



**The Effect of Environmental Regulation on Employment:
An Examination of the 1990 Clean Air Act Amendments and
its Impact on the Electric Power Sector**

Ann Ferris, Ron Shadbegian and Ann Wolverton

Working Paper Series

Working Paper # 14-03
February, 2014



U.S. Environmental Protection Agency
National Center for Environmental Economics
1200 Pennsylvania Avenue, NW (MC 1809)
Washington, DC 20460
<http://www.epa.gov/economics>

**The Effect of Environmental Regulation on Employment: An
Examination of the 1990 Clean Air Act Amendments and its
Impact on the Electric Power Sector**

Ann Ferris, Ron Shadbegian and Ann Wolverton

NCEE Working Paper Series
Working Paper # 14-03
February, 2014

DISCLAIMER

The views expressed in this paper are those of the author(s) and do not necessarily represent those of the U.S. Environmental Protection Agency. In addition, although the research described in this paper may have been funded entirely or in part by the U.S. Environmental Protection Agency, it has not been subjected to the Agency's required peer and policy review. No official Agency endorsement should be inferred.

The Effect of Environmental Regulation on Employment: An Examination of the 1990 Clean Air Act Amendments and its Impact on the Electric Power Sector

Ann Ferris, Ron Shadbegian, and Ann Wolverton^{1 2}

Abstract:

The effect of regulation on employment is of particular interest to policy-makers in times of high sustained unemployment. In this paper we use a panel data set of fossil fuel fired power plants to examine the impact of Phase I of the SO₂ trading program created by Title IV of the 1990 Clean Air Act Amendments (CAAA) on employment in the electric utility sector and a two-stage estimation technique that pairs propensity score matching with a difference-in differences estimator. Overall, we find little evidence that power plants subject to Phase I of the SO₂ trading program had significant decreases in employment during Phase I of the program relative to non-Phase I power plants. We also find that accounting for utility-level fixed effects is important when examining how electric utilities chose to comply. For instance, when using plant-level fixed effects we find significant negative employment effects for power plants that chose to comply by switching to low sulfur coal. However, utilities took advantage of the flexibility offered under the trading program by switching to low sulfur coal at a subset of the power plants they owned to generate excess allowances to meet compliance needs at other power plants. When we include utility-level fixed effects in this case, we find that the negative employment effect is no longer statistically significant, offering some evidence that utilities used the flexibility of the regulations to minimize the overall impact on employment. When we control for a more traditional NO_x rate-based standard that partially overlaps with Phase I of the SO₂ trading program, we find that employment effects associated with the SO₂ program continue to be insignificant.

Key Words: SO₂ regulations, cap-and-trade, employment effects

JEL Codes: Q52, Q53

¹ All authors are employed at the National Center for Environmental Economics at the U.S. EPA. Questions or comments on the draft paper can be emailed to the authors at Wolverton.ann@epa.gov.

² Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Environmental Protection Agency.

1. INTRODUCTION

The EPA first established emissions rate-based standards on SO₂ emissions from fossil fuel fired power plants in 1970 under the Clean Air Act as a way to reduce acid rain as well as a number of human health effects. With the Clean Air Act Amendments (CAAA), the EPA used a cap-and-trade system in 1995 to further reduce SO₂ emissions beyond what was required by the rate-based standards. A well-documented result of moving from a command-and-control to a market-based regulatory approach in this context has been the ability to meet the standard more cheaply than would have occurred otherwise.³ However, even with this flexibility, electric power generation ranked first in terms of new pollution abatement capital expenditures during this time period (1994 and 1999) for the industries included in the Pollution Abatement and Cost Expenditures (PACE) survey. Spending on capital for pollution abatement represented about 0.5 percent of revenues from total retail sales of electricity in 1999.⁴ This is roughly consistent with other heavily regulated industries: pollution abatement capital costs for U.S. manufacturing plants are roughly 0.4 percent of total shipments, while they are approximately one percent for pulp and paper, steel, and oil refining in 2005.⁵

Despite the fact that pollution abatement expenditures for power plants are relatively small when compared to sale revenues, a mantra that is sometimes reported as a basic truth by the popular press is that environmental regulation "kills jobs."⁶ The logic behind this statement seems to be that more stringent regulation leads to increased production costs, which raises prices and thus reduces demand for the output produced by the regulated sector and thereby the factors of production including labor, at least in a competitive market. However, even though this effect of regulation on employment might appear obvious at first glance, a careful microeconomic analysis shows that the effect of regulation on employment is ambiguous and therefore warrants empirical study.

³ In one particular study, Keohane (2003) estimated that the SO₂ trading program resulted in cost savings between \$150 million and \$270 million annually, compared to a uniform emissions-rate standard.

⁴ 1999 Pollution Abatement and Cost Expenditures divided by 1999 retail electricity revenues from Energy Information Agency data.

⁵ Calculated from data in the 2005 PACE survey.

⁶ For example, see <http://cnsnews.com/news/article/economic-study-shows-epa-regulations-increase-prices-kill-jobs> .

In this paper we analyze how environmental regulation of power plants has affected employment in this sector. We are particularly interested in examining the employment effects of Phase I of the 1990 Clean Air Act Amendments Title IV cap-and-trade program for SO₂ emissions. Utilizing a panel data set, we examine the impact of environmental regulation on employment using a two-stage estimation technique that pairs propensity score matching with a difference-in differences estimator. We explore specifications that control for power-plant fixed effects as well as differences in regulatory stringency and macroeconomic conditions over time. In addition, we investigate the relevance of controlling for utility-level (i.e. firm-level) fixed effects.⁷ To our knowledge, this has not been explored by other researchers in the context of strategies to comply with Phase I of the SO₂ Title IV trading program for the electricity sector. At the end of the paper, we also control for overlapping NO_x requirements that applied to some of the same plants during this time period. The employment impacts associated with NO_x regulations also serve as a potentially interesting contrast to those associated with Phase I of the SO₂ trading program, since they adhere to the less flexible emission-rate standard approach.

Overall, we find little evidence that power plants subject to Phase I of the SO₂ trading program had significant decreases in employment relative to untreated (i.e., non-Phase I) power plants. Perhaps more importantly, we find that accounting for how electric utilities made decisions regarding compliance strategies is important. More specifically, we find that accounting for utility-level fixed effects matters. For instance, when we use power-plant level fixed effects we find significant negative employment effects for power plants that chose to comply by switching to low sulfur coal. However, there is evidence to suggest that decisions regarding how to comply are made at the utility level. Electric utilities, taking advantage of the flexibility offered under the SO₂ trading program, reportedly chose to switch to low sulfur coal at a subset of power plants to generate excess allowances to meet the compliance needs at other power plants. When we use utility-level fixed effects to examine power plants that switched to low sulfur coal, we find that the negative employment effect is no longer statistically significant.

⁷ We use the term electric utility or utility to designate the firm that owns the individual power plants.

Section 2 reviews the history of regulation in the electric utility sector, focusing in particular on the change in the form and stringency of SO₂ regulation and additional compliance flexibilities that occurred with the passage of the 1990 CAAs. Section 3 reviews the relevant economics literature and presents a conceptual framework of the expected employment effects of regulation on the directly regulated sector. Section 4 discusses the data and empirical approach. Sections 5 and 6 present the summary statistics and main results, respectively. Sections 7 and 8 present sensitivity analyses and the conclusion, respectively.

2. REGULATION OF THE ELECTRIC UTILITY SECTOR

With the passage of the Clean Air Act in 1970, new power plants (i.e., those built after 1970) were subject to emissions rate-based standards, termed New Source Performance Standards (NSPS), for SO₂ emissions. These standards were defined as maximum allowable emission rates in terms of pounds of SO₂ emissions per million Btus of heat input. Older plants were largely grandfathered from these requirements. Carlson et. al (2000) note that EPA essentially required the use of flue gas desulphurization (i.e., scrubbers) by new coal-fired power plants to meet the NSPS beginning in 1978. This was a capital-intensive, expensive piece of equipment to put in place, prompting discussion by some in Congress on whether there was a less costly way to meet national air pollution reduction goals.

With the passage of the 1990 CAAs, market-based regulation became available as a tool for reducing SO₂ emissions. Title IV of the 1990 CAAs created a market for tradable SO₂ permits for power plants by setting an annual cap on the total amount of SO₂ that could be emitted nationwide. Permits were allocated to power plants (actually to boilers within power plants) on the basis of historical use of heat input. The SO₂ trading program provided power plants with more flexibility to achieve the more stringent standard compared to the NSPS: power plants could use any type of technology or production process available to reduce their emissions rate, or those with high marginal abatement costs could buy permits from those with lower marginal abatement costs, and

power plants could bank allowances for future use. While the 1990 CAAAs included specific requirements for the monitoring of SO₂ emissions, the only requirement with regard to compliance with the standard was that, at year's end, a power plant hold one allowance for every ton of SO₂ it emitted.⁸

Title IV also resulted in more stringent regulation. Its objective was to reduce SO₂ emissions from power plants to 8.95 million tons per year by 2010, roughly 50% of the 1980 level, in a cost-effective way (Chan et al., 2012). This extremely large reduction in SO₂ emissions was achieved through the implementation of two phases. During Phase I (1995-1999), EPA required 263 “Table A” units at the dirtiest 110 power plants in eastern and midwestern states to reduce SO₂ emissions by approximately 3.5 million tons per year starting in 1995.⁹ In addition, 182 units at these and other power plants that were not originally part of Phase I joined the program in 1995 as substitution or compensation units. Utilities were allowed to voluntarily “substitute” units (which could be at the same plant or a different plant) scheduled to join the program under Phase II for a Table A unit with higher-cost emission reductions. When a utility opted to include these units, allowances were provided to them based on use of historic heat input in the same way as for the originally designated Table A units. Title IV also allowed utilities to reduce generation below baseline at a Table A unit to reduce SO₂ emissions if they designated a “compensating” unit that would correspondingly increase generation. Phase II, which began in 2000, further reduced the allowable annual emissions of these large, high SO₂ emitting power plants and also imposed constraints on smaller, cleaner coal, oil, and gas fired plants. In addition, some of these same units had to comply with a new NO_x standard for certain types of coal-fired boilers in 1996. Unlike for SO₂, NO_x was controlled using a traditional emissions rate-based standard. This restricted the methods a plant could use to comply with the standard to a relatively narrow set of technologies, essentially requiring installation of a low-NO_x burner technology, though firms were allowed to average across units.

⁸ As an incentive to comply with this regulation a power plant is fined \$2,000 for each ton of SO₂ emitted for which they do not have an allowance.

⁹ A “unit” is a boiler at a power plant – this is what generates the emissions.

