual.[2] However, for a variety of reasons, the attributes of the recently implemented NBP make for a policy whose impacts are both important and possible to identify. First, the NBP had a major impact on energy production. The regulation of 2,250 gas, oil and coal-fired electric generating units plus 350 large industrial units forced NO_x emitting firms to make difficult and costly decisions on how to comply with the cap-and-trade scheme. Overall, complying with the NBP was expected to add $2.15 billion dollars of annual costs to utilities, which would largely be passed on to consumers in the form of higher prices (Palmer et al. 2001). Second, this policy was implemented in nineteen eastern states over a period of two years. Many states were not exposed to the NBP and, under certain assumptions discussed below, may be considered a valid counterfactual after controlling for preexisting differences. Finally, industries in the NBP region that require high levels of energy in their production process would be expected to be most affected by the NBP. These sources of geographic, time and industry heterogeneity form the basis for the identification strategy used to determine the impact of the NBP on manufacturing employment. To the best of my knowledge, this is the first such credible study of the employment effects of any major EPA cap-and-trade program.

One reason why empirical methods may prove particularly useful is that economic theory gives no clear intuition regarding the effect of power plant regulation on manufacturing employment outcomes (Berman and Bui 2001). Environmental regulation that causes an exogenous shock in the price of energy is likely to lead to two competing employment effects on an establishment's intensive margin. First, given that capital and energy are complements and capital and labor are substitutes, a positive shock in the price of energy may lead plant managers to employ more labor and less capital. How-

[2] For this reason, researchers have generally used structural models to estimate the costs and benefits of the Acid Rain Program (see Burtraw et al. (1998)).

ever, an increase in the price of electricity will also increase marginal costs and decrease the demand for labor due to a decline in production. Furthermore, firm owners may adjust the extensive margin as they take production costs into account when determining plant location. Plants in areas with increased energy prices are more likely to be shut down and newly constructed plants are more likely to be built in regions that did not experience a positive shock to energy prices.[3] Each of these effects will be at play whether the manufacturing plant produces its own energy and is directly regulated or chooses to purchase energy from now regulated utilities.[4]

While qualitative predictions from theory may be somewhat ambiguous, recent empirical research studying the impact of environmental regulation on employment has shown either no change (Berman and Bui 2001) or a decrease in the employment levels of regions where regulation of electric utilities has been implemented (Greenstone 2002; Kahn and Mansur 2010; Walker 2011). Furthermore, the broader literature on the impact of environmental regulation has consistently found a negative impact of regulation on plant openings and productivity levels (Becker and Henderson 2000; Dean et al. 2000; List et al. 2003; Henderson 1996; Hanna 2010; Greenstone et al. 2012).[5]

Using a similar technique as that employed by Kahn and Mansur (2010) to study manufacturing industry location, this paper takes advantage of the heterogeneity in industry energy intensity levels to perform a triple differences (DDD) analysis that estimates the NBP's impact on manufacturing employment in the regions where it was enforced. In addition to examining the impact on employment levels, this paper also looks at worker turnover, earnings and examines heterogeneity of the impact by worker age. Using County Business Patterns and Quarterly Workforce Indicator data on employment as well as NBER's Productivity Database, I am able to account for important state, county, industry and year controls as well as state, county and industry trends. These controls will prove important as industries tend toward regional agglomeration and as the manufacturing sector as a whole has experienced general geographic shifts

[3]Carlton (1983) and Kahn and Mansur (2010) find strong empirical evidence that electricity prices are a major determinant of manufacturing establishment location decisions.

[4]The 350 non-EGU regulated units belong to 140 large manufacturing plants (Author's calculation based on records from EPA's website). Not surprisingly, these plants are almost entirely in energy-intensive industries. The aggregate nature of the employment data used in the paper prevent estimating a separate effect for the manufacturing plants whose energy production was directly regulated. The need for a separate estimate is mitigated as energy production is regulated whether or not it occurs within the boundary of the firm.

[5]Of course there is also a large literature on the benefits of regulation and air quality improvements. See Chay and Greenstone (2003); Jerrett et al. (2009); Deschenes et al. (2012); Zivin and Neidell (2012) for a few representatitive examples.

in recent years. Performing the DDD analysis with a broad set of controls, I find that employment in industries with an additional percentage point of energy intensity decreased 1.38% in the region that was impacted by the NBP. Employment declines are shown to be largest among young workers. Using data on worker flows from the QWI, firms are shown to have reduced their employment levels primarily through a reduction in hiring rather than an increase in separations. Furthermore, wage offers, as measured by new hire earnings, are shown to have declined by as much as 3.5% for industries in the top quartile of energy intensity as compared to the bottom quartile. The ability to observe new hire earnings is important as firms may be unable to immeditely adjust earnings of incumbent workers. In order to evaluate the plausibility of these findings I then examine the NBP's impact on electricity prices. While not conclusive, these results suggest a rise in electricity prices within the range of ex-ante predictions. I then back out an implied manufacturing employment electricity price elasticity, which is in line with previous estimates (Kahn and Mansur 2010) and suggests a plausible causal mechanism for the impact on labor markets.

This paper adds to the literature in two important ways. It is the first paper to empirically estimate the impact of any EPA cap-and-trade program on labor markets. Given the size of the cap-and-trade programs, as well as the current policy debate over additional energy sector regulations, a better understanding of their impact is greatly needed. Second, it provides evidence of which workers were affected; uses worker and job flows to examine how the employment adjustments occurred; and examines the impact on worker earnings by focusing on the earnings of new hires, the margin on which earnings changes will most quickly adjust.

The remainder of the paper is organized as follows. Section 2 presents a brief history of the Clean Air Act Amendments and how the NBP came to be implemented. Section 3 describes conditions required for identification and Section 4 details important aspects of the data used in the analysis. Section 5 provides the econometric model, results and specification checks. Section 6 discusses the results. Section 7 performs a plausibility check and Section 8 concludes.

2 Background

Originally passed in 1963, the Clean Air Act (CAA) is the main federal law that seeks to control air pollution throughout the United States. The CAA has been amended multiple times including 1966, 1970, 1977 and 1990. Perhaps the most researched of the regulations brought upon by CAA and its amendments is the NAAQS. The NAAQS were established following the 1970 CAAAs and required polluting establishments located in counties that failed to achieve certain air quality levels to meet stricter emissions standards than establishments located in counties whose air quality was deemed acceptable. These emission regulations were by far the most important federal emissions regulations to date.

The 1977 amendments, in addition to strengthening the NAAQS, included Section 126, a provision that allowed the EPA to regulate interstate air pollution and limit the environmental harm that downwind states could impose on upwind states. The EPA did not immediately enforce Section 126, however, choosing instead to focus regulation efforts on establishments whose pollutants impacted the air quality of their local community rather than those impacting regions outside their immediate geographic region. In fact, between 1977 and 1998, the EPA never granted a petition filed under the interstate air pollution clause found in Section 126 of the CAA.

The passage of the 1990 Clean Air Act Amendments strengthened the language of Section 126 and established the first cap-and-trade programs. It was passed in response to the continued failure of many northeastern regions to meet air quality requirements despite having already restricted emissions in their local region. Title IV of the 1990 CAAA established a cap-and-trade program for SO_2. This would become known as the Acid Rain Program and in 1995 the EPA began Phase I for the dirtiest 110 power plants.

In 1998 the EPA granted its first petition under Section 126, paving the way for the NBP cap-and-trade of NO_x, an important precursor of ground-level ozone. The granting of this petition came as the result of two factors. First, the 1990 amendments had strengthened the interstate pollution protection law, calling for "reasonably available control technology" throughout an ozone transport region. Second, numerous lawsuits filed against the EPA by northeastern states requested that the EPA regulate NO_x emissions from states whose emissions directly contributed to their own levels of smog and ozone. These upwind states argued that NO_x pollution from downwind sources had not only negative health impacts on their citizens but also prevented them from meeting the NAAQS ozone non-attainment standards. In these lawsuits, a large body of scientific evidence was presented showing that NO_x gases can in fact be transported significant distances by wind currents and that NO_x emissions should therefore be subject to Section 126 of the CAAA. By granting the petition of the northeastern states, the EPA agreed to regulate and reduce the amount of NO_x emitted by electric generating units (EGU's) and large industrial plants in southern and central states.[6]

The NBP cap-and-trade program formally began for eight states and the District of Columbia in 2003 (see Figure 1). States and utilities in the Midwest and Southeast continued to fight legal battles against the EPA with varying outcomes, but in 2004 eleven additional states began compliance with the NBP, for a total of nineteen states.[7] The program would regulate 2,250 EGU's and 350 large industrial units that produced energy and heat for their own consumption.

As can be seen in Figure 2, EGU's dramatically decreased their output of NO_x on May 31, 2004, the first day in which all nineteen states began participating. Regulated establishments could choose to reduce their NO_x emissions in a variety of ways. One option was fuel switching, whereby establishments would shift away from coal and towards alternative energy sources such as natural gas that release far less NO_x into the atmosphere. Despite the additional production costs brought on by the NBP, most coal remained a cheaper source of energy than the alternatives (Fowlie 2010). Because electricity production from coal fired plants remained inframarginal, utilities largely continued to burn coal and found alternative ways to comply with the NBP. Depending on the preference of the plant, the compliance costs could be fixed and upfront or they could be variable and spread out over time. Plants opting for the high upfront cost option installed selective catalytic reduction (SCR) technology. This technology cuts NO_x emissions by up to 90% but costs the average plant $40 million dollars (Linn 2008). On the other end of the spectrum, about 30% of NBP regulated establishments chose to make no capital adjustments and simply purchased permits for every unit of NO_x they emitted (EPA 2009).[8] Regardless of which reduction technique they choose, the production costs of electric utilities will increase as a result of the NBP. Three estimates have been made that calculate the cost of the NBP to utilities. Palmer et al. (2001) estimated the program's total costs to utilities at about $2.1 billion per year. Deschenes et al. (2012) use the market price of permits to estimate the cost at $400-700 million and Linn (2010) examines utility

[6]The smaller and less restrictive program known as the Ozone Transport Commission NO_x Budget Program (not to be confused with the NBP) began for 11 northeastern states in 1999.

[7]Through negotiations and court battles, Missouri delayed compliance until 2007. Georgia, originally slated to also begin in 2007, was eventually ruled exempt from the program altogether. Additionally, deals were struck in Missouri, Alabama and Michigan which limited compliance to only certain counties.

[8]SCR is both the most expensive and most effective technology in reducing NO_x emissions (Fowlie 2010). There are a variety of less expensive and less effective technologies that utilizes chose to install. Selective non-catalytic technologies cost the average plant $10 million but only reduces NO_x by 35 percent. Additional pre-combustion and combustion technologies can decrease emissions between 15 and 50 percent depending on the specifications of the plant. Fowlie shows that regulated utilities are more likely to pursue capital intensive solutions than deregulated utilities.

stock prices to estimate a drop in expected utility profits of up to $25 billion dollars. Palmer et al. (2001) argue that the costs of the NBP will be passed on to the consumer in the form of higher electricity prices. Indeed, the EPA1997 estimated that electricity prices would rise by 1.6% as a result of the NBP and a later report (*"The NOx Challenge"* 2003) by Platts Research and Consulting predicted a $1-$3/MWh increase in the price of wholesale electricity.[9]

Carlton (1983) and Kahn and Mansur (2010) document that electricity prices are a major determinant of where manufacturing firms choose to locate their workers. Given the NBP's substantial impact on electricity and energy production costs more generally, firms that require high levels of energy in their production process may be expected to reoptimize their input mixture in response to a change in the expected costs of a crucial input. Using a DDD approach, I test whether firms with high energy requirements did in fact respond to the NBP by decreasing employment levels after the implementation of the NBP, relative to low-energy firms and relative to non-NBP control areas.

Since it began in 2003 and 2004, the NBP has changed names but the market for NO_x allowances continues to exist. In 2008 the NBP became part of the Clean Air Interstate Rule (CAIR) and in 2011the EPA announced it would replace and expand the regulations of CAIR with the new Cross-State Air Pollution Rule (CSAPR).[10] These regulations continue to be greatly debated and on August 21, 2012 the D.C. Circuit Court of Appeals vacated CSAPR leaving the future of both the SO_2 and NO_x cap-and-trade programs in doubt.

3 Identifying NBP Employment Effects

In order for a DDD methodology to accurately capture the causal effect of the NBP on manufacturing employment there are certain identification assumptions that should hold. One such crucial assumption is that control groups are not impacted by the treat-

[9]Using average wholesale prices in the Northeast ISO and the PJM in 2003 this is equal to an electricity price increase of between 2.47% and 7.41%.

[10]In 2005 the EPA announced that the Clean Air Interstate Rule (CAIR) would replace the NBP's regulation of NO_x emissions in 2009 and SO_2 in 2010. CAIR was intended to expand the number of covered states to twenty-five and further tighten emissions standards. Given that there was significant legal uncertainty surrounding CAIR when it was announced, it is unlikely that manufacturing industries would have immediately reacted. Because CAIR is a continuation of the NBP cap-and-trade program it is difficult and perhaps even unnecessary to disentangle the impact of the CAIR announcement from the implementation of NBP. In short, the interpretation of the evidence is influenced only slightly by CAIR, with all results still attributable to the overall cap-and-trade policy.

ment. Identication rests on two sources of employment change variation.The first is the variation that occurs within a state across industries and the second is the variation that occurs within an industry across states. When considering the within-state variation, it is possible that workers leave high energy industries in the county and are hired by low energy industries in the county. These local labor market spillovers are a potential source of bias as this may result in increased employment levels for low energy industries in the NBP region. To check whether these spillovers are driving the results, I consider models where the identifying variation does not come from within-state differences. The second source of variation is that which occurs across states. If firms shift production from NBP to non-NBP states then estimates may overstate the effect.

Determining an appropriate start date is another important part of determining the treatment effect. This analysis assigns start dates as the dates when the NBP went into effect (2003 for eight states and 2004 for the eleven others). While the cap-and-trade program was first approved in 1998, there was significant uncertainty until March of 2000 when the D.C. Court of Appeals ruled in *Michigan et al. vs. EPA et al.* that the program was legal. Some states and utilities continued with lawsuits after the 2000 court decision and electric utilities did not face the full costs of the NBP until sometime later, either when they purchased the abatement technology in the form of new capital, or once the program began in the form of purchasing emissions credits.[11] As seen in Figure 2, electricity production itself was not affected until much closer to the beginning of the NBP. By using the NBP start date I make the assumption that manufacturing firms remained uncertain of how the program would affect their energy costs until after permit trading began and NO_x emitters faced the full costs of the program. If manufacturing firms began decreasing employment before this date then the results would be biased towards zero, as some of the treatment effect would be attributed to the pre-treatment period.

The timing of the policy should also be checked against other simultaneously occurring events that may impact manufacturing employment. Given the controls described above, in order for such an event to drive the results it would have to have a different impact on the NBP region than the non-NBP region and it would have to differentially impact industries based on their energy intensity levels. One potential event was a change in NAAQS nonattainment standards which caused 408 counties across the country to enter nonattainment status in 2004. These counties were disproportionately located in

[11]When examining the NBP's impact on expected future utility profits, Linn (2010) uses 2000 as the beginning date.

the NBP treated region, and non-attainment designation is likely to differentially impact industries based on their energy intensity. To control for this, I examine county-industry data and allow for high energy industries in new NAAQS non-attainment industries to experience a separate employment effect. My results are robust to controls for these potential NAAQS effects.

A second potential concern is that certain regions may be particularly sensitive to changes in fuel prices. If an increase in the price of oil raises energy prices uniformly across the country, this will be picked up with the industry-year fixed effects. However, certain regions rely heavily on one particular fuel source for their electricity production. To account for the possibility that regional electricity prices may be differentially impacted by changing relative fuel prices, I obtain average oil, natural gas and coal prices for the years 1998-2008 as well as the percent of electricity that is derived from that source in each North American Electric Reliability Corporation (NERC) region in the country. Interacting the fuel price with the percent of electricity derived from that fuel in the state's NERC region and the industry indicator variables allows for the fact that certain industries in certain regions may be particularly sensitive to a change in fuel prices. My results are insensitive to these controls as well.

All time-invariant differences are absorbed by full sets of state-industry (or county-industry) fixed effects. I use state, county and industry-specific linear time trends to control for pre-existing trends; in other models industry-year and state-year indicator variables flexibly account for nonlinear trends. Even after accounting for state and industry specific trends, results could still be driven by pre-existing trends whereby high energy industries in the NBP region are trending differently than high energy industries in the non-NBP region. For example, if energy intensive industries in the east and west have different employment trends as a result of the NAAQS or some other feature of earlier CAAA's that disproportionately impacted the NBP region, then these pre-existing trends could be mistakenly attributed to the NBP. To assure that separate pre-existing trends are not being picked up, I allow for each industry to trend differently based on the region in which its employment is located. If the implementation of the NBP causes employment levels and trends to differ by region, then adding these separate trends is likely to absorb some of the impact of the NBP. This again will result in a conservative estimate of the NBP's impact on employment. Thus, my results are identified off of very weak assumptions, which allow for state and industry-specific non-parametric trends and pre-existing east-west differences in industry-specific trends.

4 Data

4.1 County Business Patterns

The two employment datasets used to analyze the impact of the NBP on manufacturing employment are the Census Bureau's County Business Patterns (CBP) and the Quarterly Workforce Indicators. The CBP is a yearly data product released by the Census Bureau that provides sub-national economic data by industry. Data can be obtained at the national, state, county and metropolitan levels and include the total number of establishments and workers by industry in a geographic area. The source of the CBP is the Business Register, Census' Company Organization Survey and other economic censuses and surveys such as the Census of Manufactures and the Annual Survey of Manufactures. Using CBP data from 1998-2008, I create panel data sets at both the state-industry and county-industry level.

In 1998 the Census Bureau switched its industry classifier variable to account for the changing face of the American economy. The Standard Industrial Classification (SIC) system was abandoned in favor of the newer North American Industry Classification System (NAICS). The change from SIC codes to NAICS codes in 1998 creates some difficulties in consistently estimating industry employment across time periods. Given that the NBP was implemented in 2003, it is logical to use 1998 as a starting year for the data and avoid any inconsistencies that may arise from merging previous years with different industry definitions. All data between 1998 and 2008 use NAICS codes which are consistent across time periods. Following previous literature, this paper uses three-digit NAICS codes as the industry level of observation (Greenstone 2002; Kahn and Mansur 2010).

While the CBP has the distinct advantage of being publicly available, it also has the disadvantage of having to undergo a thorough review process to prevent the release of any data that would disclose the exact records of any single establishment. Therefore, if very few establishments are located in a particular county or state-industry, then employment data will be suppressed for that observation. The primary results of this paper use state-industry data which has limited cell suppression for employment. In the state-industry dataset 76% of state-industry cells are observed directly. These cells represents 93% of all manufacturing employment in the United States. For those cells that are suppressed, I perform an imputation method similar to that used by Kahn and Mansur (2010) and Mian and Sufi (2012), which takes advantage of the CBP's establishment-size cell count variables and imputes employment for the suppressed cells by multiplying the number of establishments in each establishment-size cell by the midpoint establishment size of that category.[12] The same imputation method is used for the county-industry level data used in the robustness checks.[13]

4.2 Quarterly Workforce Indicators

Like the CBP, the QWI is a publicly available dataset that contains sub-national employment data by industry. The underlying microdata for the QWI is the Longitudinal Employer Household Dynamics (LEHD) program at the U.S. Census Bureau, which uses state unemployment insurance data as its primary input (see Abowd et al. (2006) for a complete description of the QWI and the LEHD). In recent years a number of papers have begun to use this data to evaluate the labor market impacts of the housing crash, changes to minimum wage laws and workplace mandates (Abowd and Vilhuber 2012; Gittings and Schmutte 2012; Dube et al. 2011; Curtis et al. 2013).[14] The QWI has both strengths and weaknesses compared to the CBP but there are two primary reasons to use QWI data. First, it contains detailed cuts of the data by worker characteristic. That is, the QWI provides not only total employment within a state-industry, but also breaks down this employment by age group and gender. This data can then be used to understand the heterogeneity of the treatment effect along a number of dimensions that are not available in the CBP. The second reason is that the QWI provides data not only on employment levels but also on worker flows (hires and separations) and job flows (creations and destructions). Finally, the quarterly nature of the data also provides more frequent snapshots of employment variables and thus gives a better feel for the dynamics at play and the impact of the program over time.

The disadvantages of the QWI lie in its coverage and its data suppression. Most states have now agreed to share UI data with the LEHD system but the historical data they provide differs by state. As a result, I use data from the 40 states whose data goes back until at least 2000 so that there is a reasonable pre-treatment period for each state. Of additional concern is the lack of information on suppressed cells. When cells

[12] All state and county-industry observations contain the number of establishments in narrowly defined employee size categories (1-4, 5-9, 10-19, . . . , 5,000+). For the 5,000+ category employment is top coded at 6,000. See (Kahn and Mansur 2010) for a full explanation of the imputation method.

[13] As discussed later in the paper, regressions weight state-industry observations by their pre-NBP employment level. This is common in the literature and mitigates concerns about imputation related bias as imputed cells are smaller and thus given less weight in the regressions.

[14] Data on job flows has been available for longer. See Davis and Haltiwanger (2001) for an example of research examing job flows.

are suppressed in the CBP they can be imputed using the employee size categories, which are available for every observation. In the QWI there are no additional variables which allow for the imputation of suppressed cells. Because of these disadvantages, the benchmark and primary employment specifications use the CBP. Importantly, results using the QWI not only confirm findings using the CBP but are able to paint a more complete picture of how the NBP impacted labor markets. See the data appendix for additional details on both the QWI and the CBP.

4.3 NBER Productivity Database

After obtaining annual (or quarterly) state-industry labor data, I merge in three-digit-industry energy intensity data from the 1998 NBER Productivity Database. This database contains total energy expenditure by industry in the given year and is based off of the Census of Manufactures and the Annual Survey of Manufactures. To construct an energy intensity index for the 21 different 3 digit manufacturing industries, I divide total industry energy expenditure by total value of shipments for the industry.[15] As seen in Table 1, energy intensity in the manufacturing sector varies from a low of 0.6% in the computer and electronic product industry to a high of 5.5% in the primary metal manufacturing industry.

5 Econometric Model and Results

In order to motivate the econometric analysis, provide summary statistics and preview the results, it is informative to begin by viewing the raw employment data between 1998 and 2008. Based on the energy-intensity index in Table 1, I split the 21 industries into three separate groups. The seven industries with the highest energy intensity measures are defined as "high intensity industries", the seven industries with the lowest energy intensity measures are defined as "low intensity industries" and the middle seven are defined as "medium intensity industries".

Figure 3 plots out the east-west employment difference for each industry energy intensity grouping from 1998-2008. Specifically, the figure plots the percentage change in employment in the east minus the percentage change in employment in the west for each industry grouping, using 1998 as the baseline year. This plot suggests a potential

[15] Because the NBP regulated NOx emissions from heat, steam and electricity production I use energy intensity rather than electricity intensity.

effect of the NBP on employment, but it also reveals that pre-existing trends may be present that, if unaccounted for, could bias DDD estimates. The plot shows that between 1998 and 2008 the east-west employment difference falls most prominently for the high intensity industries while there is little change in the east-west difference for the low-energy industries. Employment in medium intensity industries tracks closely with high intensity for the first four years, but starting in 2003, the east-west employment difference begins to fall for high energy industries relative to medium energy industries. This gap widens substantially between 2004 and 2008. Vertical lines are drawn in 2001, the year after the courts determined the NBP's legalitiy, and in 2003, the start date of the program. Examining and accounting for any pre-existing trends will be important to the internal validity of the DDD estimates.

5.1 Baseline Regression Specification

The identification strategy of this paper takes advantage of the geographic, time and industry heterogeneity found in the data in order to estimate the impact of the NBP on manufacturing employment outcomes. As a first step towards exploiting this heterogeneity, I consider the following DDD model:

$$\begin{aligned} y_{gkt} = \beta_T(Post_{gt} \times East_g \times EnInt_k) + \beta_{pe}(Post_{gt} \times East_g) + \\ + \beta_{pen}(Post_{gt} \times EnInt_k) + \theta x_{gkt} + \delta_{gk} + \xi_t + \epsilon_{gkt} \end{aligned} \tag{1}$$

In this model, y_{gkt} is the employment outcome of interest (logged employment, hiring rate, separation rate, etc.) in geographic region g, in industry k in period t. All findings listed in the following section maintain the same definitions for the variables $Post_{gt}$, $East_g$ and $EnInt_k$. $Post$ is an indicator variable equal to one if the date is after the start of the NBP.[16] The NBP began in May of 2003 in eight states and in May of 2004 for eleven states. For the eight states beginning in 2003, $Post$ equals one for years 2004 and later. For all other states $Post$ equals one for years 2005 and later.[17] The baseline specification

[16]See Figure 1 for a map of the NBP region and the year which each state began compliance. Missouri is dropped because it did not begin until 2007 and only certain Missouri counties were required to comply.