A large number of studies have estimated the costs of the cap-and-trade approach to reducing SO₂ emissions, finding that electric utilities made ample use of the flexibility built into the cap and trade program including the banking mechanism to smooth costs over time – over-abating in the early years and then banking permits for use in later years. The first year of Phase I, 1995, resulted in actual SO₂ emissions that were almost forty percent lower than the allowable emissions level (U.S. EPA 1996). Title IV compliance flexibility enabled a wider range of abatement approaches for electric utilities than would have been allowed under the previous command-and-control regime, and cap-and-trade ultimately cost much less to meet the emission goals than initially anticipated (e.g., Burtraw et. al 1998; Carlson et al. 2000; Ellerman et. al 2000; Harrington et al. 2000; Keohane 2003, Popp 2003; and Burtraw and Palmer 2004) For instance, many utilities complied with Title IV by switching to low sulfur coal instead of installing flue gas desulfurization. For our purposes, it is interesting to note that these technologies and production changes also likely vary in labor intensity.

Compliance Flexibility

Title IV afforded greater flexibility to power plants with regard to how they complied with SO₂ regulations. Based on what was reported in utilities' compliance plans, they chose one or a combination of methods to reduce SO₂ emissions under Phase I of the cap-and-trade program: fuel switching and/or blending with lower sulfur coal, obtaining additional allowances (beyond those allocated by the program), installing flue gas desulfurization equipment (scrubbers), using previously implemented controls, retiring units, boiler repowering, substituting Phase II units, or compensating with Phase II units.

In 1995, the majority of Table A units – approximately 52 percent -- chose to switch to or blend with lower-sulfur coal to comply with Title IV (EIA 1997, p. 6). Railroad deregulation in the mid-1980s significantly lowered transportation costs associated with low sulfur coal from the Powder River Basin in Wyoming and Montana, allowing many units in the Midwest to contemplate this as a viable compliance option (Ellerman and Montero 1998).^{10 11}

¹⁰ Units at two Phase I plants in Arkansas switched to using low sulfur coal in the late 1980s/early 1990s

The next most-used method for complying with Title IV of the 1990 CAAs was to obtain additional allowances. Approximately 32 percent of the Table A units chose this option as their primary compliance method. However, Swift (2001) observes that utilities had a tendency to pursue an ‘autarkic or ‘comply on your own’ strategy.’” Utilities made use of the flexibilities afforded by the cap-and-trade program across units they own but rarely traded with other utilities. He notes that smaller utilities are at a particular disadvantage under this strategy since they only have a few units over which they can spread compliance costs.

Table 1: Allocation of Phase I Allowances by Type in 1995

Allocation	Number of Allowances	Percent of Total
Units based on historic utilization	5,550,231	63.4
Units that reduce emissions by 90 percent	1,350,068	15.4
Substitution units	1,220,044	13.9
Early reduction credits	314,248	3.6
Auctioned	150,000	1.7
Compensating units	109,116	1.2
Small diesel fuel refiners that produce and desulfurize fuel	37,558	0.4
Units that undertook efficiency or renewable energy measures	12,816	0.1
Total	8,744,081	100

Source: EPA (1996), Exhibit A.

By way of example, Table 1 shows that the vast majority of Phase I allowances allocated to power plants in 1995 were based on historic (1985-1987 average) heat input. Bonus allowances were given to units to reduce emissions by 90 percent relative to this baseline

for financial reasons alone, prior to the implementation of the SO₂ trading program. The units are still included in Phase I, however, because their baseline SO₂ emissions in 1985 – prior to switching to low sulfur coal - are high.

¹¹ Table A and non-Table A units at plants in Minnesota, New Hampshire, and Wisconsin also may have switched to lower sulfur coal prior to 1995 to meet state environmental regulations or as part of State Implementation Plans (SIPs) enacted under the Clean Air Act prior to the 1990 amendments (EIA 1997, p. 33; Ellerman and Montero 1998). According to Ellerman and Montero (1998), New York, Michigan, and Massachusetts also had enacted regulations but “they were not applicable to coal-fired units in 1993.”

to encourage the installation of flue gas desulfurization equipment (U.S. EPA 1996); to voluntarily reduce emissions after enactment of the 1990 CAAAs but prior to 1995, to desulfurize fuel produced by small diesel fuel refineries, and to undertake efficiency or renewable energy measures. Another 14 percent of total available allowances were allocated to “substitution” units. These units played an important role in many utilities’ compliance strategies, accounting for about 20 percent of the SO₂ emission reductions achieved under Phase I (Swift 2001). A far smaller number of allowances also were allocated to compensating units. Finally, a few allowances were auctioned. In addition, many utilities reduced emissions by more than required in Phase I at their plants to bank allowances for later use during the more-stringent second phase.

Only 10 percent of Table A units (with a capacity of about 16,000 MW) chose to install new flue gas desulfurization systems (scrubbers) in Phase I. While innovations in scrubber design and increased utilization decreased per-ton costs, consideration of bonus allowances and the ability to bank a large number of allowances for use in Phase II drove many of these investments. A number of smaller companies also elected to install scrubbers, even though it is a relatively expensive option (about \$295/KW in 1995 compared to \$50-\$75/KW for low-sulfur coal), to avoid trading with other companies (Swift 2001).

Seven Table A units were retired in Phase I – most of these were outdated and small capacity units in the midwest (EIA 1997; Swift 2001).¹² A number of small capacity substitution units were also retired (Swift 2001). Compliance reports indicate that firms complied with the Title IV requirement to maintain baseline utilization, averaged across their units (in other words, retired generation was made up at existing units). Eight units were repowered with natural gas, fuel oil, or an integrated gasification combined-cycle generator (EIA 1997).¹³

¹² Wisconsin Electric Power Company removed four units from service at North Oak Creek in 1988 and 1989, Indiana-Michigan Power’s Breed plant shut down in March 1994 and is undergoing asbestos removal and may be used again in the future. Cleveland Electric Illuminating’s Avon Lake unit 8 was retired in November 1987, and Iowa Power’s Des Moines unit 7 was reportedly placed out of service (though it could be brought back into service in 180 days).

¹³PSI Energy Inc.’s Wabash River Station unit 1 was repowered with an integrated gasification combined-

3. THE RELATIONSHIP BETWEEN REGULATION AND EMPLOYMENT

While the question of how environmental regulations affect plant operations is not new, few papers specifically examine the effect of environmental regulations on employment.¹⁴ Berman and Bui (2001a) developed a unique plant-level data set to estimate the effect of air pollution regulations on labor demand in the South Coast Air Quality Management District (SCAQMD) of southern California. They find evidence suggesting that air quality regulations designed to bring the area into compliance with various NAAQS, did not reduce the demand for labor in the SCAQMD. Cole and Elliot (2007) estimate a similar model to Berman and Bui (2001a) but use panel data on 27 industries from the United Kingdom. They also find that environmental regulation had no statistically significant effect on employment.

Greenstone (2002) uses a difference-in-difference model to examine the effect of a county being designated by the EPA as out of attainment for criteria air pollutants on employment. Plants located in counties that are out of attainment face stricter environmental regulations than plants that are located in attainment counties. Greenstone finds that nonattainment counties (relative to attainment ones) lost roughly 600,000 jobs over a 15 year time period.¹⁵

Morgenstern, Pizer, and Shih (2002) estimate the effect of abatement spending (their proxy for environmental regulation) on employment for four highly polluting/regulated industries (pulp and paper, plastic, petroleum refining, and steel). They also find evidence that increased regulation does not cause a significant change in employment. More

cycle generator. Using new technology, the plant burns high-sulfur coal, reduces SO emissions, and increases the plant capacity by approximately 155 megawatts. One unit each at Illinois Power's Vermilion plant and Ohio Edison's Edgewater plant were switched to natural gas. Two units at the Long Island Lighting Company's Port Jefferson plant and three units at North Port plant are using No. 6 fuel oil.

¹⁴ For instance, studies have examined the effect of environmental regulation on productivity (e.g., Färe et al. 1986; Boyd and McClelland 1999; Berman and Bui 2001a; and Shadbegian and Gray 2005, 2006), investment (e.g., Gray and Shadbegian 1998, Greenstone 2002), and environmental performance (e.g., Magat and Viscusi 1990; Laplante and Rilstone 1996; and Shadbegian and Gray 2003, 2006).

¹⁵ This is a gross effect and not a net effect, thus Greenstone's result does not mean that there is less aggregate employment due to environmental regulation, it simply suggests that the relative growth rate of employment in some sectors may differ between attainment and non-attainment areas.

recently, Gray et. al (2013) analyze how EPA's Cluster Rule affected employment in the pulp and paper industry, finding some evidence of small employment declines (on the order of 3 percent to 7 percent) associated with the adoption of the Cluster Rule, but these effects are not always statistically significant.

In sum, most past studies using plant-level data have not found large negative impacts of stricter environmental regulation on labor demand, and many have found no statistical effect.

Conceptual Framework

In Berman and Bui's (2001a) theoretical model, the change in a firm's labor demand arising from a regulation is decomposed into two main components: output and substitution effects.¹⁶ First, by changing the marginal cost of production, regulation affects the profit-maximizing quantity of output. An environmental regulation can be interpreted as an increase in demand for a specific type of output: environmental quality. To meet this new demand firms in the regulated sector – in this case, power plants - may increase their demand for various factors of production such as capital (e.g., the purchase of new equipment such as a new scrubber), labor (e.g. to install abatement equipment, monitor the abatement capital, and fill out paperwork), or other inputs (e.g., switching to low-sulfur coal).

At the same time, if the regulation increases production costs (which in most cases, it will) the plant reduces output, thereby reducing demand for factor inputs such as labor. However, a change in the demand for environmental quality also often requires new pollution abatement technologies, some of which may be more or less labor intensive (see Berman and Bui 2001b, and Morgenstern et. al 2002), leading to a shift in the factors of production utilized: the substitution effect. It is not possible ex-ante to predict which of

¹⁶ The authors also discuss a third component, the impact of regulation on factor prices, but conclude that this effect is unlikely to be important for large competitive factor markets, such as labor and capital. Morgenstern, Pizer and Shih (2002) use a similar model but break the employment effect into three parts: the demand effect; the cost effect; and the factor-shift effect.

these effects will dominate. In other words, the net effect of environmental regulation on employment in the regulated sector could be positive, negative, or near zero.¹⁷

The flexibility available in Phase I of the SO₂ trading program, in terms of utilities now able to utilize the lowest-cost compliance method, indicates a distribution of potential labor demand effects, ranging from positive to negative, including zero. In addition, these heterogeneous effects may vary over time. The dynamics of net employment impacts by compliance strategy over time are important to consider. Some compliance methods, such as using previously implemented controls to meet more stringent state emissions standards, may imply no change in employment during Phase I. Other compliance methods, such as obtaining additional allowances or fuel switching or blending, may also have few net impacts. Retiring units should have a negative impact on labor demand at those power plants. However because of the requirement that net generation be maintained labor demand at other power plants is likely increasing.