[17]Results using QWI assign the start time likewise as the quarter following the NBP. Linn (2010) performs a similar analysis in which he estimates the impact of the NBP on electric utility profits. His empirical work, which examines the impact on utility stock prices, uses 2000, the year the Federal Court of Appeals confirmed the NBP would be implemented the first date in which the policy was known to be occurring with certainty. As can be seen in Figure 2, electricity production itself was not altered until the

also uses a broad definition of the variable *East*. The definition of the treated region is informed by the mechanism through which a change in employment would occur. I set *East* equal to one for all states whose electric utility provider is impacted by the NBP. This expands the treated region to include areas that were not directly regulated by the NBP. The results are shown to be insensitive to alternate specifications that restrict the treated region to only those states and counties that were directly regulated. Figure 4 provides a map of the treated region and provides additional details. Finally, the variable $EnInt_k$ is a time-invariant measure of the industry's energy intensity as defined by total energy expenditure divided by total value of shipments for the entire industry in 1998. The primary employment specifications will be at the state level such that gkt is state-industry-year data. Robustness checks will use county level data where a unit of observation is at the county-industry-year level and g refers to county rather than state.

The main coefficient of interest is β_t, which captures the triple interaction of an observation being in the treatment group, after the treatment has been applied and allowing for differences by industry based on their energy-intensity level. The variables δ_{gk} and γ_t represent full sets of state-industry and year indicator variables in order to control for time-invariant differences across state-industries and any shock that occurs in a given year and is common to all manufacturing industries and all states. A vector of variables, represented by x_{gkt} is included in the robustness checks to ensure the results are not being driven by omitted variables. This specification does not capture secular state and industry trends that are unrelated to the NBP but are likely a driving source of the employment change within a state-industry.[18]

These and other concerns are addressed shortly, but before discussing the results and additional specifications, a few important details and assumptions bear mention. First, as is common in the literature, observations are weighted by their pre-treatment 1998 employment levels to ensure that state-industries with little or no employment do not drive the results (Greenstone 2002; Walker 2011). Second, while using aggregate state-industry data reduces the computational burden and accounts for *some* of the inference concerns raised by (Bertrand et al. 2004), given that the indentifying variation occurs at a level higher geographic level than the state, it is crucial to account for serial and spatial correlation of the error term to avoid understating the size of the standard errors. The standard errors are clustered at the NBP region-Industry level to address these

NBP was implemented. For the purposes of this analysis I assume that manufacturers did not react to the policy until electricity production was actually altered and the price of permits had been established.

[18]Note that the variable $East_g \times EnInt_k$ is absorbed by higher order fixed effects.

concerns.[19]

A final feature to note is that the model assumes a linear effect in energy intensity. When logged employment is the outcome variable, the triple difference coefficient should be interpreted as the percentage change in employment that occurs for every additional percentage point in energy intensity. While there are strong theoretical reasons to believe that the most energy intensive industries will experience the greatest impact from the NBP, deciding exactly how to model this differential impact is not immediately straightforward. Imposing a linear assumption allows for results to be obtained in a single DDD coefficient which is easily interpretable for the manufacturing sector as whole. However, I provide additional results which calculate the impact of the NBP separately on each of the 21 industries. This is done by replacing the triple interaction variable with 21 separate industry-specific triple interaction terms. Results from this non-parametric approach, presented later, reveal that the linearity assumption is not unreasonable.

5.2 Employment Results Using County Business Patterns

The main employment results of the paper are found in Table 2 and use employment data from the CBP, which contains data for every state-industry-year cell between 1998 and 2008. Column 1 reports estimates from the base model with additional controls being added in each subsequent column to account for potential confounding factors. The preferred specification is reported in column 6 and includes state trends, industry trends as well as separate industry by region trends. Tables 3 and 4 report a variety of robustness checks using a similar table layout but alternate samples and control variables. Each column represents a separate regression where the dependent variable is logged employment plus one. In each table the coefficient on the primary triple difference variable $Post_{gt} \times East_g \times EnInt_k$ is reported as well as the coefficient on the $Post_{gt} \times East_g$ variable for those specifications in which it is not absorbed by higher order fixed effects. The coefficient on the triple difference variable signifies the percentage change in employment that occurred for industries with an additional percentage point of energy intensity after the policy was enacted, in the states to which it applied. The coefficient on the variable $Post_{gt} \times East_g$ represents any change that occurred to all manufacturing employment in the NBP region relative to the non-NBP region apart from the differen-

[19] In separate results not presented here, standard errors were clustered at a variety of other levels including state-industry, state and NBP region. Clustering at the NBP region-Industry level accounts for spatial and serial correlation within an indstry over time and proved the most conservative of all clustering methods.

tial impact by energy intensity. This coefficient, when it is not absorbed by State-Year fixed effects, is close to zero in all of the specifications, thus supporting the case that the primary mechanism through which the policy impacted employment was through a heterogeneous treatment effect that varied by the industry's energy intensity.

The results of the baseline model, listed in the column 1 of Table 2, show a large but statistically insignificant impact of the NBP on manufacturing employment in higher energy industries. The imprecise nature of these estimates is unsurprising given that much of the employment change within a state industry is likely to be driven by exogenous state and industry trends which are not controlled for in this base model.

Because state and industry trends are likely to be important, I consider two methods of controlling for this source of variation. The first takes a non-parametric approach by including a full set of industry-year and state-year indicator variables. The inclusion of industry-year indicator variables accounts for any industry-specific shock that is common to all states in a given year while the inclusion of state-year indicator variables accounts for any shock that is common to all manufacturing employment in a state in a given year. The non-parametric specification provides the model with a high degree of flexibility but the large set of indicator variables, particularly the set of state-year dummies, are quite demanding of the data as they dramatically reduces the degrees of freedom in the regression.[20] Results using the fully non-parametric approach with sets of state-year and industry-year indicator variables are given in column 2 of Table 2. The triple difference coefficient is negative and statistically significant, implying that employment in high energy industries fell relative to low energy industries and relative to other high energy industries in non-treated states.

If we are willing to relax these non-parametric trend assumptions we may continue to account for state specific time trends by including a full set of state indicator variables that have been interacted with a linear time trend variable. While not as flexible as the non-parametric approach, these state-specific linear trends are likely to capture a substantial portion of the year to year change as employment levels rarely make discontinuous jumps within a state from one year to the next. Column 3 reports results using these state specific linear trends. As expected, assuming linear state trends does not impact the magnitude of the triple difference coefficient but does shrink the standard errors, thus providing a slightly tighter confidence interval. These smaller standard errors

[20]The dataset used in the main specification contains 11,319 observations. Including state-year indicator variables adds 539 additional variables to the regression. Using linear instead of non-parametric state trends reduces the number of state-trend variables from 539 to 49.

will prove useful in upcoming specifications.

Columns 4-6 repeat the specifications used in columns 1-3 with one exception. Columns 4-6 now include separate region-by-industry trends to account the possibility that, for example, high energy industries in the east may be trending differently than their counterparts in the west before the start of the program. If high energy industries in the east are trending down faster than high energy industries in the west then failing to account for these trends will overstate the impact of the policy. While including these trends is essential to accurately estimating the impact of the policy, it is also important to note that if the NBP lead high energy industries to experience a change in employment levels as well as a change in trends, then including these trends variables may pick up some of the impact of the program thus masking its full effect.

The size of the estimates in Columns 4-6 are lower than their counterparts in columns 1-3 and suggest that accounting for these separate region by industry trends is important. The new coefficients, now ranging from -1.46 to -1.16 imply a smaller but still sizable impact of the NBP. The coefficient in column 5, which includes the full set of 539 state-year indicator variables is no longer statistically different from zero. Including this set of fixed effects on top of the 1,029 State-Industry fixed effects is quite demanding of the data. If we are willing to replace these state-year fixed effects with a more parsimonious set of forty-nine state-specific trend variables then the coefficient becomes statistically significant with little impact on its magnitude. This model, written out fully in equation (2) below represents the preferred specification of the paper.

$$\begin{gathered} y_{gkt} = \beta_T(Post_{gt} \times East_g \times EnInt_k) + \theta x_{gkt} + \delta_{gk} + \alpha_{kt} + \\ + \sum_{g=1}^{G} \beta^g_{trend}[trend_t \times 1(State_s = g)] + \\ + \sum_{k=1}^{K} \beta^k_e[trend_t \times East_g \times 1(Ind_i = k)] + \sum_{k=1}^{K} \beta^k_w[trend_t \times West_g \times 1(Ind_i = k)] + \epsilon_{gkt} \end{gathered} \tag{2}$$

As described above, this model controls for industry trends non-parametrically by including α_{kt}, a set a industry-year fixed effects. It accounts for state trends through the use of forty-nine state-specific linear trend variables as represented in the first summation term. And finally, the last two summation terms represent separate region by industry trends, which allow for industry employment to trend differently based on the region in which it is located. The results from this model, are listed in column 6 of Table 2. The coefficient is -1.38 and is statistically significant at the 5 percent level. This implies that for each additional point of energy intensity, an industry's employment level declined by 1.38 percent.

Interpreting the results requires careful consideration of the assumptions at hand and the identifying variation used in the analysis. The coefficient on the triple difference variable represents the percent employment change that occurred in an industry for every additional percentage point of that industry's energy intensity level. One straightforward way to calculate the overall employment loss due to the policy is to take the pre-NBP employment levels in each industry, multiply them times their energy intensity level and then times the triple difference coefficient. This calculation suggests the transition or loss of 192,000 manufacturing jobs in the NBP region as a result of the program. While it is interesting to calculate this simple, back-of-the-envelope estimate it is, for a number of reasons, likely to give an upper bound for the overall employment effect for two reasons. First, as previously discussed, some workers who leave high intensity industries and NBP states will find employment in low energy industries and non-NBP states. Thus, if the intention is to calculate overall declines in employment in the NBP region, then this may overestimate as the number of jobs, particularly in the long run once separating workers have found new employment. Section 5.4 explores this possibility by examining worker flows and periods of non-employment for separating workers.

The second reason is specific to the assumptions made by the triple difference model. With the exception of column 5, which includes state-year fixed effects, all specifications report the coefficient on the *PostxEast* variable. This coefficient captures any overall shift that may have occurred to all manufacturing employment in the NBP region relative to the non-NBP region after the policy went into effect. Deciding whether or not this coefficient should be included as part of overall employment effect depends on the assumptions the reader is willing to make. While it is possible that an overall shift occurred due to the NBP, there are reasons to think this was not the case. First, the coefficient on the postxeast variable is small, often positive and statistically insignificant in all specifications. That this coefficient should come close to zero is not surprising given the large set of industry and state trends in the specification. If the positive coefficient were to be interpreted as part of the NBP's employment effect, this would suggest a smaller impact as the differential employment loss associated with each industry's energy intensity would be slightly offset by an overall rise in all manufacturing employment. Given that there are other factors likely to be driving an overall shift in manufacturing that are unrelated to the NBP, a more reasonable interpretation will allow for overall shifts to occur and limit the impact of the NBP to be only the employment change that varies by energy intensity. This allows for the identifying variation to come from differential shifts that occurred across industries based on the industries' energy intensity level. Any event that impacts all manufacturing employment in the east relative to the west will be controlled for in the results.

Given that the NBP did not result in an overall shift in all manufacturing employment, a decision must still be made regarding the appropriate baseline against which the triple difference coefficient should be judged. The back of the envelope method used above assumes that an industry with zero energy intensity experienced zero employment change. Because all industries have energy intensity greater than zero, this would imply that all industries experience some employment loss due to the NBP. A more conservative interpretation would compare the overall employment loss in the most energy intensive industry with the loss in the least intensive industry. Under this method, we see that the NBP caused the loss of 5.8% of employment in the most energy intensive industry relative to the least intensive industry. More conservative yet, the triple difference coefficient implies that the average employment loss of an industry in the top energy intensive quartile lost 4.4% of employment relative to the average industry in the bottom quartile of the energy intensive measure.

5.3 Alternative Employment Specifications

To examine the robustness of the result in the preferred specification, I consider models which: (1) include additional controls in the regression model; (2) more narrowly define the treated region; (3) examine county-level rather than state-level data and (4) relax the assumption of a linear employment effect in energy intensity. Panel A of Table 3 reports results from identical models as Table 2 but controls for exogenous changes in fuel prices that may impact the energy costs of certain regions more than others. Any shock in energy prices that is common to the entire country will be picked up in the industry-year fixed effects but if, for example, there is a shock to the price of oil and certain regions heavily rely on oil for electricity, then it is possible that these regions will see an increase in the cost of energy that is unrelated to the NBP.[21] To ensure that shocks

[21] Annual oil, natural gas and coal prices are the Brent Price, the Henry Hub Natural Gas spot price and the EIA total average coal price and were downloaded from `http://www.eia.gov/totalenergy/data/annual/showtext.cfm?t=ptb0709`. NERC region resource mix data comes from EPA's eGRID summary tables. Georgraphic boundaries of NERC regions do not always correspond to state boundary lines. For those states which belong to more than one NERC, I assign them a fuel-intensity level equal to a weighted sum of the fuel-intensity level of the NERCs to which they belong where the weight is the percent of

to oil, natural gas and coal prices are not driving the results I gather data on annual fuel prices for each of these fuels and interact these prices with the percent of electricity that is derived from that fuel in the NERC region to which that state belongs. This variable is then interacted with the energy intensity variable to control for the fact that high energy industries in certain regions may be particularly sensitive to certain fuel price shocks. The results are shown to be insensitive to these controls.