Other compliance methods could have some effect on labor demand. For example, the installation of flue gas desulfurization units could require more workers in the initial years of the program as equipment is installed and tested but relatively fewer workers in later years for ongoing operation and maintenance. Likewise, repowering of boilers may involve short term increases in labor, but may result in lower demand for ongoing operation and maintenance, depending on the technology used. The use of substitution units should shift employment away from the Table A affected unit towards those Phase II units that are being included.

¹⁷ We focus only on the direct employment effects in the regulated sector. However, it is likely that regulation also changes employment in sectors that produce pollution control equipment, for example. Mapping out these effects to estimate the net effects on employment economy-wide while accounting for the temporal dimension of labor markets is complex. We do not attempt such an exercise in this paper.

4. DATA AND EMPIRICAL FRAMEWORK

To estimate potential employment impacts from the Title IV SO₂ trading program, controlling for key plant level observables, we compile an unbalanced panel dataset of 526 fossil-fuel fired power plants from 1988-1999, for a total of 6,265 plant-year observations.¹⁸ Our main power plant data comes from three main sources: EPA (1996, 1997), the Utility Data Institute's (UDI) O&M Production Cost Database and Energy Information Agency (EIA) Forms 423 and 767.¹⁹ From EPA we obtain data indicating whether a power plant has a Table A boiler covered by Phase I of the SO₂ trading program, as well as which power plants voluntarily agreed to have non-Table A boilers regulated during Phase I – so-called substitution and compensation units. We refer to plants with Table A and substitution and compensation units collectively as Phase I plants (versus non-Phase I plants). We also obtain data on whether or not a power plant has a boiler covered by the NO_x provisions of Title IV, which due to litigation did not begin until 1996.

The UDI's 2012 O&M Production Cost Database is comprised of information from the Federal Energy Regulatory Commission (FERC), which collects data annually for investor-owned utilities on its FERC Form 1, and similar data for municipally owned power plants and rural electric cooperatives on the EIA's Forms 412 and the Rural Utilities Services Forms 7 and 12, respectively. The UDI data set provides us with the average annual number of employees at each of the 526 power plants in our sample.²⁰ We had a difficult time using the UDI data for power plants with multiple owners. In many cases power plants with multiple owners had their employee data misreported as some multiple of actual employment by UDI from 1988 to 1996 or 1997. We contacted McGraw-Hill Platts (MHP) about this issue, but they could not figure out the source of the misreporting issues. However, we were able to confirm that earlier versions of the same database did not have the same misreporting issue; it had been introduced into the

¹⁸ 1988 is the first year in which we could obtain power plant employment data from the Utility Data Institute.

¹⁹ Note that the data on the portions of the EIA767 form we used are boiler level data that we aggregated to the power plant level.

²⁰ This data set was purchased from McGraw-Hill Platts and was also used by Fabrizio, Rose and Wolfram (2007).

database over time. MHP allowed us to use earlier versions of the data that had been sold to other researchers that pre-dated the misreporting issue and we were able to resolve most of these issues.

EIA Form 423 provides us with information on the type and quantity of purchased fuel used by each power plant and its Btu and sulfur content. EIA Form 767 provides us with data on plant gross capacity, net electricity generation, plant age (we base this on the year that the oldest boiler was installed), heat input (in Btus), primary fuel type and the installation date for FGD units for plants that installed them which we aggregated to the plant level. EIA also provides us with historical SO₂ emissions. Unfortunately the UDI and EIA data do not have a common numerical identifier, therefore we merged these two data sets based on power plant name, owning utility name and state.²¹ We augment our main data sources with data from the Bureau of Labor Statistics (BLS) and Denny Ellerman. From BLS we get the annual average utility worker wage by state. Denny Ellerman provided us with the distance from each power plant to the Powder River Basin, the main source of low sulfur coal.

Our sample began with 806 power plants, which in principal had data in EIA767 and were in operation prior to 1990. Sixteen of these power plants were not part of Title IV, so they were dropped from our sample. Of the remaining 790 plants we dropped 18 more due to data reporting issues (17 of them stopped reporting data to EIA sometime around 1987 and one never reported output data (kwhs)). Since Title IV focuses on fossil fuel fired power plants we also dropped three nuclear plants. We then merged these 769 electricity generating plants with the UDI data. We were able to match all 110 Table A Phase I plants and 603 Phase II plants. We then dropped an additional 187 plants due to missing data for heat input (133 plants) or employment (10 plants), poorly reported employment data (18 plants), and missing employment data after 1990 (26 plants). Our final data set includes 133 Phase I plants: 103 of the 110 Table A power plants and 30

²¹ Only a few power plants have the same name and no power plant within a state has the same name so we are confident that we merged the data sets properly. The two data sets have some common variables that we used to double-check in cases when we were uncertain about a particular matched power plant.

power plants that opted into Phase I as compensation or substitution units.²² Three of the 103 Table A plants and two of the 30 compensation-substitution units closed during Phase 1; we include them in our data set with zero employment after shut down. The remaining 393 power plants in our final data set are electricity generating plants included in Phase II of the Title IV SO₂ program. Our final data set is quite comprehensive, accounting for 95 percent of the kilowatt hours reported to EIA in 1990.

The last year of our data set is 1999 for two reasons. First, in 2000 nearly all fossil-fuel fired power plants become part of the Title IV SO₂ trading program. The only exceptions are very small and older peaking plants that were grandfathered and therefore did not have to comply with Title IV SO₂ requirements. Both of these groups are sufficiently different from the Phase I plants that they likely would not serve as a reasonable counterfactual. Second, the reporting of employment data drops off tremendously in the early 2000s and this would alter our sample.

Empirical Methodology

We are particularly interested in identifying the effect of Phase I of the Title IV SO₂ cap-and-trade program on plants' employment relative to the less stringent command-and-control regulatory regime. To accomplish this, we employ a difference-in-difference estimator. The difference-in difference approach takes advantage of the fact that only 110 plants were part of Phase I of the cap-and-trade program, 103 of which are in our final sample. As previously mentioned, another 30 plants opted into Phase I (from Phase II) as compensation-substitution units. We use the remaining Phase II plants to substitute for a true counterfactual for the Phase I plants - we do not know what employment would have been for Phase I plants (i.e., the treated group) if the prior regulatory program had remained in place (i.e. they had remained untreated).

Difference-in-difference estimation is most appropriate when the treatment (in our case

²² We dropped seven Table A plants (identified here by their EIA code) for reasons mostly related to their reported employment data: 1) 1083 because it stopped operating in the mid-80's; 2) 1091 because we could not resolve the multi-owner problem described above; and 3) 1295, 2049, 2835, 2836, and 2837 because they did not report employment data for over half the years in our data set (mostly in the middle years), even though they were still using typical levels of other inputs and producing normal levels of output.

being designated a Phase I plant) is random, or observable characteristics can be used to control for treatment. However, as we noted above units were not randomly selected to be Phase I plants. Furthermore, power plants in our control group (Phase II plants) may not be very similar, based on average observable characteristics, to our Phase I plants, thus our difference-in-difference estimator may be biased. Rubin (2008) argues that we can approximate a randomized experiment by choosing a suitably-matched control group to eliminate or at least reduce this bias. To obtain approximately unbiased estimates we need a control group that is not systematically different from the Phase I plants in our sample (Stuart and Rubin 2007). We use a version of the propensity score matching technique (developed by Rosenbaum and Rubin 1983) based on pre-Title IV attributes – aside from the outcome variable, employment - to select a statistically defensible comparison group from non-treated (Phase II) plants. Then, we use a difference-in-difference estimation technique to investigate how implementation of the SO₂ trading program under Title IV affected Phase I plant employment (1995 – 1999) relative to what occurred prior to the program as well as for our control group.²³

Combining the propensity score matching and difference-in-difference estimation techniques allows us to match a Phase I plant with its closest Phase II neighbor and then compare employment across the two sets of plants.²⁴ The propensity score matching estimation uses a probit regression (the dependent variable is equal to one for Phase I plants and 0 otherwise) where the independent variables are pre-treatment (pre-1995) characteristics that may affect a plant’s “propensity” to participate in the Title IV SO₂ trading program. Plants are “matched” with their nearest neighbor using the propensity score, which is a scalar summary of the included pre-treatment characteristics from the probit regression. The matching objective is to control for pre-existing differences between the treated and untreated groups such that the observed covariate distributions

²³ A regression discontinuity approach will not work here for several reasons: (1) Emission and size criteria to qualify for Phase I were specified at the boiler, not the plant level. However, we perform our analysis at the plant level (employment is only available at the plant level). (2) Plants had the option to opt-in to Phase I early by designating Phase II units as substitution units.

²⁴ Heckman et al. (1997), Heckman et al. (1998), and List et al. (2003a) employ a similar matching difference-in-difference approach.

are only randomly different from each other, replicating a natural experiment.²⁵ Using this technique, we assemble a matched sample that consists of the treatment group and its nearest untreated neighbors.

The difference-in-difference technique then estimates the average treatment effect of Phase I of the SO₂ trading program on employment. Our basic specification is as follows:

$$\ln EMP_{pt} = \beta_0 + \beta_1 PHASE I + \beta_2 YR9599_t + \beta_3 PHASE I * YR9599_{pt} + u_{pt} \quad (1)$$

where p indexes plants and t indexes years; $\ln EMP$ is the log of employment; $PHASE I$ equals one for a plant that must comply with the first stage of the cap-and-trade under the 1990 CAAA, and zero otherwise; $YR9599$ equals one for the Phase I years 1995-1999, and zero otherwise; and $PHASE I * YR9599$ is the interaction term between the $PHASE I$ dummy and the $YR9599_t$ dummy, which captures the change in employment at Phase I plants relative to Phase II plants during the Phase I years. Its coefficient β_3 thus measures the difference-in-difference effect where

$$B_3 = (\underline{PHASE I * YR9599}_{Phase I=1, PostPolicy=1} - \underline{PHASE I * YR9599}_{PHASE I=1, PostPolicy=0}) - (\underline{PHASE I * YR9599}_{Phase I=0, PostPolicy=1} - \underline{PHASE I * YR9599}_{Phase I=0, PostPolicy=0}) \quad (2)$$

$Phase I * YR9599$ is underlined to denote that the parameter measures the expected value (or average) difference-in-differences across the two groups.