Panel B of Table 3 continues to use state level data but uses an alternate definition for the treated region. Rather than use the broad definition described in Figure 4, this specification defines the NBP region as only those states that were directly impacted by the program. Using this narrower definition, North Dakota, Minnesota, Iowa, Wisconsin, Maine, New Hampshire and Vermont are dropped from the sample. Results using this new definition find the coefficient on the triple difference variable to be -1.57, thus suggesting a slightly greater impact of the NBP than was found in the baseline specification.

Table 4 provides new estimates of the NBP's employment impact using county rather than state level data. The use of county level data comes with both benefits and costs. The primary reason why county data may be useful is that it allows the specifications to control for county level changes to the NAAQS non-attainment standards that occurred to 408 counties in 2004, the first year that the NBP was fully implemented. These counties were located across the United States but a disproportionate number were located east of the Mississippi. Using county data also allows for additional robustness checks that allow the NBP treatment group to be defined along county rather than state lines. While the NBP was a state level program, certain counties in both Alabama and Michigan were ruled exempt of the regulation and can be excluded from the treated region when using county level data.[22] The major downside of the county level data is the increased prevalence of data suppression. Due to disclosure concerns, employment for 51% of all observations is suppressed. This represents 72% of all manufacturing employment. When suppressed, employment is imputed by assigning the midpoint of the employment range provided in the data. This imputation process reduces the true employment variation in the data and raises other concerns regarding when cells may fall into and out of suppression status. These concerns are mitigated with the use of state data which is suppressed far less often Panel A of Table 4 reports results using a

manufacturing employment in the state that lies in that NERC region.

[22]As discussed earlier these counties were likely still impacted by the policy as their utility providers produced electricity in NBP regions.

nearly identical set of regression specifications in Table 2 but at the county rather than state level. County-industry fixed effects control for any time invariant differences in employment levels and industry, state and industry by region trends are controlled for in an analogous manner. Reassuringly, Panel A reveals the county level results to be of similar magnitude to their state level counterparts. Results are no longer statistically significant for columns 2 and 3 but are statistically significant at the 5% level in column 6, the preferred specification which includes separate region by industry trends.

Panel B of Table 4 restricts the treated region to only the counties that were directly impacted by the NBP. The results for the preferred specification suggest a slightly smaller impact than suggested by the state level counterpart, but the coefficient falls well within the confidence interval of the state level results. The final important use of the county level data is to control for changes in county level NAAQS non-attainment standards. To isolate the impact of the NBP from any impact of the NAAQS, I include a term in the econometric model which interacts 2004 county NAAQS nonattainment status with the *Post* and *EnergyIntensity* variables.[23]

As a final specification, I return to the state-level data and consider the effect that the NBP had separately on each specific manufacturing industry. To do this, the main triple difference interaction variable is replaced with 21 different industry-specific interaction variables.[24] The coefficients on the industry specific interaction variables represent the estimated effect of the NBP on each industry. Figure 5 plots each of the resulting coefficients on the y-axis and the industry's energy intensity on the x-axis. As expected, the higher the energy intensity of the industry, the greater it was impacted by the NBP. There are no clear outlier industries that would be driving the results in the benchmark specification. Low energy industries generally experienced no clear impact while high energy industries experienced a negative impact. The apparent linearity of the relationship between an industry's energy intensity and their NBP related employment change supports the identification strategy used in the triple difference models.

[23]Specifically, a NAAQS non-attainment indicator variable is set equal to one for all counties that enter non-attainment for any criterion pollutant in 2004. This variable is then interacted with the *PostxEast* variable and included in the model. Panel C shows that the results in the preferred specification are robust to including these controls.

[24]That is, I drop the $Post_{st} \times East_s \times EnInt_k$ variable and replace it with $Post_{st} \times East_s \times Ind1_k$, $Post_{st} \times East_s \times Ind2_k, \ldots, Post_{st} \times East_s \times Ind21_k$ where $Post_{st} \times East_s$ has been separately interacted with each of the twenty-one industry indicator variables ($Ind1_k, \ldots, Ind21_k$).

5.4 Worker Flows and Earnings by Age Group: Results from the Quarterly Workforce Indicator Data

While employment is one outcome of interest, estimating changes in industry employment levels does not fully capture the impact of regulation on workers. Of great interest is how these employment shifts occurred, who was impacted and whether workers experienced a drop in earnings. These are all crucial, policy-relevant outcomes that are not captured by simply measuring changes in industry employment levels. To capture these important outcomes, I turn to data from the QWI. As discussed in Section 4 and the data appendix, the QWI has advantages and disadvantage in comparsion to the CBP. As a validity check, it is useful to begin by replicating the primary employment regression descriced in equation (2) using the QWI data. Column of Table 5 shows the DDD estimate using QWI data to be smaller (-0.871 as compared to -1.38) but it remains negative and statistically significant at the 5% level. That these two estimates differ is not altogether surprising given that the QWI contains only forty states.

Columns 2 through 5 of Table 5 estimate the impact of the regulation on job flows, using a similar specification to that described in equation (2). The hiring rate, defined as the number of new quarterly hires divided by the total employment, is shown to have declined for energy intensive industries. The coefficient of -0.225 implies that for every additional percentage point in energy intensity, an industry decreased their hiring rate by .225 percentage points. The average hiring rate for all industries over this time period was 6.9 per 100 employees. This coefficient would imply that there were 1.24 fewer hires per 100 workers in the most energy intensive industry. The programs impact on separation rates is not precisely estimated. Coefficients points negative but are statistically insignificant.

Regression results are also reported for job creation and destruction rates. Job creation is defined as the employment increase at expanding establishments while destruction is defined as the employment decrease at contracting establishments. Previous work by Davis and Haltiwanger (2001) has shown destruction rates increase in response to energy price increases. Walker (2011) finds that regulation both increases destruction and decreases creation. While the coefficient estimates in columns 4 and 5 are imprecisely measured, it is of interest that they both point negative.

Table 6 reports the major findings for each of the key variables in the QWI by age group. Each cell reports the results of a separate regression using a different outcome-population combination. All regressions use the main specification listed in equation (2).

The first row of Table 6 uses the entire population of workers. The first three columns of this row are therefore identical to the first three columns of Table 5. Columns 4 through 6 now report results examining the impact of the program on three different outcome variables: Periods of non-employment for separating workers, average earnings of new hires and average earnings of all workers. The QWI calculates periods of non-employment for separators calculating the average length of time separated workers are not observed in UI data following their separations. Each separated worker is defined as having a quarter of non-emoloyment if they are not observed working for any firm in the quarter following their separation. Each separated worker is tracked for four quarters following the separation. This variable, while somewhat crude, provide a sense of how workers fare after separating from their jobs. For example, over the recession, average quarters of non-employment for separating workers jumped from 1.6 to 2.1. If workers who separated from regulated industries see no increase in periods of non-employment, this suggests they may have shifted quickly to other jobs. An increase in this variable, particularly for middle aged workers, suggest regulation may have forced people into unemployment. See the Data Appendix for more details on this variable.

Regression results are unable to capture any impact on periods of non-employment. The coefficient points negative for all workers, is positive for young workers and turns negative for older workers. While it is of interest that there is no clear increase in periods of non-employment for separating workers in regulated industries, given the noisiness of this variable no meaningful conclussions can be drawn.

Earnings are also crucially important to an understanding of how regulations impact labor markets. Columns 5 and 6 report results using the same mthodology to examine the regulation's impact on earnings and specifically new hire earnings. As noted by Curtis et al. (2013), the effects of labor demand shocks will be more evident in new hire flows rather than the overall wage levels which will dominated by incumbents. The QWI provides data not only on average earnings of all workers, but also on the average earnings of all new hires. Due to existing wage contracts, firms that experience negative shocks, such as an increase in energy costs, may be unable to reduce the wages of incumbent workers. New hire earnings provide a margin over which firms are most likely to be able to adjust. Regression results reported in columns 5 and 6 support this claim. While the NBP appears to have no distinguishable effect on earnings of all workers, there is a statistically significant decline in the earnings of new hires. This decline is common across age categories with a coefficient of -1.277 for all workers. This implies a decline in new hire earnings of 4.2% for industries in the top quartile of the energy intensity index compared with those in the bottom quartile. No decline is observed when examing average earnings of all workers which consists primarily of incumbent workers whose contracts were negotiated before the regulation.

The final important aspect of these findings is the break down of the regulation's effect on each age grouping. Column 1 shows that the young workers, particularly those in the 19-21 and 22-24 age groupings, experienced the largest percent employment declines. Losses are less severe and statistically insignificant for older workers in the 45-54 and 55-64 categories but increase again for workers 65 and over. Separation rates decline for the youngest two categories, none of the age groups experience an increase in periods of non-employment following a separation and the earnings of new hires are shown to decline for all ages. While the 95% confidence intervals overlap for the employment age group coefficients, they suggest a heterogeneous impact along the age dimension.

5.5 Event-Time Models

To understand how the effects of the policy evolved over time it is useful to perform an event-time study. In the context of the above models, an event-time study tracks how the coefficient on the variable $East_g \times EnInt_k$ changes throughout the study period. I use the second quarter of 2003 to be the start date and normalize the coefficient for this period to be zero. Specifically I estimate the following model:

$$\begin{aligned} y_{gkt} = & \sum_{t=1}^{44} \beta_T^t [1(Qtr_q = t) \times East_g \times EnInt_k)] + \beta_{pen}(Post_{st} \times EnInt_k) + \\ & + \sum_{g=1}^{G} \beta_{trend}^g [trend_t \times 1(State_s = g)] + \theta x_{gkt} + \delta_{gk} + \alpha_{kt} + \epsilon_{gkt} \end{aligned} \tag{3}$$

This model mirrors that of equation (2) but removes the industry specific east-west trends and replaces the triple interaction variable with forty-four event-time coefficients. This model allows us to view trends both before and after the implementation of the NBP. Figure 6 plots he event-time coefficents on the logged employment model. The coefficients before the policy are slightly positive but none are statistically different from zero. The coefficients become negative after the NBP's implementation with zero falling outside of their confidence intervals.

As a comparison, Figure 7 reports the same coefficients where the left hand side variable is periods of non-employment for separating workers. The pattern here is quite distinct from that seen in Figure 6. Here, the coefficients consistently hover around zero and are not statistically significant before or after the NBP. The NBP appears to have had no noticeable impact on periods of non-employment for separating workers. This anaysis is a particularly important compliment to the the DDD model for this variable. DDD models takes the average of the pre-period and compares it to the average outcome in the post period. This model may not uncover an impact of the policy if the impact is short lived. By contrast event-time studies capture quarter-specific policy effects. The plot in Figure 7 shows that there was no clear positive spike in quarters on non-employment for separating workers in any of the quarters following the NBP.

6 Discussion

Taken as a whole, the results paint a picture of how labor markets in energy intensive industries respond to energy sector regulation. The findings show that employment in industries in the top quartile declined 4.4% compared to manufacturing employment in the bottom quartile. This employment decline provides evidence of a labor demand shock, but there are important, policy relevant measures that the employment decline does not capture. To gain a more complete picture of the labor market impact, it is useful to look at how the declines, occurred, which workers were impacted and how wages changed in response to the policy.

Results suggest that the employment drop occurred primarily through a decrease in hiring rates rather than an increase in separation rates. It is unsurprising then that young workers, a group with high turnover, experienced the largest employment declines as a result of the policy. These workers are likely to have less firm and industry-specific capital and have high turnover rates.[25] While young workers experience a decline in separation rates, they continue to separate at higher rates than older workers. The decline in firm hiring rates thus results in these workers not being replaced.[26]

Separating workers, regardless of their age, did not experience increased periods of non-employment following their regulation. While more research and better data is needed, this, together with the decline in the hiring rate suggests that firms may have

[25] For example, quarterly turnover rates, defined as $(hires_q + separations_q)/(employment_q \times 2)$ were roughly 0.13 compared to 0.05 for workers 45-54.

[26] Jacobson et al. (1993) and Walker (2012) both find that young workers experience faster earnings recovery following job displacement than older workers.

reduced employment through natural separations, such as job-to-job transitions rather than layoffs. Given the presence of labor adjustment costs, such as severance packages and increased unemployment insurance taxes, it is not surprising that firms will choose to adjust to small demand shocks primarily through natural separations.[27] Distinguishing between layoffs and natural separations (job-to-job transitions and retirements) is important. A variety of studies have found that workers who experience mass displacement events experience extended periods of unemployment, large earnings losses and increased mortality rates (Sullivan and von Wachter 2009; Davis and von Wachter 2011; Farber 2011).

Earnings results show that firms reduced their wage offers to new workers but that average earnings of all workers experienced little to no change. This, together with the worker flow findings, suggest that incumbent workers do not bear the brunt of regulation's impact. The effects are felt by future workers who receive lower wages, and potential future workers, who are no longer hired due to the regulation. These effects are important to labor markets and the economy but are quite distinct from the traditional job loss story that portrays long-time incumbent workers as the primary losers from environmental regulation.

7 Plausibility Check: Electricity Results

In order to evaluate the plausibility of the employment loss figures, it is useful to revisit the causal mechanisms that connect the NBP to declines in manufacturing employment. First, large manufacturing plants that were directly regulated by the program may have made employment adjustments on their intensive margin in response to the new costs imposed upon them. Second, energy intensive firms may have chosen to locate new large industrial plants in regions where they would not be subject to these new direct regulations. A third way in which the NBP affected manufacturing employment is through electricity prices. Firms may have adjusted their input bundle in response to either increased uncertainty regarding the future price of electricity or to an actual price increase.

These causal mechanisms suggest two primary strands of literature which may be useful in interpreting the magnitude and plausibility of the findings. The first is to simply compare these findings with previous work on environmental regulation and

[27]There is a large literature on firm adjustment costs and their implications. See Hamermesh (1989), Caballero and Engel (1993), Bloom et al. (2007) and Cooper and Willis (2009) for a few examples.

manufacturing employment. The second is to evaluate the NBP's impact on industrial electricity prices and then determine if the employment results fit in line with previous estimates of the employment electricity price elasticity. Labor markets may also have reacted to energy price uncertainty brought on by the regulation. While there is little empirical work on the employment effect of energy price uncertainty, to the best of my knowledge, only two ex-ante simulations of the NBP estimated the impact it would have had on electricity prices and there have been no ex-post studies. The first, performed by the EPA (1999) predicted an increase in electricity prices of 1.6%. The second, performed by Platts Research and Consulting titled *"The NOx Challenge"* (2003), predicted an increase of $1-$3/MWh in the price of wholesale electricity or an increase of between 2.47% and 7.42%. Fowlie (2010) notes that in every state where electricity is regulated, firms successfully petitioned for rate base adjustments in order to cover the compliance costs of the NBP. Linn (2010) and Deschenes et al. (2012) estimate the NBP's cost on utilities, but to my knowledge, there has been no ex-post evaluation of the program's impact on electricity prices.

Using a similar differences-in-differences technique as that employed in the previous section, I can examine whether the predicteded increases in the price of electricity occurred in states that were subject to the NBP. Table 7 presents the results of these electricity price regressions. As discussed in the previous section, controlling for the price of fuels used in electricity production is important if certain regions in the country have a relatively high reliance on certain fuels to produce electricity. Using the same technique as in the robustness check, I interact the average annual fuel price with the percent of electricity that is derived from that fuel in the NERC region and include these terms in the regression. This allows for the electricity price in regions with high reliance on certain fuels to vary with the price of those fuels. Column 2 contains the results with a full set of state and year fixed effects. The coefficient on the PostxEast variable indicates that the NBP increased industrial electricity prices by approximately 5.8% in states impacted by the NBP. Column 3 includes separate East and West trends and while the coefficient is less precisely estimated it remains positive and of an economically significant magnitude. In short, the results in Table 7 suggest that the NBP led to an increase in the price of electricity and while the estimates are not always precisely estimated, they fall within the range provided in ex-ante simulations.

Using this figure together with the employment loss calculated in the previous section allows us to estimate the implied employment electricity price elasticity associated with the NBP and to compare it to recent studies which have sought to estimate this elasticity.

Deschenes (2010) estimates an employment electricity price elasticity of -0.10. Using only manufacturing employment, Kahn and Mansur (2010) find an employment electricity price elasticity ranging between -0.15 for the computer products industry to -1.17 for the energy intensive primary metals industry. Using results from columns 2 and 3 of Table 7 along with the primary employment findings in column 6 of Table 2, this paper suggest an employment electricity price elasticity of between -0.12 and -0.20 for the computer products industry and between -1.12 and -1.88 for the primary metals industry. The implied elasticities associated with the NBP in this study fall more in line with the Kahn and Mansur range.

Although imprecisely estimated, the electricity price results together with a reasonable employment electricity price elasticity adds credence to the main employment findings. To be clear, this elasticity estimate is meant to serve as a plausibility check to the employment finding in the previous section and should not be interpreted as a stand-alone, well identified employment electricity price elasticity. Again, the NBP is likely to have had other impacts on employment in energy intensive manufacturing industries that did not occur through its' impact on electricity prices. Most obviously, the NBP directly regulated the energy production (in the form of heat, steam and electricity) of 140 large manufacturing plants. Furthermore, energy intensive manufacturing firms may have responded to the uncertainty the NBP created in the markets rather than an actual increase in electricity prices themselves. Attributing the entire employment change to electricity prices will overstate the magnitude of the elasticity. Nevertheless, as a plausibility check, the estimates are reassuringly in line with previous estimates, lending confidence to the overall employment effects.

8 Conclusion

This paper has examined the impact of the NOx Budget Trading program on a variety of labor market outcomes. Employment is seen to have declined among energy intensive manufacturing industries. This decline occurred largely as a result of firms decreasing hiring rates. Employment declines were most severe for young workers, but even for this group there is no evidence suggesting increased periods of non-employment for separating workers. Wage offers to new workers fell following the regulations, but incumbent workers say no decline in earnings. These findings also suggest that natural separations are an important mechanism by which firms adjust to demand shocks.

Looking forward, future research should perform ex-post evaluations of energy sector regulation's impact on electricity prices. While a number of papers have examined the impact of energy prices on employment (Davis and Haltiwanger 2001; Deschenes 2010; Kahn and Mansur 2010), there have been surprisingly few ex-post studies of regulation's impact on energy prices. More precise estimates of these impacts will permit a better understanding of the causal mechanisms at play.

References

Abowd, J. M. and Vilhuber, L. (2012). Did the housing price bubble clobber local labor market job and worker flows when it burst?, *American Economic Review: Papers and Proceedings* **102**(3): 589–93.

Abowd, J., Stephens, B. and Vilhuber, L. (2006). The lehd infrastructure files and the creation of the quarterly workforce indicators, *Longitudinal Employer-Household Dynamics Technical Papers 2006-01*, Center for Economic Studies, U.S. Census Bureau.

Becker, R. and Henderson, V. (2000). Effects of air quality regulations on polluting industries, *Journal of Political Economy* **108**(2): 379–421.

Berman, E. and Bui, L. T. M. (2001). Environmental regulation and labor demand: evidence from the south coast air basin, *Journal of Public Economics* **79**(2): 265–295.

Bertrand, M., Duflo, E. and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates?, *The Quarterly Journal of Economics* **119**(1): 249–275.

Bloom, N., Bond, S. and Reenen, J. V. (2007). Uncertainty and investment dynamics, *Review of Economic Studies* **74**(2): 391–415.

Burtraw, D., Krupnick, A., Mansur, E., Austin, D. and Farrell, D. (1998). Costs and benefits of reducing air pollutants related to acid rain, *Contemporary Economic Policy* **16**(4): 379–400.

Caballero, R. J. and Engel, E. M. R. A. (1993). Microeconomic adjustment hazards and aggregate dynamics, *The Quarterly Journal of Economics* **108**(2): 359–83.

Carlton, D. W. (1983). The location and employment choices of new firms: An econometric model with discrete and continuous endogenous variables, *The Review of Economics and Statistics* **65**(3): 440–49.

Chay, K. Y. and Greenstone, M. (2003). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession, *The Quarterly Journal of Economics* **118**(3): 1121–1167.

Cooper, R. and Willis, J. (2009). The cost of labor adjustment: Inferences from the gap, *Review of Economic Dynamics* **12**(4): 632–647.

Curtis, M., Hirsch, B. and Schroeder, M. (2013). Evaluating workplace mandates with flows versus stocks: An application to california paid family leave, *Technical report*, Georgia State University.

Davis, S. J. and Haltiwanger, J. (2001). Sectoral job creation and destruction responses to oil price changes, *Journal of Monetary Economics* **48**(3): 465–512.

Davis, S. and von Wachter, T. (2011). Recessions and the costs of job loss, *Brookings Papers on Economic Activity* **43**(2 (Fall)): 1–72.

Dean, T. J., Brown, R. L. and Stango, V. (2000). Environmental regulation as a barrier to the formation of small manufacturing establishments: A longitudinal examination, *Journal of Environmental Economics and Management* **40**(1): 56–75.

Deschenes, O. (2010). Climate policy and labor markets, *NBER Working Papers 16111*, National Bureau of Economic Research, Inc.

Deschenes, O., Greenstone, M. and Shapiro, J. S. (2012). Defensive investments and the demand for air quality: Evidence from the nox budget program and ozone reductions, *Working Paper 18267*, National Bureau of Economic Research.

Dube, A., Lester, T. W. and Reich, M. (2011). Do frictions matter in the labor market? accessions, separations and minimum wage effects, *IZA Discussion Papers 5811*, Institute for the Study of Labor (IZA).

EPA (1999). Regulatory impact analysis for the final section 126 petition rule, *Technical report*, U.S. Environmental Protection Agency.

EPA (2009). Nox budget trading program: 2008 highlights, *Technical report*, U.S. Environmental Protection Agency.

Farber, H. (2011). Job loss in the great recession: Historical perspective from the displaced workers survey, 1984-2010, *IZA Discussion Papers 5696*, Institute for the Study of Labor (IZA).

Fowlie, M. (2010). Emissions trading, electricity restructuring, and investment in pollution abatement, *American Economic Review* **100**(3): 837–69.

Gittings, K. and Schmutte, I. (2012). Getting handcuffs on an octopus: Minimum wages, employment and turnover, *Technical report*, University of Georgia.

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures, *Journal of Political Economy* **110**(6): 1175–1219.

Greenstone, M., List, J. A. and Syverson, C. (2012). The effects of environmental regulation on the competitiveness of u.s. manufacturing, *Working Paper 18392*, National Bureau of Economic Research.

Hamermesh, D. S. (1989). Labor demand and the structure of adjustment costs, *American Economic Review* **79**(4): 674–89.

Hanna, R. (2010). Us environmental regulation and fdi: Evidence from a panel of us-based multinational firms, *American Economic Journal: Applied Economics* **2**(3): 158–89.

Henderson, J. V. (1996). Effects of air quality regulation, *American Economic Review* **86**(4): 789–813.

Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993). Earnings losses of displaced workers, *American Economic Review* **83**(4): 685–709.

Jerrett, M., Burnett, R. T., Pope, C. A., Ito, K., Thurston, G., Krewski, D., Shi, Y., Calle, E. and Thun, M. (2009). Long-term ozone exposure and mortality, *New England Journal of Medicine* **360**(11): 1085–1095. PMID: 19279340.

Kahn, M. E. and Mansur, E. T. (2010). How do energy prices, and labor and environmental regulations affect local manufacturing employment dynamics? a regression discontinuity approach, *NBER Working Papers 16538*, National Bureau of Economic Research, Inc.

Linn, J. (2008). Technological modifications in the nitrogen oxides tradable permit program, *The Energy Journal* **0**(Number 3): 153–176.

Linn, J. (2010). The effect of cap-and-trade programs on firms' profits: Evidence from the nitrogen oxides budget trading program, *Journal of Environmental Economics and Management* **59**(1): 1–14.

List, J. A., Millimet, D. L., Fredriksson, P. G. and McHone, W. W. (2003). Effects of environmental regulations on manufacturing plant births: Evidence from a propensity score matching estimator, *The Review of Economics and Statistics* **85**(4): 944–952.

Mian, A. R. and Sufi, A. (2012). What explains high unemployment? the aggregate demand channel, *Working Paper 17830*, National Bureau of Economic Research.

Palmer, K., Burtraw, D., Bharvirkar, R. and Paul, A. (2001). Restructuring and cost of reducing nox emissions in electricity generation, *Discussion Papers dp-01-10-rev*, Resources For the Future.