We also take advantage of the panel nature of our data set by adding a power-plant specific fixed effect, a_p , and quadratic time trend to the basic specification in (1). The inclusion of the time and plant-specific fixed effects means that we can now control for general macroeconomic factors that affect all plants over time as well as plant-specific characteristics that are time-invariant. However, it also implies that we can no longer independently identify the coefficient on $PHASE I$, so it drops out of the specification:

²⁵ See Fowlie et al. (2012) for an example of this method applied in a different environmental context, the evaluation of Southern California's RECLAIM NO_x trading program.

$$\ln EMP_{pt} = \beta_0 + \beta_2 YR9599_t + \beta_3 PHASE1 * YR9599_{pt} + a_p + t + t^2 + u_{pt} \quad (3)$$

Ideally, we would test several alternative measures of the dependent variable, employment, at the plant level – for example, to distinguish between production and non-production workers since the effects of regulation are likely different across these groups. A regulation that requires paperwork and procedural compliance may imply the need to hire additional non-production workers, while changes in compliance costs, the need for new abatement equipment, or changes to the production process may affect production worker employment. However, the only labor data we have available from FERC, EIA and RUS are for average total employment at the plant on an annual basis so we are not able to explore differences in employment by type of worker.

5. SUMMARY STATISTICS

Our analytic dataset consists of annual plant-level data for U.S. electric utility generating plants that utilize fossil-fuels. As described earlier, our data are derived from EPA, EIA, FERC and RUS as well as from the BLS which provides data on state-level wages for the utility sector, and also from Ellerman et al (2000) for distance to PRB.

Table 2 presents summary statistics for the 526 power plants, in our sample, separately into Phase I and non-Phase I plants under Title IV. We report averages for these sets of power plants in 1988 to describe their characteristics before Phase I, and also to roughly correspond with the allocation method for Phase I, which was based on historical heat input in 1985 to 1987. Our primary empirical analysis focuses on coal plants, with Table 2 reporting summary statistics for our treatment group, which is composed of those plants subject to Title IV Phase I in 1995 (column 1). The untreated group is composed of power plants (including some powered by natural gas and oil) that were not subject to Phase I (column 2). Recall that we also use matching estimation techniques to refine the sample of non-Phase 1 plants (column 3). The method used to arrive at the matched sample of non-Phase 1 plants is discussed in detail in section 6.

Table 2 – Summary of Average Power Plant Characteristics, 1988

Variables:	<i>Phase I Plants (Treated)</i>	<i>Non-Phase I Plants (Untreated)</i>	
	<i>(N = 131)</i>	<i>Full Sample (N = 393)</i>	<i>Matched Sample (N = 131)</i>
<i>Number of employees</i>	222 (160)	159 (139)	226 (170)
<i>Nameplate capacity (MW)</i>	958 (761)	780 (651)	942 (786)
<i>Plant age (years)</i>	29 (11)	26 (13)	28 (10)
<i>Percent coal (Coal BTUs/Total BTUs)</i>	96.9 (14.1)	22.7 (34.7)	96.1 (13.3)
<i>Distance to Powder River Basin (in miles)</i>	1,120 (263.6)	1,089 (405.6)	1,053 (400.8)
<i>SO₂ emissions / heat input in 1985 (lbs/mmBtus)</i>	3.6 (1.7)	0.82 (0.86)	1.42 (0.81)
<i>Percent with flue gas desulfurization (FGD) installed at least one boiler (1988 – 1999)</i>	22 (41)	18 (38)	20 (40)
<i>Year FGD first installed</i>	1987 (7.7)	1982 (5.5)	1983 (7.0)
<i>Percent fuel-switch (1990 – 1995)</i>	18.1 (38.7)	5.5 (23.0)	6.2 (24.3)
<i>Number of boilers at plant</i>	3.58 (2.04)	3.23 (1.81)	3.58 (1.58)

Notes: Means, with standard deviations in parentheses, are for 1988 unless otherwise noted.

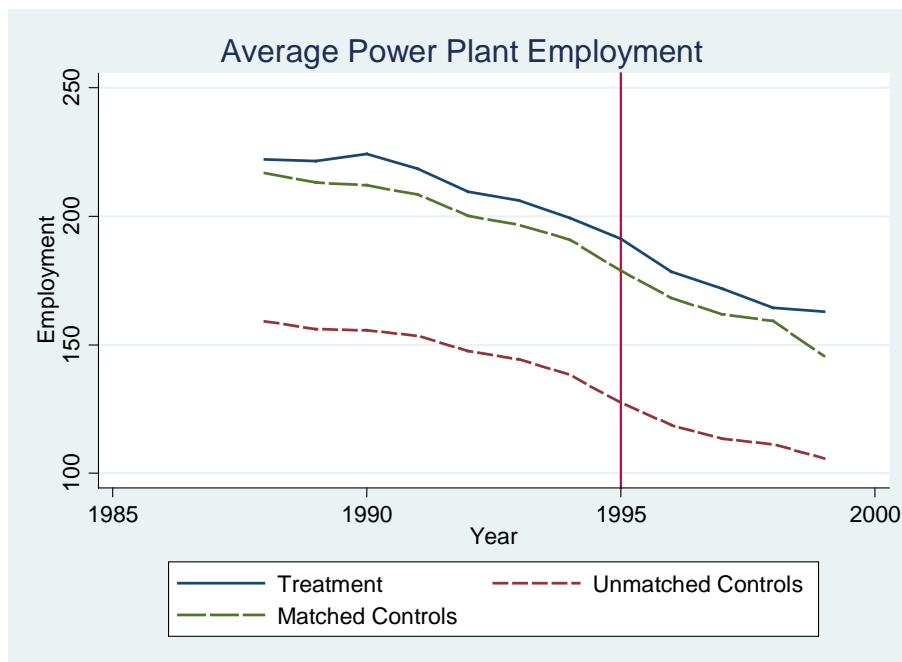
The summary statistics show that Phase I plants had a larger number of employees in 1988 on average compared to the full sample of non-Phase I plants (222 versus 159). When comparing Phase I plants to the matched sample, however, they are very similar with regard to average employment. Phase I plants were also larger and slightly older in 1988 than non-Phase I plants, on average (958 megawatts (MW) in nameplate capacity and 29 years old versus 780 MW and were 26 years old in the unmatched sample). Likewise, the share of coal Btus is much higher for Phase I plants than it is for the full sample of non-Phase I plants (96.9 percent versus 22.7 percent). By design, the non-Phase I plants in the matched sample are more similar to Phase I plants in terms of these characteristics: on average, they had a nameplate capacity of 942 MW, were about 28 years old, and about 96.1 percent of total Btus came from coal in 1988. Phase I plants are quite similar to non-Phase I plants in both the full and matched samples with regard to average distance to low-sulfur coal in the Powder River Basin and the average number of boilers at a plant.

Given the targeting of the Phase I SO₂ cap-and-trade program, it isn't surprising that plants in Phase I had a lower rate of flue gas desulfurization (FGD) units installed as of 1988 but this trend is reversed for the sample period, 1988-1999, though the three groups are fairly similar. However, the year of installation varies dramatically. Specifically, plants in both the unmatched and matched non-Phase I groups have an average installation year of 1983 for FGDs – well before the 1990 CAAs. In comparison, Phase I plants installed FGDs installed later, in 1987, on average. When substitution or compensation units are excluded, the average year a FGD is first installed by a Phase I plant is 1994 with a standard deviation of 1.1 years. Finally, Phase I plants had notable higher SO₂ emissions per unit of heat input: 3.6 lbs/mm Btus, in 1985, compared to non-Phase I plants (an average of 0.82 lbs/mmBtus for the full sample and 1.42 lbs/mm Btus for the matched sample).

The difference-in-difference estimator has the advantage of differencing out pre-existing variation between Phase I and non-Phase I plants to reduce selection bias while also controlling for other potentially confounding factors that may have changed around the

time of the SO₂ trading program and would have affected both sets of power plants similarly. However, this estimator requires the strong identifying assumption that, without Title IV, the average employment for the Phase I plants and control groups would have exhibited similar trends over time. Figure 1 presents average employment at coal plants, by our treatment and comparison groups, for the time period of our sample: 1988 to 1999. The trends in average employment prior to Phase I in 1995 are similar, separated by a level difference of approximately 4 employees or so, for the matched control group (see Table 2). These similar trends in average employment prior to Phase I support our empirical approach, in that our comparison group would likely have continued a similar trend to the treatment group in absence of the regulation.

Figure 1 – Average Employment at Phase I and Non-Phase I Plants: 1988 - 1999



6. MAIN RESULTS

For comparison purposes, Table 3 presents naïve pooled and fixed effects difference-in-difference estimations, conducted without first selecting the untreated sample of power plants based on propensity score matching. These results ignore the possibility that inherent differences between Phase I and other power plants may be misattributed to the SO₂ cap-and-trade policy.

Table 3: Naïve Difference-in-Difference Results: No Matching
(Dependent Variable = Log of Employment)

	Pooled	Utility-Level Fixed Effects
	(1)	(2)
Phase I Plants	-0.02 (0.02)	
Phase I Period	-0.30 *** (0.02)	-0.11 *** (0.03)
Phase I Plants x Phase I Period	0.09 *** (0.03)	0.12 (0.09)
Plant Age	0.006 ** (0.002)	
Nameplate Capacity (logs)	0.50 *** (0.02)	
Coal share of Btus	0.74 *** (0.02)	
Plant age squared	0.0002 *** (0.00)	
Capacity squared	0.0 *** (0.00)	
Plant age x capacity	0.00 (0.00)	
Constant	1.09 *** (0.14)	4.90 *** (0.05)
Quadratic Time Trend	Yes	Yes
Observations	6,212	6,265
Groups	--	203

Column (1) presents the basic pooled difference-in-difference and includes the non-time varying variables used for matching purposes as regressors. Column (2) presents the results of a panel regression using utility-level fixed effects (note, non-time varying variables drop out of the regression), though results are similar when plant-level fixed effects are utilized. The pooled regression performs fairly well, explaining about 52

percent of the variation in log employment. The results indicate that older, larger, more coal-intensive plants have higher employment on average. Phase I plants do not have statistically distinguishable differences in employment compared to non-Phase I plants over the entire sample period, but all plants have significantly lower employment during Phase I. The coefficient of interest – the interaction between Phase I and YR9599 - is positive and significant in the pooled regression, indicating that plants subject to Phase I of the SO₂ program are shedding workers less rapidly than the average decrease already occurring during the Phase I period. However, even before matching, this finding disappears when utility fixed effects are included.

To ensure that the results in Table 3 are not an artifact of differences between Phase I and non-Phase I plants that have nothing to do with the SO₂ trading program, we move to a difference-in-difference estimation based on a matched sample.

Propensity Score Matching

To isolate the effect of Phase I on employment, it is important to control for pre-treatment or time-invariant plant characteristics that might affect employment differently between the Phase I and Phase II plants. We try a variety of approaches to assemble our matched sample, including matching with and without replacement (i.e., a non-treated observation can be selected more than once if it is the best match for multiple treated plants vs. only being selected once) for each specification. Matching with replacement is expected to improve the closeness of the match between treated and untreated plants with regard to pre-policy characteristics and therefore result in a sample of untreated plants that most closely resemble treated plants. However, this advantage must be balanced against the decrease in the number of untreated observations serving as counterfactuals, which can affect the precision of coefficient estimation (Stuart and Rubin 2007). We also vary what covariates are included in the propensity matching estimation and the maximum distance of each match for our preferred specification (see Table 4). In particular, we adjust the maximum distance allowed for a match by using a 0.25 standard deviation of the logit transformation of the propensity score per Stuart and Rubin (2007). We also examine two other calipers, one more and one less precise.