Sullivan, D. and von Wachter, T. (2009). Job displacement and mortality: An analysis using administrative data, *The Quarterly Journal of Economics* **124**(3): 1265–1306.

"The NOx Challenge" (2003). *Technical report*, Platts Research and Consulting.

Walker, R. (2012). The transitional costs of sectoral reallocation: Evidence from the clean air act and the workforce. UC Berkeley Mimeograph.

Walker, W. R. (2011). Environmental regulation and labor reallocation: Evidence from the clean air act, *American Economic Review: Papers and Proceedings* **101**(3): 442–47.

Zivin, J. G. and Neidell, M. (2012). The impact of pollution on worker productivity, *American Economic Review* **102**(7): 3652–73.

Figures and Tables

Figure 1: NBP Compliance Region

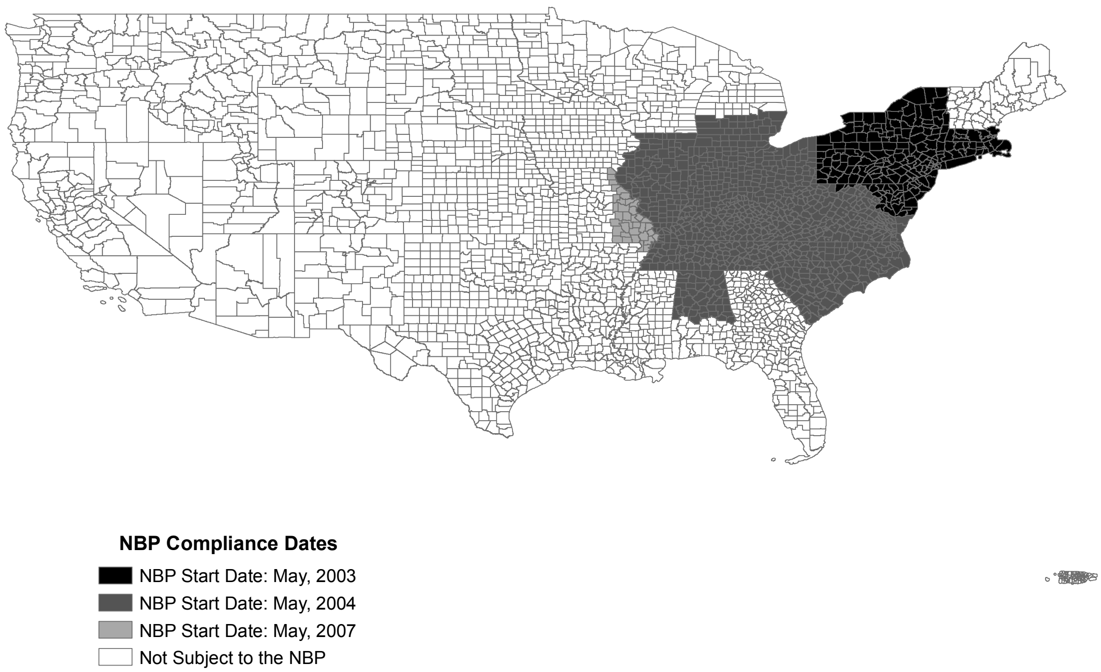

Figure 2: NO_x Emissions From all NBP Affected States

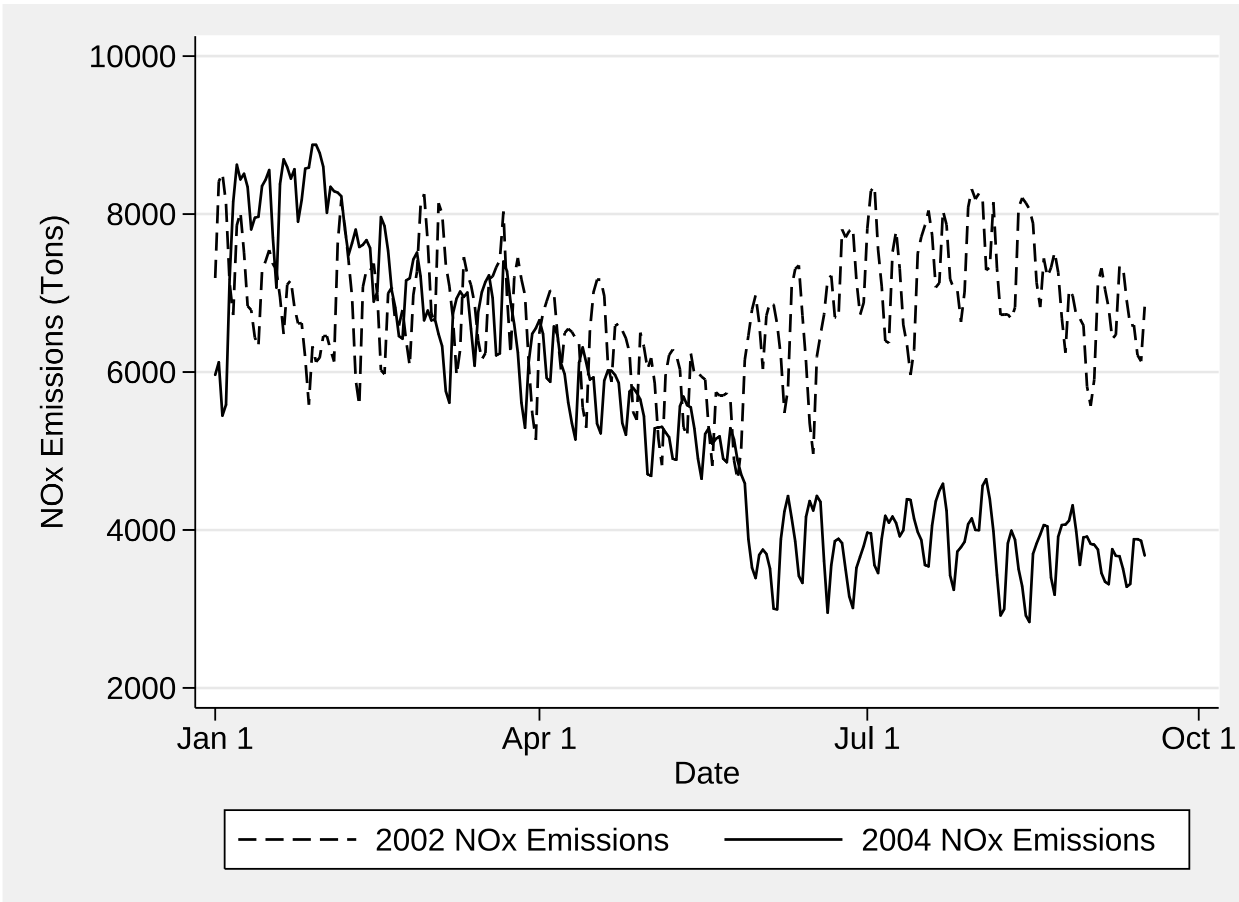

While some northeastern states began in 2003, the program started in full on May 30, 2004. This graph plots daily NO_x emissions in 2002, when no states were participating, and 2004 for the nineteen participating states. There is a visible reduction in the amount of NO_x emissions in NBP states beginning on the start date that is not present in 2002.

Figure 3: East-West Difference in Employment by Energy-Intensity Grouping

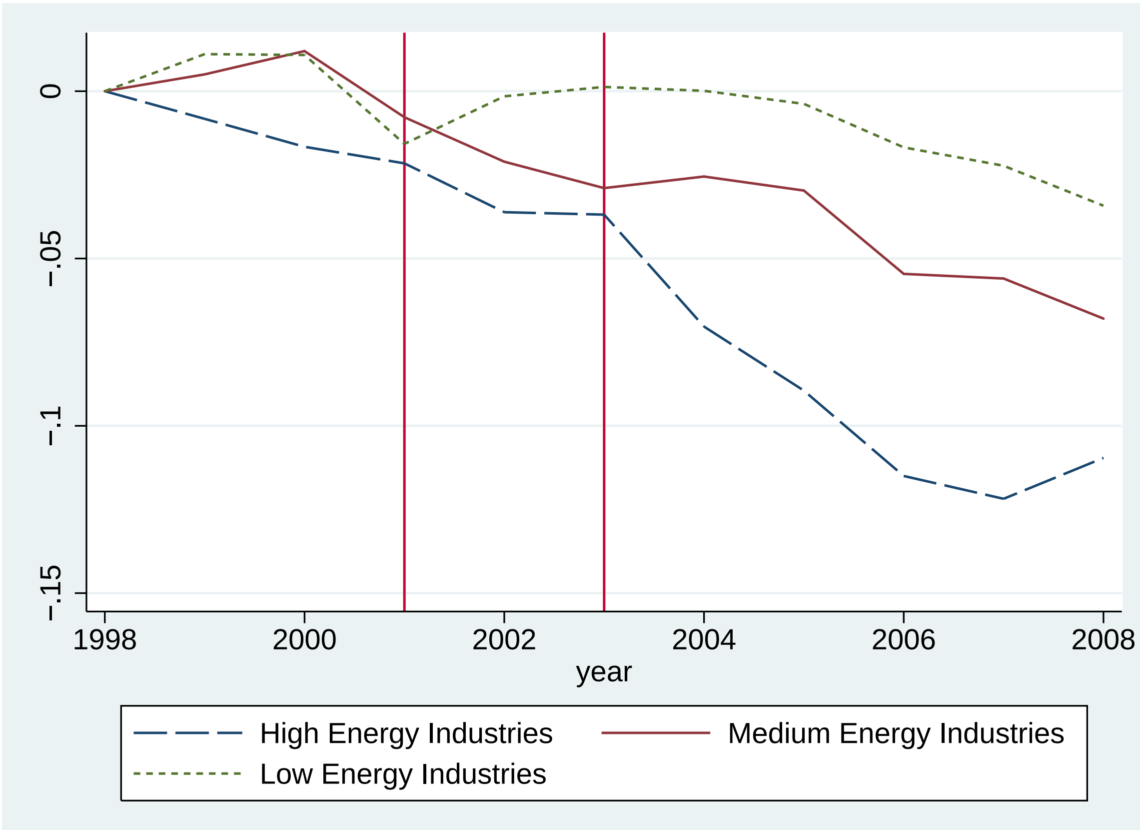

Note: Each point is calculated by $\frac{emp_{g,t,east}}{emp_{g,98,east}} - \frac{emp_{g,t,west}}{emp_{g,98,west}}$. This shows the percent employment change in the east minus the percent employment change in the west. Here t is the year (1998-2008), g is energy intensity group (low, medium, high) and $emp_{g,t,east}$ is the total employment for industry grouping g in east (treated) region in year t. All east-west differences assume the 1998 difference to be the baseline difference, set to zero, against which future differences can be compared. The vertical lines are drawn at 2001, after the policy was approved and 2003, the year the NBP went into effect.

Figure 4: NBP Treated Region

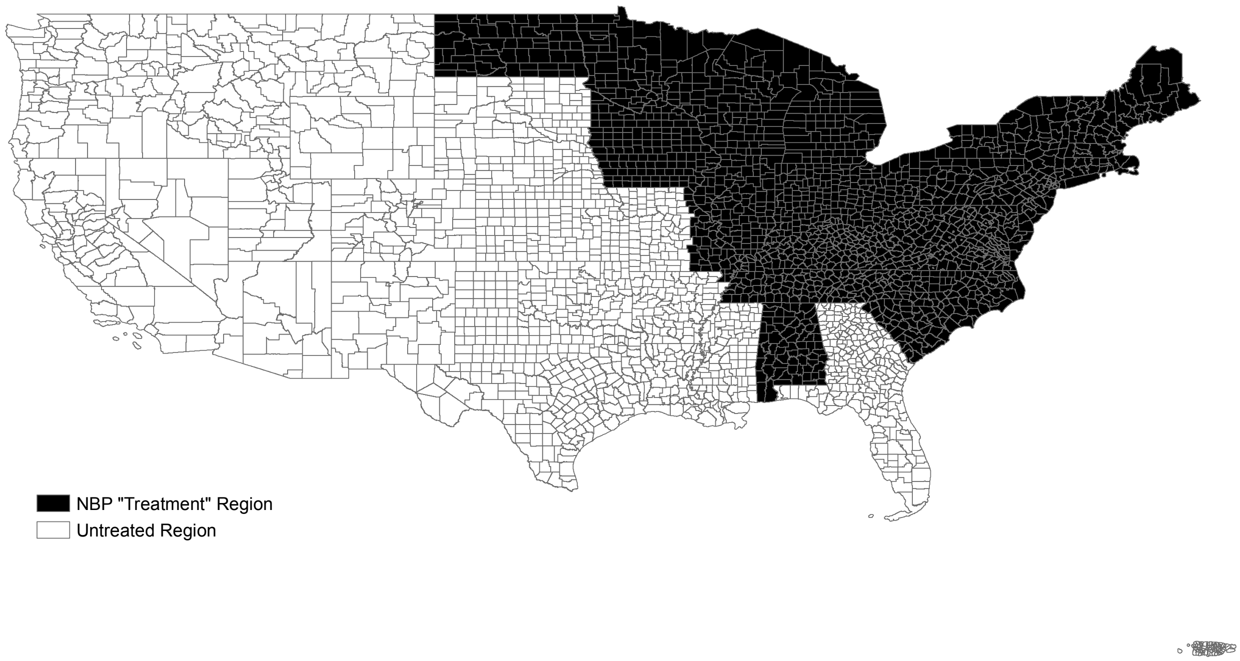

Note: An area is defined as treated if its electricity provider is part of an ISO with coal-burning power plants that were subject to the NBP. For example, Iowa, Minnesota, Wisconsin and North Dakota, while not part of the geographic area of the NBP are part of the Midwest Independent System Operator whose geographic region includes Indiana, Illinois and Michigan.