Plant characteristics such as age – measured as observation year minus the year the first boiler was operational – and size – measured by nameplate capacity (gross MWs) - play an important role in determining emission levels. These characteristics may determine whether a plant is subject to Phase I and may also have a direct effect on employment levels in the pre-treatment period. Reliance on coal is another important characteristic of Phase I plants; we include the share of total Btus from coal in the regression.

Table 4: Propensity Score Matching Specifications, With and Without Replacement

<i>Independent Variables</i>	1	2	3	4	5	6	7	8	9	10	11	12
Plant age	√	√	√	√	√	√	√	√	√	√	√	√
Nameplate capacity	√	√	√	√	√	√	√	√	√	√	√	√
Coal share of Btus	√	√	√	√	√	√	√	√	√	√	√	√
Distance to PRB		√			√		√	√	√	√	√	√
NOx dummy						√	√	√				
Squared terms for plant age, capacity			√	√	√				√			
Plant age x capacity				√					√			
Square term for PRB					√			√	√			
Max distance 0.25 std dev.										√		
Max distance 0.125 std dev.											√	
Max distance 0.375 std dev.												√

Ellerman and Montero (1998) also point out that “virtually all” SO₂ emission reductions between 1985 and 1993 occurred at Midwest plants, which had access to cheap low-sulfur coal from the Powder River Basin (PRB) due to their relative proximity. They note that most of the plants that significantly increased their use of PRB coal or began to newly purchase PRB coal during this time period did not have units subject to Phase I of Title IV. Since the use of PRB coal is an important compliance strategy and distance to the PRB is a good proxy for transportation costs, we also consider distance to PRB coal as a possible variable in the propensity score estimation. Finally, we create a dummy variable to identify plants which, due to their historic NOx emissions rate and boiler-type, are subject to rate-based standards for NOx emissions under Title IV (beginning in 1996). This dummy variable accounts for differences in boiler technology across coal plants.

Table 5 briefly describes whether a specification eliminates statistically measurable bias evident in the full sample (using pairwise t and F statistics) and the number of observations in the matched control group. The overall fit of the matching estimation for the entire sample of treated and untreated plants ranges from a pseudo R^2 of 0.24 to 0.36 with and without replacement. Without replacement, fewer than 133 observations indicate that some observations are dropped due to a lack of common support between the treated and untreated distributions.²⁶ With replacement, the number of observations indicates how many unique non-treated plants remain in the matched sample (i.e., some may also drop because of lack of common support).

Table 5: Propensity Score Matching Estimation - Bias and Control Group Observations

	1	2	3	4	5	6	7	8	9	10	11	12
<i>Without replacement</i>												
Bias?	No	No	No	No	No	No	Yes	Yes	No	No	No	No
Control group observations	132	132	131	131	132	132	132	125	127	132	130	132
<i>With replacement</i>												
Bias?	No	No	No	No	Yes	No	No	Yes	Yes	No	No	No
Control group observations	85	88	78	83	74	74	74	78	72	88	88	88

Of the specifications that do not identify a maximum distance for a match, 1 – 4 and 6 eliminate bias from the full unmatched sample both with and without replacement. Specifications 1 and 2 successfully match all but one of the Phase I plants (i.e., they have common support) without replacement. With replacement, these specifications also draw on a larger sample of non-Phase I plants than other specifications. Specifications 3 and 4 drop one additional plant due to lack of common support but also draw on a fairly broad

²⁶ “Common support” refers to the areas where there is distributional overlap (areas of the covariate space that include both treated and control units). Performing analyses only in areas with common support will result in more robust inference (Stuart and Rubin 2008).

sample of untreated plants when matching is conducted with replacement. Specifications 10 – 12 are the same as specification 2 except that they specify the maximum distance allowed for a match. The results are not sensitive to the caliper used.

Table 6 presents the propensity score matching estimation results for specification 4. While several specifications perform well, we elect to use specification 4 as our main matching estimation because it relies on variables that are mainly related to whether a plant is likely subject to the CAAA Phase I SO₂ requirements versus the method of compliance a plant may use (e.g. distance to PRB).

Table 6: Propensity Score Matching Estimation to Select Sample of Non-Treated Plants for Specification #4

			Without Replacement	With Replacement
Variables	Coefficient Estimate (Standard Error)	Percent Bias Before Match	Percent Bias After Match	Percent Bias After Match
Plant age	0.12 *** (0.03)	19.3 *	1.8	2.6
Nameplate capacity (logs)	0.0004 (0.0005)	25.2 ***	-0.5	-1.2
Coal share of Btus	2.33 *** (0.30)	123.9 ***	2.3	-7.6
Plant age squared	-0.002 *** (0.0005)	11.1	2.6	-6.1
Capacity squared	-0.00 (0.00)	23.4 ***	-3.4	-7.4
Plant age x capacity	-0.00 (0.00001)	36.2 ***	0.7	-4.7
Constant	-4.62 *** (0.60)			
Pseudo R ²	0.27			

*** indicates a coefficient that is significant at the 1 percent level; * indicates significance at the 10 percent level.

While we do not want to overemphasize the results of the propensity score matching estimation, since their main purpose is to identify an appropriate control group, we find it useful to confirm that they have the expected signs. As expected, older plants are more

likely to be part of Phase I of the SO₂ program. Likewise, larger plants – proxied by nameplate capacity – and more coal-intensive plants are more likely to be part of Phase I. With the exception of squared plant age, there is statistically significant bias evident in the regressors prior to matching. This bias is removed with matching. Figures 2 and 3 illustrate graphically how matching based on the propensity score reduces these biases. Prior to matching the overall mean bias for specification 4 was 39.9 percent. After matching without replacement, mean bias decreased to 1.9 percent. After matching for the same specification with replacement, mean bias decreased to 4.9 percent.

Figure 2: Bias Before and After Matching: Specification #4 without Replacement





Figure 3: Bias Before and After Matching: Specification #4 with Replacement



Unlike many prior studies of the SO₂ cap-and-trade program, we rely on utility-level fixed effects due to the reticence to trade with other utilities, compliance strategies were reportedly often coordinated across plants under the same ownership. Swift (2001) notes that of 51 utilities subject to Title IV between 1995 and 1999, 67 percent adopted low-sulfur coal at a subset of units to generate excess allowances for use by other units they owned. Another 31 percent installed a scrubber at one plant to generate excess allowances for other units. Media reports about specific utilities suggest similar behavior with regard to employment: utilities primarily transferred and relocated employees among plants, rather than relying solely on layoffs, when faced with plant closures.²⁷ This fits with observations in labor economics describing firms' retention of skilled workers as an efficient choice, when comparing retention to potentially significant labor adjustment costs, e.g. hiring and firing costs, and unemployment insurance rate increases.

²⁷ Washington Post, Sept. 29, 2012 GenOn's Potomac River Generating Station was permanently, and "Most of the 120 GenOn employees have accepted transfers or chosen retirement."

Table 7: Descriptive Statistics: Plant and Utility-Level Employment Growth

	Plant-level			Utility-level		
	Treatment	Unmatched Controls	Matched Controls	Treatment	Unmatched Controls	Matched Controls
	(1)	(2)	(3)	(4)	(5)	(6)
Employment growth 1988 - 1994	-2.8% (32)	-2.5% (30)	-1.0% (36)	-2.4% (0.07)	-1.4% (0.03)	-1.8% (0.03)
Employment growth 1995 - 1999	-4.5% (32)	-7.3% (34)	-6.1% (29)	-4.3% (0.06)	-4.3% (0.05)	-4.8% (0.04)
Ave. plants per utility				2.1 (1.6)	2.0 (1.3)	1.9 (1.0)
Number of plants	131	393	131			
Number of utilities	61	216	76	40	40	29

Averages, with standard deviations in parentheses. Note: utility-level statistics are limited to those utilities with at least one Phase I plant, and at least one plant in the selected control group (unmatched or matched).

Table 7 presents descriptive statistics for employment growth and levels, before (1988 – 1994) and during the Phase I period (1995 – 1999), at the plant and utility levels. Consistent with the skilled worker retention story, utilities that have plants in both the treatment and control groups also experience less of a decline in employment, on average, during the Phase I period. Specifically, reductions in employment growth are noticeably less at the utility level for non-Phase I plants, -4.3 percent or -4.8 percent employment growth in the Phase I period compared to employment when utility level effects are ignored, -7.3 percent and -6.1 percent. The summary statistics are consistent with the notion that utilities with multiple plants were able to shift employment among their plants to retain skilled workers and reduce declines in employment growth.

Table 8 presents the difference-in-difference results based on a matched sample with utility-level fixed effects based on plant ownership. Columns (1) and (2) present the pooled and fixed effect estimations based on a sample matched without replacement, while columns (3) and (4) present pooled and fixed effect estimations based on a sample

matched with replacement. As previously mentioned, matching is based on specification #4 described in Table 6. However, results for the policy variable of interest are not sensitive to the matching specification utilized. All fixed effect regressions include a quadratic time trend. We correct our standard errors for heteroskedasticity and within plant autocorrelation in all specifications.

The results differ from the naïve difference-in-difference results in several respects.²⁸ In one of the pooled regressions we find that Phase I plants do not have statistically distinguishable employment (expressed in logs) relative to the matched set of untreated plants; in the second it is significant and positive at the 10 percent level, indicating that Phase I plants have higher employment than non-Phase I plants over the entire study period. Across the four specifications, we find that all plants in the sample (treated and untreated) have lower employment in the post-policy period, which is consistent with what we observed in the naïve case and in Figure 1. However, the post-policy period is only significant at the 10 percent level in one of the pooled specifications. When we interact the dummy variable for Phase I plants with the post-policy-period dummy, we find no statistically distinguishable difference in employment between Phase I and non-Phase I plants in the post policy period. Plant age is now negative and significant in the pooled difference-in-difference regressions that rely on a matched sample, indicating that older plants in both the treated and untreated samples tend to have higher employment.²⁹

²⁸ We ran the regressions excluding two Phase I plants that use oil as their main fuel – all other Phase I plants are coal units. Likewise, we ran the regressions without two Phase I plants identified by Ellerman and Montero (2003) as switching to Powder River Basin low sulfur coal prior to 1990. Our results are not sensitive to either exclusion.

²⁹ Local labor market conditions can affect a plant’s hiring decisions. Because we know the location of each plant, we can capture differences in labor market conditions. However, wages at the local level are likely endogenous to employment decisions. To examine the robustness of our results to differences in wages, we included the natural log of average annual state-level wages for power plants in our regressions. Wages were not significant in the main specifications and do not appear to have any effect on the coefficient estimates for the other regressors.

**Table 8: Matched Difference-in-Difference Results
(Dependent Variable = Log of Employment)**

	Dif-in-Dif with Matched Sample (Without Replacement)		Dif-in-Dif with Matched Sample (With Replacement)	
	Pooled	Utility Level Fixed Effects	Pooled	Utility Level Fixed Effects
	(1)	(2)	(3)	(4)
Phase I Plant	0.003 (0.04)		0.08 * (0.04)	
Phase I Period	-0.05 (0.04)	-0.05 (0.05)	-0.03 (0.05)	-0.05 (0.06)
Phase I Plant x Phase I Period	-0.06 (0.06)	0.01 (0.11)	-0.08 (0.07)	0.01 (0.10)
Plant Age	-0.02 *** (0.001)		-0.02 *** (0.001)	
Constant	5.89 *** (0.05)	5.18 *** (0.05)	5.75 *** (0.05)	5.16 *** (0.06)
Quadratic Time Trend		Yes		Yes
Observations	3,153	3,153	2,580	2,580
Groups	--	129	--	116

*** indicates a coefficient that is significant at the 1 percent level; * indicates significance at the 10 percent level. Standard errors are presented in parentheses.

Using plant-level fixed effects (265 groups) produces similar results to those using plant level fixed effects: neither the post 1994 period nor the interaction term between Phase I plants and the Phase I period are statistically significant. However, while our main results do not appear sensitive to this nuance, we show later in the paper that specifying plant versus utility level fixed effects matters for a number of our follow-on analyses.

Timing of Investment

The first year in which plants had to comply with the aggregate SO₂ cap set by Title IV of the CAAA was 1995. Information on the likely cost of allowances became available through EPA auctions held in 1993 and 1994. In anticipation of the new more stringent standard, it is possible that some plants did not wait until 1995 to adopt measures that reduce SO₂ emissions. In particular, some types of compliance strategies required that

new equipment be installed, which may well have employment implications prior to the first compliance year of the program. Moreover, Swift (2001) reports that plants that installed scrubbers on boilers as a compliance strategy did so in 1994, the year prior to the start of the program. For this reason, we examine whether redefining the Phase I dummy variable to begin in 1994 instead of 1995 allows us to capture employment effects associated with Title IV that are missed by our main specification. We also tested an alternate specification where the Phase I period was defined as 1997-1999. We did this because a number of Phase I plants received bonus allowances in the initial years of the SO₂ cap-and-trade program, which would have loosened the relative stringency of the program. In general, using an alternate definition for the post policy period does not alter the sign or significance of the main results we already presented in Table 8.

Employment Effects over Time

Thus far, we have examined employment effects associated with the Phase I period as a whole. However, it is possible that impacts on employment in the first years of the SO₂ trading program are of a different sign and magnitude than later years.³⁰ This could occur, for instance, if plants expend resources – and hire employees – to install certain types of pollution equipment (e.g. scrubbers or calibration of equipment for the use of low sulfur coal) during the first few years of the program but need only a few employees to monitor ongoing compliance in the subsequent years of the program once these changes have been made. We examine this possibility by interacting the Phase I dummy with separate year dummies for each year (1994 – 1999) of the SO₂ trading program that we evaluate. To be as inclusive as possible we begin with 1994, a year prior to what is included in the main specification to include the installation of scrubbers (recall that plants that installed scrubbers on boilers as a compliance strategy did so in 1994). We present results using both plant-level and utility-level fixed effects.

The results in Table 9 indicate that employment is lower in Phase I plants in each year relative to non-Phase I plants. However, this effect is only statistically significant for

³⁰ We also ran a series of pooled long-difference regressions in which we examined effects of Phase I on employment from 1988 to 1995, 1988 to 1997, and 1988 to 1999. Our main results presented in Table 8 did not change.

1994 when using plant-level fixed effects. When we account for the possibility that utilities may make employment decisions at the utility level, varying how they shift workers between plants in response to market conditions, we find that this effect is no longer significant and that the point estimate is 10 times smaller. Employment also is not statistically different for Phase I plants in any of the subsequent years for the four specifications. Year dummies (not reported in the table) are negative and significant for each year after 1992, indicating a decrease in employment at all plants – Phase I and non-Phase I – throughout most of the 1990s.

**Table 9: Panel Fixed Effect Year-By-Year Difference-in-Difference with Matching
(Dependent Variables = Log Employment)**

	Without Replacement		With Replacement	
	Plant-Level	Utility-Level	Plant-Level	Utility-Level
Phase I Plant x Year 1994	-0.10 * (0.05)	-0.01 (0.10)	-0.13 ** (0.06)	-0.01 (0.11)
Phase I Plant x Year 1995	-0.07 (0.07)	0.01 (0.11)	-0.10 (0.08)	0.01 (0.11)
Phase I Plant x Year 1996	-0.08 (0.07)	-0.00 (0.11)	-0.12 (0.08)	-0.01 (0.11)
Phase I Plant x Year 1997	-0.08 (0.07)	-0.00 (0.11)	-0.13 * (0.08)	-0.02 (0.11)
Phase I Plant x Year 1998	-0.08 (0.08)	0.01 (0.12)	-0.09 (0.10)	0.02 (0.12)
Phase I Plant x Year 1999	-0.06 (0.9)	0.05 (0.13)	-0.10 (0.11)	0.05 (0.13)
Constant	5.14 *** (0.03)	5.15 *** (0.04)	5.12 *** (0.04)	5.12 *** (0.04)
Year Dummies	Yes	Yes	Yes	Yes
Observations	3,153	3,153	2,580	2,580
Groups	264	129	216	116

*** indicates a coefficient that is significant at the 1 percent level; * indicates significance at the 10 percent level. Standard errors are presented in parentheses.

To ensure that our main results - which suggest that during Phase I of the SO₂ trading program Phase I plants did not experience employment losses that were statistically

different from those experienced by non-Phase I plants during the 1995-1999 time period - are robust we perform a series of sensitivity analyses in the following two sections.

7. EMPLOYMENT EFFECTS BY COMPLIANCE OPTION

As previously noted employment effects at the plant-level are expected to vary with the compliance strategy chosen. We attempt to disaggregate the net employment effect by compliance strategy to understand if the small, insignificant effect presented in Table 8 masks larger employment effects for particular types of power plants.

Switching to low-sulfur coal

One of the most common compliance strategies pursued by Phase I plants was switching to low sulfur coal. A total of 37 Phase I plants in our sample made use of this option in 1995 - 1997. To examine whether plants that switched to low sulfur coal as a result of Phase I experienced employment effects noticeably different from the overall average, we re-do the propensity score matching estimation for Phase I plants that switched to low sulfur coal and non-Phase I plants and then re-estimate the main set of difference-in-difference regressions.³¹ We identify Phase I plants that used low-sulfur coal (those that use coal with less than 1.2 pounds of SO₂ per mmBtus, on average) in 1990, 1995, and 1997. We then limit our sample of Phase I plants to those that had used coal with higher sulfur content in 1990 but were using coal with low sulfur content in either 1995 or 1997 (when bonus allowances were no longer available).

We find little evidence of differential employment effects associated with the implementation of Phase I for plants that pursued switching to lower sulfur coal as a compliance strategy (Table 10). However, once again the use of utility-level instead of plant-level fixed effects matters. When only accounting for plant fixed effects it appears that switching to low sulfur coal is associated with a statistically significant (at the 10 percent level) negative employment effect for Phase I plants during the Phase I period. However, Swift (2001) reports that about 67 percent of utilities subject to Phase I

³¹ Propensity score matching still substantially reduces bias in the sample. Prior to matching the mean bias was 57 percent. Matching reduces the bias to 10.1 percent.

switched to low-sulfur coal at a subset of their boilers to create excess allowances that met their compliance needs at other boilers they own. Thus, not accounting for a cross-plant compliance strategy could misestimate the average treatment effect of the SO₂ trading program on employment for coal-switching plants. When we use utility-level fixed effects in the low-sulfur coal regression, we find that the interaction between Phase I Plant and Phase I Period is no longer significant.

Table 10: Switching to Low-Sulfur Coal (Dependent Variable = Log Employment)

	Dif-in-Dif with Matched Sample (Without Replacement)		Dif-in-Dif with Matched Sample (With Replacement)	
	Plant-Level Fixed Effects	Utility-Level Fixed Effects	Plant-Level Fixed Effects	Utility-Level Fixed Effects
Phase I Period	0.06 (0.10)	0.02 (0.12)	0.08 (0.12)	0.04 (0.14)
Phase I Plant x Phase I Period	-0.33 * (0.18)	-0.27 (0.22)	-0.34 * (0.19)	-0.30 (0.24)
Constant	5.12 *** (0.11)	5.14 *** (0.13)	5.14 *** (0.12)	5.16 *** (0.14)
Quadratic Time Trend	Yes	Yes	Yes	Yes
Observations	878	878	806	806
Groups	74	59	68	55

*** indicates a coefficient that is significant at the 1 percent level. Standard errors are presented in parentheses.

Utilities with fewer boilers

As previously mentioned, utilities largely restricted trading of allowances to plants they owned, rarely trading across companies during Phase I (Swift (2001) reports that inter-firm trading in Phase I represented less than 3 percent of total emissions). For utilities that own plants with many boilers the cap-and-trade system still afforded them substantial flexibility with respect to what compliance option they pursued compared with a rate-based emission standard approach. However, Swift (2001) points out that utilities with fewer boilers pursued compliance strategies that were substantially more expensive (i.e., they had far fewer boilers over which to spread the costs of compliance) to avoid trading

with other companies. In particular, several small utilities elected to install scrubbers. In this case, we might worry that plants owned by utilities with fewer boilers could experience differential, potentially negative employment impacts.

We explore this possibility by splitting the data set into two samples: Phase I plants owned by a utility with less than the median number of boilers, and Phase I plants owned by a utility with more than the median number of boilers. (About half of the Phase I plants are owned by a utility with fewer than 14 boilers.) We then re-estimate the matching model, amending it to add the number of boilers at a utility as a predictor of whether a plant is subject to Phase I. We make this modification to the propensity score matching estimation since the split sample is now biased – all Phase I plants owned by utilities with many boilers are in one sample while all Phase I plants owned by utilities with few boilers are in the other - to ensure we match to an appropriate non-Phase I plant. We then re-estimate our fixed effect regressions.