Figure 5: Industry Coefficients Vs. Energy Intensity

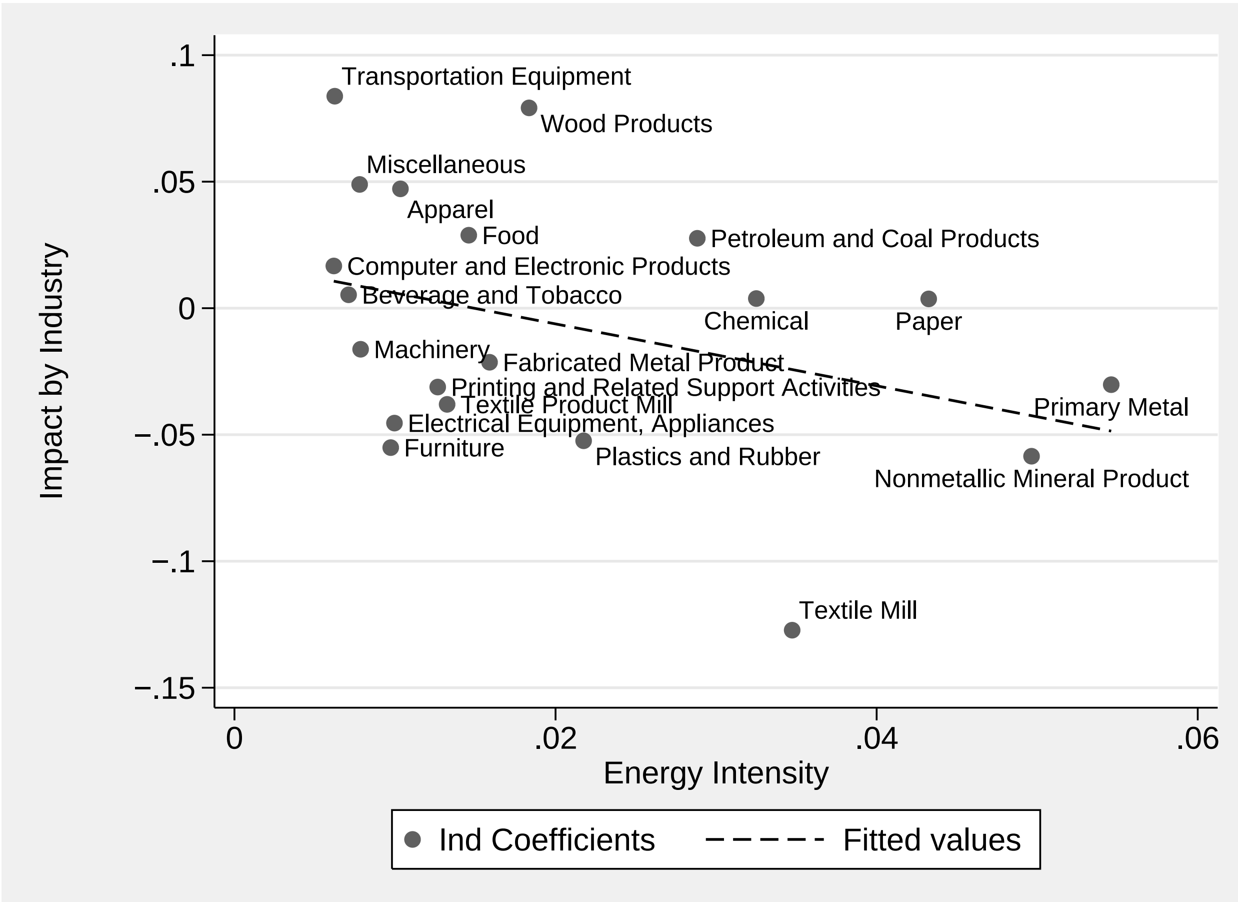

Note: This chart plots the each industry specific triple difference coefficient against that industry's energy intensity measure. The coefficients are obtained by replacing the $PostxEastxEnInt$ variable with 21 industry specific triple interaction variables ($PostxEastxInd1$, $PostxEastxInd2, \ldots, PostxEastxInd21$) and then running the specification described in equation (2). Other model specifications yield similar plots.

Figure 6: Impact of NBP on Employment by Quarter

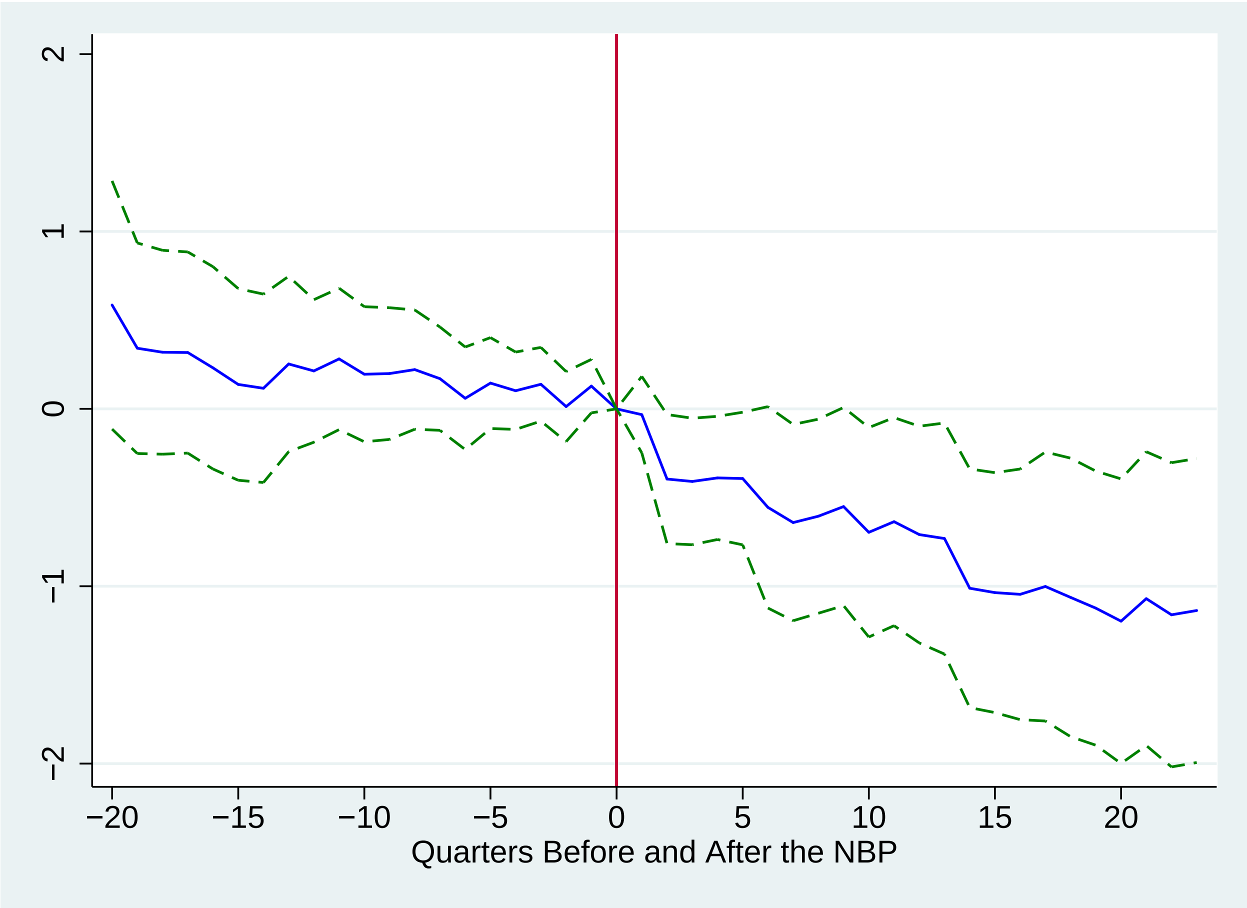

Note: This chart plots the coefficient estimates from a version of equation 3 where the outcome variable is logged employment. Specifically, it plots the coefficients on the event time indicator variables which demonstrate how the impact of the policy on employment evolved over time. The dashed lines represent the 95% confidence intervals. Before the policy the coefficients are slightly above zero but statistically insignificant. After the policy the coefficients fall below zero and their confidence intervals do not include zero.

Figure 7: Impact on Periods of Non-employment for Separating Workers by Quarter

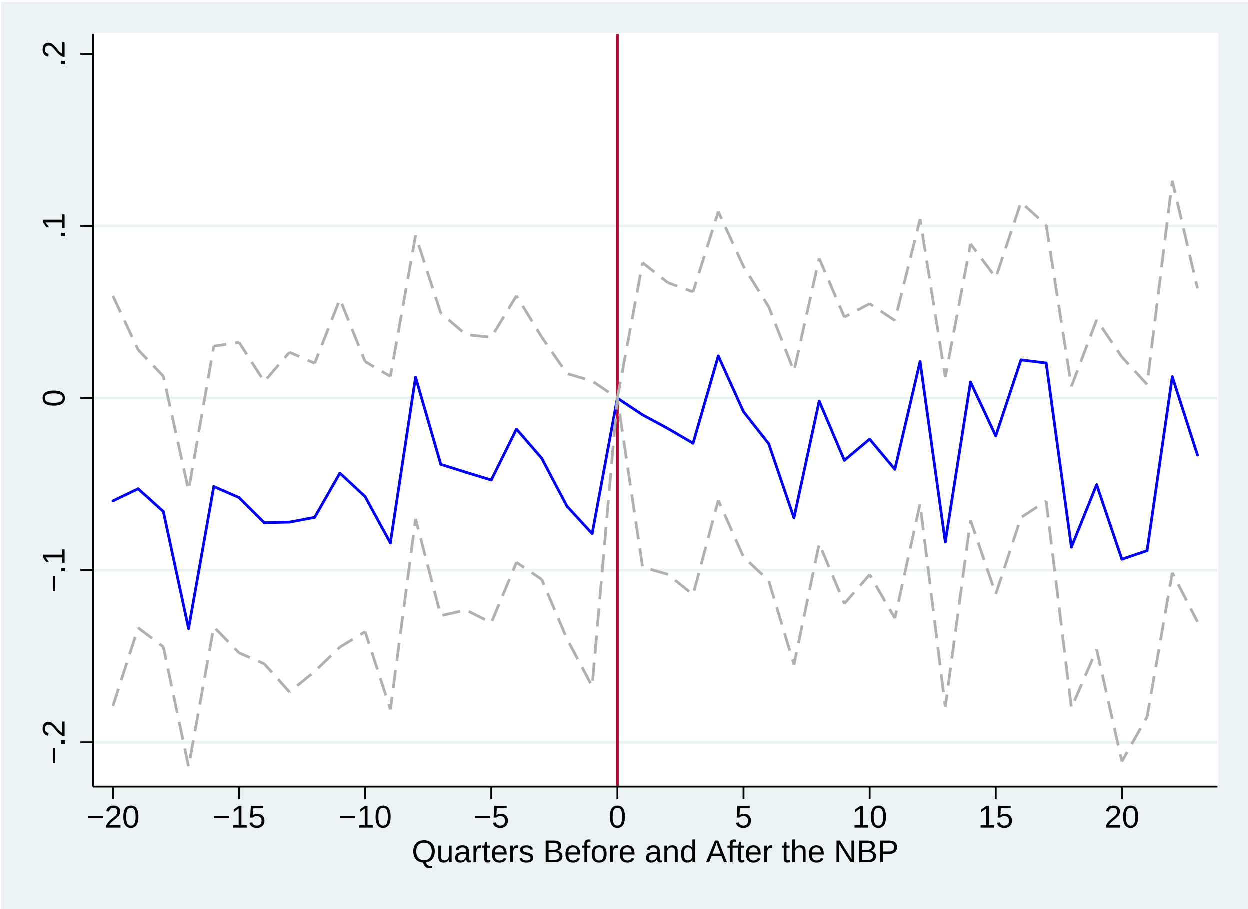

Note: This chart plots the coefficient estimates from a version of equation 3 where the outcome variable is periods of non-employment for separating workers. Specifically, it plots the coefficients on the event time indicator variables which demonstrate the impact of the policy on periods of non-employment for separating workers evolved over time. The dashed lines represent the 95% confidence intervals. The coefficients are always close to zero and are never statistically distinguishable from zero.