Table 11: Panel Fixed Effects for Plants Split by Number of Boilers
(Dependent Variable= Log of Employment)

	Plants Owned by Utilities With Fewer Than 14 Boilers		Plants Owned by Utilities With 14 or More Boilers	
	Plant-Level Fixed Effects	Utility-Level Fixed Effects	Plant-Level Fixed Effects	Utility-Level Fixed Effects
Phase I Period	0.02 (0.06)	-0.05 (0.06)	-0.03 (0.03)	-0.07 * (0.04)
Phase I Plant x Phase I Period	-0.13 * (0.08)	0.001 (0.12)	0.02 (0.10)	0.04 (0.11)
Constant	5.44 *** (0.07)	5.45 *** (0.09)	4.88 *** (0.05)	4.91 *** (0.05)
Quadratic Time Trend	Yes	Yes	Yes	Yes
Observations	1,620	1,620	1,536	1,536
Groups	136	56	129	96

*** indicates a coefficient that is significant at the 1 percent level; ** indicates significance at the 5 percent level, and * indicates significance at the 10 percent level. Standard errors are presented in parentheses.

When we use plant fixed effects we find some evidence for the hypothesis that Phase I plants owned by utilities with fewer boilers experience statistically significant, negative employment during the Phase I period relative to non-Phase I plants (Table 11). However, when we account for utility-level fixed effects to account for the fact that utilities often make firm-wide decisions that can potentially result in inter-plant shifting of employment we find that this result disappears. In fact, the point estimate changes from -0.13 using plant-level fixed effects to 0.001 with utility-level fixed effects. This illustrates how imposing a plant-level decision model may miss ways in which a utility is able to mitigate plant-specific shocks by spreading them over multiple plants (i.e., it is efficient for firms to find ways to retain skilled workers).

The interaction between Phase I plants and the post-policy period is not statistically significant for plants owned by utilities with more than the median number of boilers in either specification. However, the model picks up on a statistically significant decrease in employment for all plants in the post policy period when we use utility-level fixed effects. The results for both samples continue to hold when we define the post-policy period as beginning in 1994. However, the results disappear when we use a matching strategy that allows for replacement.

Compensation and Substitution Units

Another strategy that was frequently pursued was to bring a compensation or substitution unit into the Phase I program. Our expectation is that these plants were selected by the utilities for early inclusion in the SO₂ cap-and-trade because they have relatively low marginal abatement costs. Thirty plants in our dataset entered Phase I purely as compensation or substitution units (substitution units were also utilized by plants already subject to Phase I). Our main results do not change when we exclude these 30 plants from the estimation.³²

Retirement

Recall that utilities are not allowed to decrease net generation to comply, so any drop in

³² Results available upon request.

generation due to retirement of units at one plant must be compensated for within the same utility (e.g. installing a new boiler with the same or larger capacity or by bringing a compensation unit into Phase I). Thus, while we expect the net effect on employment to be close to zero for the utility as a whole, the effect of these changes on plant-level employment is likely negative for the plant that shuts down and positive for another plant owned by the same utility. We examine the sensitivity of our main results by dropping the five Phase I plants (and their respective matches) in our dataset that retired during the Phase I period.³³ Using utility-level fixed effects, Table 12 shows negative and significant employment effects for all plants still in operation during the Phase I period. However, as expected, Phase I plants still in operation experience positive and significant employment effects during the Phase I period once retired units are omitted.³⁴

Table 12: Utility-Level Fixed Effects Without Retirement Units (Dependent Variable = Log Employment)

	Without Replacement	With Replacement
Phase I Period	-0.11 *** (0.04)	-0.11 * (0.06)
Phase I Plant x Phase I Period	0.16 ** (0.07)	0.13 * (0.08)
Constant	5.14 *** (0.04)	5.16 *** (0.04)
Quadratic Time Trend	Yes	Yes
Observations	3,153	2,580
Groups	264	216

8. ACCOUNTING FOR NO_x STANDARDS

Because the Title IV NO_x standards overlap with the SO₂ trading program, we examine the robustness of the difference-in-difference estimation to implementation of the NO_x standard. Unlike for SO₂, NO_x was controlled using a traditional rate-based standard that required coal-fired units with specific types of boilers to install a low-NO_x burner

³³ We identified 10 plants that were retired over the study period, but only five of them had enough data to be included in the final dataset.

³⁴ When we use plant-specific fixed effects the interaction between Phase I and the Phase I period remains statistically insignificant.

technology. When this type of technology was unavailable for a boiler type, they were exempt from the first phase of the standard.³⁵ Utilities were allowed to average emissions across their units but not allowed to average (or trade) across utilities.

Qualifying boilers that were part of Phase I or had opted in as substitution units were required to comply with the NOx standard in 1996. In addition, plants subject to Phase II of the NOx standard, which would be in place in 2000, could opt to comply with the NOx standards early. About half of the units eligible for early opt-in chose to adopt Phase I NOx standards (Swift 2001). In return for meeting the standard they were not required to make further reductions to meet the more stringent Phase II NOx emission requirements until 2008.

In our sample, 73 percent of plants subject to Phase I of the SO₂ trading program were part of or opted into Phase I of the NOx program, while 31 percent of non-Phase I plants in terms of SO₂ were part of Phase I of the NOx program. We make use of the variation in timing of the two requirements and in whether plants were subject to one or both Phase I requirements to separately identify the effects of the NOx standard on employment. The panel estimation is now a difference-in-difference-in-differences:

$$\ln EMP_{pt} = \beta_2 YR9599_t + \beta_3 PHASE1 * YR9599_{pt} + \beta_4 YR9699_t + \beta_5 NOx * YR9699_{pt} + \beta_3 PHASE1 * NOx * YR9599_{pt} + a_p + d_t + u_{pt} \quad (4)$$

where *NOx* is a dummy variable equal to one when a plant is subject to the Title IV NOx standard and zero otherwise; and *YR9699* is a dummy variable equal to one for the NOx standard time period, 1996 – 1999, and zero otherwise.

Table 13 shows that our previous results remain unchanged when accounting for NOx standards: We still find no evidence of a differential employment effect during Phase I for plants participating in the SO₂ cap-and-trade program. This result is not sensitive to the way in which we define the post-policy period for Phase I of the SO₂ trading program

³⁵ These controls were expected to cost far less than those required to meet the SO₂ standard and applied mostly to units in Western states.

(i.e., post-1993, post-1994 or post 1996) when using utility-level fixed effects.³⁶ We find that plants subject to Phase I of the NOx standard reduced employment by significantly less in the post-policy period than other plants in one of the two specifications that use utility-level fixed effects. In fact, the net employment effect for plants subject to the NOx standards in the post policy period is positive in all of the fixed effects regressions.

Table 13: Controlling for NOx Standards (Dependent Variable= Log Employment)

	Dif-in-Dif with Matched Sample (Without Replacement)		Dif-in-Dif with Matched Sample (With Replacement)	
	Plant-Level Fixed Effects	Utility-Level Fixed Effects	Plant-Level Fixed Effects	Utility-Level Fixed Effects
	(1)	(2)	(3)	(4)
NOx Policy Period (1996-1999)	-0.06 * (0.04)	-0.16 ** (0.05)	-0.09 * (0.05)	-0.13 ** (0.07)
NOx Phase I x NOx Policy Period	0.07 (0.05)	0.20 ** (0.08)	0.11 (0.07)	0.15 * (0.09)
SO ₂ Policy Period (1995-1999)	0.01 (0.04)	-0.2 (0.05)	0.04 (0.05)	-0.02 (0.06)
SO ₂ Phase I x SO ₂ Policy Period	-0.12 (0.13)	-0.13 (0.24)	-0.12 (0.13)	-0.12 (0.23)
SO ₂ Phase I x NOx Phase I x SO ₂ Period	0.06 (0.14)	0.14 (0.24)	0.02 (0.15)	0.13 (0.24)
Constant	5.17 *** (0.05)	5.20 *** (0.05)	5.15 *** (0.05)	5.17 *** (0.06)
Quadratic Time Trend	Yes	Yes	Yes	Yes
Observations	3,153	3,153	2,580	2,580
Groups	264	129	216	116

*** indicates a coefficient that is significant at the 1 percent level; * indicates significance at the 10 percent level. Standard errors are presented in parentheses.

³⁶ The interaction term between post policy and Phase I for SO₂ is statistically significant at the 10 percent level when the post-policy period is defined as 1994-1999, we use plant fixed effects, and we match without replacement.

9. CONCLUSION

While public discourse often asserts that environmental regulations have significantly large negative employment effects we find little evidence of this for fossil fuel fired power plants subject to Phase I of the SO₂ trading program. This finding is robust to matching estimation strategy, plant or utility-level fixed effects, and the way that the Phase I period is defined. An examination of employment effects associated with particular compliance strategies demonstrates the importance of accounting for utility-level fixed effects. Utilities largely adhered to intra-utility approaches to compliance and avoided trading with other utilities. For instance, we find significant negative employment effects for Phase I plants that switched to low sulfur coal in the post-policy period when we use plant fixed effects. However, utilities reportedly chose to switch to low sulfur coal at a subset of plants to generate excess allowances that met their compliance needs at other Phase I plants. Thus, not accounting for a cross-plant compliance strategy could misestimate the average treatment effect. When we use utility-level fixed effects in this case, we find that the negative employment effect is no longer statistically significant and that the point estimate is 10 times smaller. Finally, we control for the implementation of a NO_x rate-based standard that overlaps to some degree with Phase I of the SO₂ trading program. We again find that employment effects associated with the SO₂ program are insignificant even after controlling for NO_x compliance.

REFERENCES

Berman, E., and L. Bui, (2001a) Environmental Regulation and Labor Demand: Evidence from the South Coast Air Basin. *Journal of Public Economics*, 79: 265 – 295.

Berman E., and L. Bui, (2001b) Environmental regulation and productivity: evidence from oil refineries. *Review of Economics and Statistics* 83: 498–510.

Bertrand, M, E. Duflo, and S. Mullainathan, (2004) How Much Should We Trust Difference-in-Differences Estimates? *Quarterly Journal of Economics* 119(1): 249-275.

Boyd G., and J. McClelland (1999) The Impact of Environmental Constraints on Productivity Improvement in Integrated Paper Plants. *Journal of Environmental Economics and Management* 38:121–142.

Burtraw, D., A. Krupnick, E. Mansur, D. Austin and D. Farrell. 1998. Costs and Benefits of Reducing Air Pollutants Related to Acid Rain. *Contemporary Economic Policy* 16(4): 379-400.

Burtraw, D., and K. Palmer. 2004. SO₂ Cap-and-Trade Program in the United States: A ‘Living Legend’ of Market Effectiveness. In *Choosing Environmental Policy: Comparing Instruments and Outcomes in the United States and Europe*, ed. W. Harrington, R. Morgenstern, and T. Sterner. Washington D.C.: Resources for the Future.

Carlson, C., D. Burtraw, M. Cropper, and K. Palmer. 2000. SO₂ Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy* 108(6): 1292-326.