Table 1: Energy Intensity of 3-Digit NAICS Manufacturing Industries

NAICS 3-Digit Code	Industry Description	Energy Intensity Level
311	Food Manufacturing	1.46%
312	Beverage and Tobacco Product Manufacturing	0.71%
313	Textile Mill	3.47%
314	Textile Product Mill	1.33%
315	Apparel Manufacturing	1.03%
316	Leather and Allied Product Manufacturing	0.98%
321	Wood Product Manufacturing	1.84%
322	Paper Manufacturing	4.32%
323	Printing and Related Support Activities	1.27%
324	Petroleum and Coal Products Manufacturing	2.88%
325	Chemical Manufacturing	3.25%
326	Plastics and Rubber Products Manufacturing	2.18%
327	Nonmetallic Mineral Product Manufacturing	4.96%
331	Primary Metal Manufacturing	5.46%
332	Fabricated Metal Product Manufacturing	1.59%
333	Machinery Manufacturing	0.79%
334	Computer and Electronic Product Manufacturing	0.61%
335	Electrical Equipment, Appliance, and Component Manufacturing	1.00%
336	Transportation Equipment Manufacturing	0.63%
337	Furniture and Related Product Manufacturing	0.97%
339	Miscellaneous Manufacturing	0.78%

Note: The energy intensity measure is created by dividing the industry's total energy expenditure by their total value of shipments. These variables are obtained from the NBER Productivity Database and use 1998 values.

Table 2: Employment Results: County Business Patterns

	(1) ln(Emp)	(2) ln(Emp)	(3) ln(Emp)	(4) ln(Emp)	(5) ln(Emp)	(6) ln(Emp)
PostxEastxEnInt	-3.861	-2.293**	-2.301**	-1.458*	-1.164	-1.385**
	(2.991)	(0.966)	(0.869)	(0.798)	(0.729)	(0.631)
PostxEast	0.0024		0.0155	0.0113		0.0043
	(0.0799)		(0.0166)	(0.0151)		(0.0164)
State-Ind FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes			Yes		
Ind-Year FE		Yes	Yes		Yes	Yes
State-Year FE		Yes			Yes	
State Linear Trend			Yes			Yes
E / W Ind Trends				Yes	Yes	Yes
Observations	11,319	11,319	11,319	11,319	11,319	11,319
R-Squared	0.985	0.994	0.994	0.993	0.994	0.994

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table reports the main employment results using versions of equations (1) and (2). Robust standard errors are reported in parentheses and are clustered at the NBP Region-industry level. Results are not sensitive to clustering at other levels including, but not limited to, state, industry and state-industry. Column 1 gives the results using equation (1). Column 2 includes industry-year and state-year fixed effects and Column 3 includes state linear trends and industry-year fixed effects. Columns 4 through 6 repeat the specificatios in columns 1 through 3 but now each industry is allowed to trend differently based on it's location. For example, steel industries in the east have a separate trend than steel industries in the west. The coefficient on the *PostxEast* variable drops whenever State-Year fixed effects are included.

Table 3: Employment Results: State Robustness Checks

	(1) ln(Emp)	(2) ln(Emp)	(3) ln(Emp)	(4) ln(Emp)	(5) ln(Emp)	(6) ln(Emp)
Panel A: Fuel Price / Composition Controls						
PostxEastxEnInt	-4.405	-2.509**	-2.459***	-1.247	-1.379*	-1.326**
	(3.203)	(1.037)	(0.904)	(0.909)	(0.711)	(0.633)
PostxEast	0.0100		0.0184	0.0083		0.0039
	(0.0816)		(0.0162)	(0.0165)		(0.0164)
Observations	11,319	11,319	11,319	11,319	11,319	11,319
Panel B: Restricted NBP Region						
PostxEastxEnInt	-3.925	-2.185**	-2.247***	-1.739**	-1.207*	-1.575***
	(3.329)	(0.866)	(0.798)	(0.821)	(0.664)	(0.554)
PostxEast	-0.0100		0.0130	0.0135		0.00397
	(0.0901)		(0.0137)	(0.0152)		(0.0142)
Observations	9,240	9,240	9,240	9,240	9,240	9,240
State-Ind FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes			Yes		
Ind-Year FE		Yes	Yes		Yes	Yes
State-Year FE		Yes			Yes	
State Linear Trend			Yes			Yes
E / W Ind Trends				Yes	Yes	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: These results use the same specifications as those reported in table 2. Panel A controls for changes in fuel prices that may disproportionately affect certain regions. Panel B limits the NBP region to only states which are directly regulate. States which were part of the NBP region but are not directly regulated are dropped from the specification. See the text for additional details.

Table 4: Employment Results: County Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	ln(Emp)	ln(Emp)	ln(Emp)	ln(Emp)	ln(Emp)	ln(Emp)
Panel A: Base Regressions						
PostxEastxEnInt	-3.223	-1.510	-1.456	-1.546*	-1.465	-1.250**
	(2.847)	(0.898)	(0.932)	(0.896)	(0.889)	(0.597)
PostxEast	0.0156		-0.0003	0.0108		-0.0033
	(0.0803)		(0.0159)	(0.0161)		(0.0153)
Observations	374,356	374,356	374,356	374,356	374,356	374,356
Panel B: Restricted NBP Region						
PostxEastxEnInt	-3.270	-1.416	-1.318	-1.689*	-1.550	-1.211*
	(2.997)	(0.897)	(0.913)	(0.955)	(0.943)	(0.655)
PostxEast	-0.0293		-0.0084	0.00980		-0.0102
	(0.0852)		(0.0156)	(0.0168)		(0.0161)
Observations	320,323	320,323	320,323	320,323	320,323	320,323
Panel C: NAAQS Controls						
PostxEastxEnInt	-3.326	-1.573*	-1.523	-1.651*	-1.540*	-1.345**
	(2.825)	(0.914)	(0.935)	(0.915)	(0.912)	(0.593)
PostxEast	-0.0161		0.0083	0.0096		-0.0025
	(0.0806)		(0.0162)	(0.0157)		(0.0152)
Observations	374,356	374,356	374,356	374,356	374,356	374,356
State-Ind FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes			Yes		
Ind-Year FE		Yes	Yes		Yes	Yes
State-Year FE		Yes			Yes	
State Linear Trend			Yes			Yes
E / W Ind Trends				Yes	Yes	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Note: See text for additional details.

Table 5: Job and Worker Flows

VARIABLES	(1) ln(Emp)	(2) Hiring Rate	(3) Separation Rate	(4) Job Creation Rate	(5) Job Destruction Rate
PostxEastxEnint	-0.871**	-0.225**	-0.187	-0.140	-0.104
	(0.411)	(0.0999)	(0.128)	(0.0919)	(0.118)
Observations	33,596	33,256	33,272	33,511	33,511

*** p<0.01, ** p<0.05, * p<0.1

Note: Robust standard errors in parentheses are clustered at the NBP region-industry level. Each column represents the regression coefficient on the triple interaction variable from the model in equation (2) using a different outcome variable.

Table 6: Employment, Flows and Earnings by Age Group

	(1) ln(Emp)	(2) Hires Rate	(3) Seps Rate	(4) Seps: Qtrs of Non-Emp	(5) ln(New Hire Earnings)	(6) ln(Avg Earnings)
All	-0.871**	-0.225**	-0.187	-0.516	-1.277***	0.435
	(0.411)	(0.0999)	(0.128)	(0.636)	(0.470)	(0.299)
Age 19-21	-1.939**	-0.0559	-0.477**	0.719	-0.689*	-0.387
	(0.722)	(0.254)	(0.179)	(0.585)	(0.352)	(0.274)
Age 22-24	-2.121***	-0.288*	-0.358**	0.688	-0.737*	-0.417
	(0.647)	(0.171)	(0.163)	(0.738)	(0.400)	(0.287)
Age 25-34	-0.860	-0.144	-0.154	0.0354	-0.982**	0.141
	(0.588)	(0.122)	(0.126)	(0.699)	(0.436)	(0.217)
Age 35-44	-0.861**	-0.222**	-0.149	-0.170	-1.214**	0.370
	(0.398)	(0.0868)	(0.142)	(0.723)	(0.482)	(0.332)
Age 45-54	-0.718	-0.156	-0.0986	-0.724	-0.887	0.434
	(0.521)	(0.117)	(0.136)	(1.216)	(0.550)	(0.327)
Age 55-64	-0.304	-0.113	-0.214	-0.984	-1.877**	0.303
	(0.387)	(0.105)	(0.138)	(1.370)	(0.723)	(0.328)
Age 65-99	-0.748*	-0.0391	-0.0635	-1.829	-1.664	0.361
	(0.400)	(0.0927)	(0.0888)	(1.111)	(1.267)	(0.484)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Electricity Price Regressions

	(1) ln(Elec Price)	(2) ln(Elec Price)	(3) ln(Elec Price)
PostxEast	0.0456	0.0587*	0.0350
	(0.0321)	(0.0322)	(0.0319)
PerCoalxCoalPrice	0.0228***	0.0121	0.0122
	(0.0054)	(0.0122)	(0.0121)
PerOilxOilPrice	0.0091***	0.0088***	0.0089***
	(0.0008)	(0.0010)	(0.0009)
PerNatGasxNatGasPrice	0.0663***	0.0811***	0.0819***
	(0.0126)	(0.0231)	(0.0229)
State FE	Yes	Yes	Yes
Year FE		Yes	Yes
E / W Trends			Yes
Observations	539	539	539
R-Squared	0.918	0.924	0.925

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: See text for details on fuel prices. Standard errors are robust and clustered at the state level. The dependent variable is log of average industrial electricity prices in a state-year.

A Data Appendix

County Business Patterns does not report observations for which there is no employment in a given year. If firms shutdown and employment in a state-industry goes from 1,000 one year to zero the following year then this drop will not be observed because there will be no record for the zero employment (the same issue could occur in reverse. whereby employment growth that is attributable to new establishments locating in a state-industry with zero previous employment will not be captured because the zero employment wasa not observed in the prior years. To address this concern, I create a balanced panel of every state-industry between 1998 and 2008 by adding zeros when there is no record listed in the CBP. Data that is suppressed for disclosure purposes is imputed by the method used by Kahn and Mansur (2010) as described in the text. Other methods of imputation were explored, but results were not sensitive to the use of other methods.

Employment data is observed for 49% of county-industry pairings. Observed cells contain 68% of all employment in the U.S. A large percent of the overall employment remains because employment is only suppressed for observations with few establishments. Also, some noise is infused in observations for which there are few establishments, but these noise infusions are always less than 5%, sum to zero at the state level and are made in fewer than 5

The QWI is the seond source of data in this paper. It is built from state Unemployment Insurance records and contains 98% of all private-sector, non-agriculture employment at high levels of demographic, geographic and industry detail. Importantly, it contains data on job and worker flows. A job is a relationship between a worker and an establishment where the worker receives positive earnings from that establishment in a quarter. Unlike the CBP, there is no way to impute suppressed data.

Quarterly Workforce Indicator data is also suppressed at times though this is likely to bias against finding a result. If small drops in employment lead to an observation becoming suppressed, then it is possible that employment declines resulting from the NBP will not be observed if the observation becomes suppressed. This would bias against finding a result. The ten states not included in the QWI results are Alabama, Arkansas, Arizona, Kentucky, Massachusetts, Michigan, Mississippi, Missouri, New Hampshire and Wyoming.

Hiring, Separations, Creations and Destructions are all refer to stable jobs. A job is considered stable if the worker receives positive earnings from the establishment for three consecutive quarters. Periods of non-employment for separating workers is obtained by tracking each worker that separates from their firm for the following four quarters. If they are not observed working at any other employer then they are assigned four quarters of non-employment. If, in the quarter following their separation they are observed working at another firm, then they are assigned zero quarters of non-employment. Currently, the QWI is only able to track workers if they find employment in the same state in which they separated. Using state-industry fixed effects accounts for, among other things, time-invariant differences that may arise due to a state's size.

EPA's website provides a list of regulated plants in the NBP. Based on the author's calculation, 93% of regulated manufacturing facilities are in the high intensity industry grouping as defined in section 5.