Chan, G., R. Stavins, R. Stowe, and R. Sweeney (2012). The SO₂ Allowance Trading System and the Clean Air Act Amendments of 1990: Reflections on Twenty Years of Policy Innovation. NBER Working Paper 17845.

Cole, M., and R. Elliott. (2007) Do Environmental Regulations Cost Jobs? An Industry-Level Analysis of the UK. *The B.E. Journal of Economic Analysis and Policy* vol 7. issue 1 (Topics).

Ellerman, A., P. Joskow, R. Schmalensee, J. Montero, and E. Bailey. 2000. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge: Cambridge University Press.

Ellerman, D., and P. Montero (1998). The Declining Trend in Sulfur Dioxide Emissions: Implications for Allowance Prices. *Journal of Environmental Economics and Management* 36: 26-45.

Energy Information Administration (1994). *Electric Utility Phase I Acid Rain Compliance Strategies for the Clean Air Act Amendments of 1990*, DOE/EIA-0582. Washington, DC.

Energy Information Administration (1997). *The Effects of Title IV of the Clean Air Act Amendments of 1990 on Electric Utilities: An Update*. DOE/EIA-0582(97). Washington, DC.

Fabrizio, K., N. Rose, and C. Wolfram (2007) Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency. *American Economic Review* 97(4): 1250 – 1276.

Färe, R., S. Grosskopf and C. Pasurka Jr. (1986). Effects on Relative Efficiency in Electric Power Generation Due to Environmental Controls. *Resources and Energy* 8 (2): 167-184.

Fowlie, M., S. Holland, and E. Mansur (2012). What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program. *American Economic Review* 102(2): 965-93.

Gray, W. and R. Shadbegian (1998). Environmental Regulation, Investment Timing, and Technology Choice. *Journal of Industrial Economics* 46: 235-56.

Gray, W. and R. Shadbegian (2003). Plant Vintage, Technology, and Environmental Regulation. *Journal of Environmental Economics and Management*, 384-402.

Gray, W., R. Shadbegian, C. Wang, and M. Cebi (2013). Do EPA Regulations Affect Labor Demand? Evidence from the Pulp and Paper Industry, Working paper.

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.

Harrington, W., R. Morgenstern, and P. Nelson (2000) On the Accuracy of Regulatory Cost Estimates. *Journal of Policy Analysis and Management* 19(2): 297-322.

Heckman, J., H. Ichimura, J. Smith, and P. Todd (1998). Characterizing selection bias using experimental data. *Econometrica* 66.

Heckman, J., H. Ichimura, and P. Todd (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Review of Economic Studies* 64: 605–654.

Keohane, N. (2003). What Did the Market Buy? Cost Savings under the U.S. Tradable Permits Market Program for Sulfur Dioxide. Yale Center for Environmental Law and Policy Working Paper ES #33.

Laplante, B., and P. Rilstone (1996). Environmental Inspections and Emissions of the Pulp and Paper Industry in Quebec. *Journal of Environmental Economics and Management* 31: 19-36.

- List, J., D. Millimet, P. Fredriksson, and W. McHone (2003). Effects of environmental regulations on manufacturing plant births: Evidence from a propensity score matching estimator. *Review of Economics and Statistics* 56: 944–952.
- Magat, W., and W. Viscusi (1990). Effectiveness of the EPA's Regulatory Enforcement: The Case of Industrial Effluent Standards. *Journal of Law and Economics* 33: 331-360.
- Morgenstern, R., W. Pizer, and J. Shih (2002). Jobs Versus the Environment: An Industry-Level Perspective. *Journal of Environmental Economics and Management* 43: 412–436.
- Popp, D. (2003). Pollution Control Innovations and the Clean Air Act of 1990. *Journal of Policy Analysis and Management* 22(4): 641-660.
- Rosenbaum, P., and D. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70: 41–55.
- Rubin, D. (2008). For objective causal inference, design trumps analysis. *The Annals of Applied Statistics* 2: 808–840.
- Shadbegian, R., and W. Gray (2003). "What Determines Environmental Performance at Paper Mills? The Roles of Abatement Spending, Regulation, and Efficiency" *Topics in Economic Analysis & Policy*.
- Shadbegian, R., and W. Gray (2005). Pollution Abatement Expenditures and Plant-Level Productivity: A Production Function Approach. *Ecological Economics* 54: 196-208.
- Shadbegian, R., and W. Gray (2006). Assessing Multi-Dimensional Performance: Environmental and Economic Outcomes. *Journal of Productivity Analysis* 26: 213-234.
- Stuart, E., and D. Rubin (2008). Best Practices in Quasi-Experimental Designs: Matching Methods for Causal Inference in *Best Practice in Quantitative Methods*. Ed., J. Osborne. CA: Sage Publications, 155-176.
- U.S. Environmental Protection Agency (1997). *1996 Compliance Report: Acid Rain Program*.
- U.S. Environmental Protection Agency (1996). *1995 Compliance Results: Acid Rain Program*. EPA/430-R-96-012. July.
- Washington Post (2012). GenOn power plant in Alexandria is set to close. September 29. http://www.washingtonpost.com/local/genon-power-plant-in-alexandria-is-set-to-close/2012/09/29/daa355ea-08d7-11e2-858a-5311df86ab04_story.html

APPENDIX A: ATTEMPTING TO DIFFERENTIATE EFFECTS BY REGULATORY FORM AND STRINGENCY

The Title IV SO₂ program allowed plants more flexibility in how they comply relative to a rate-based standard. However, it also aimed to significantly reduce SO₂ emissions relative to the prior regulatory approach. In this specification, we attempt to distinguish employment effects associated with changes in relative stringency of SO₂ regulation in the Phase I period from changes in the form of regulation. We do this by interacting the difference-in-difference measure, *Phase I*YR95*, with measures of regulatory form and stringency.

STRINGENCY is defined as emissions per unit of heat input (in 1985) divided by a measure of expected regulatory stringency in 1995, the first year in which an aggregate cap for SO₂ is set. For Phase I plants, the expected stringency in 1995 is the rate at which allowances were allocated to units, 2.5 pounds of SO₂ per million Btus of heat input. For plants in the untreated group defining expected stringency in 1995 is a bit more challenging. We do not set the ratio of relative stringency equal to one because of evidence that state regulations and changes in NSPS standards may imply some reductions in emission rates among Phase II plants relative to 1985. Likewise, we do not use emissions per unit of heat input in 1995 to define expected stringency because of the possibility that Phase II plants are already beginning to anticipate the stricter SO₂ cap in 2000, to which they will be subject. Instead, we lag emissions per unit of heat input several years – to 1992 - such that the measure captures state and NSPS regulations but is free from possible endogeneity in emission rates leading up to Phase II. We interact this variable with the dummy variable for the post-policy period, as well as the interaction between post-policy and the Phase I dummy.

To capture the change in regulatory form for plants participating in Phase I of the SO₂ trading program, we follow Kerr and Newell (2003). Prior to the Title IV program plants essentially only had one technology option available for compliance, the installation of a scrubber. Establishment of the cap-and-trade program allowed plants to select

compliance strategies based on marginal abatement costs. Economic theory suggests that a plant that can lower its emissions at an abatement cost less than the allowance price will do so and sell its extra allowances. On the other hand, a plant that has marginal abatement costs higher than the allowance price will not invest in the expensive compliance strategy and instead purchase allowances. Kerr and Newell (2003) estimate the predicted probability that a plant is a net seller of allowances based on a set of variables associated with compliance costs. We take this same approach.

We estimate a probit where the dummy variable equals one when a plant is a net banker or seller (meaning they incur costs to reduce emissions) in 1997 and zero otherwise. We choose 1997 because of bonus allowances in the early years of the program that may obfuscate the true marginal abatement cost of the plant. The independent variables included in this regression are: plant age, overall production cost per megawatt hour, amount of labor utilized at the plant in hours (in logs), whether a plant is located in a state with higher than median coal production in 1988, distance to the Powder River Basin (PRB) and its square, whether a plant is located in a state with pre-existing SO₂ limits (i.e., Minnesota, new Hampshire, or Wisconsin), average heat rate in 1985, and the sulfur content of fuels used in 1985.³⁷

Older plants may find it more costly to reduce emissions because of outdated equipment, making them less likely to sell allowances (more likely to purchase). We expect that plants with higher per megawatt hour production costs also face higher abatement costs, making them less likely to sell allowances. We include labor hours at a plant to identify base versus peak load plants, hypothesizing that base load plants face different abatement costs than peak plants, though we can't predict the sign.³⁸ Plants that are located in states with higher than median coal production may face pressure to continue to use coal, which would eliminate a potential lower-cost strategy, making them less likely to sell

³⁷ We also considered including a measure of the plant's complexity, for example by including the number of boilers or fuels used at a plant; whether a plant had a pre-existing scrubber; and whether it was located in a county out of attainment for SO₂. However, these variables were not significant, nor did they significantly improve the fit of the probit regression.

³⁸ In an alternate specification, we substituted nameplate capacity and found it had little effect on our results.

allowances. Plants located closer to the Powder River Basin likely face lower abatement costs if they switch to low-sulfur coal, making them more likely to sell allowances. We hypothesize that plants located in states with SO₂ limits prior to the Title IV trading program are more likely to sell allowances since abatement to meet the state standards occurred after the baseline year used to allocate allowances (1985-1987), making it more likely that a plant will have more allowances than it needs. If plants that use dirtier coal have to do more to come into compliance they are less likely to sell allowances, while less efficient power plants – measured by the average heat rate – may have more low-hanging fruit with regard to abatement, making them more likely to sell.

Table A1: Likelihood of Selling Allowances

Variables	Coefficient (Standard Error)
Median Coal Production	-1.30 ** (0.52)
Distance to PRB	0.01 (0.01)
Distance to PRB ²	-0.00 (0.00)
State SO ₂ Regulations	0.99 (1.19)
Plant Age	-0.05 ** (0.02)
Labor Hours (logs)	-0.38 (0.37)
Sulfur Content of Fuel	-0.29 (0.26)
Production Cost per mwh	-0.001 (0.001)
Average Heat Rate	0.00 (0.00)
Constant	3.37 (6.52)

** indicates significance at the 5 percent level. Standard errors are presented in parentheses.

Two variables are significant at the 5 percent level with regard to the likelihood of selling SO₂ allowances: median coal production and plant age (Table A1). Plants in states with higher than median coal production and older plants are both, as expected, less likely to sell allowances. While the remaining variables are insignificant, they have the expected

sign. We use the coefficient estimates to calculate the predicted probability that a Phase I plant will be a net seller, *NET SELL*. We then interact this with the dummy variable for the post-policy period. Regulatory form does not change for the non-treated group of plants.

The difference-in-difference panel estimation now is:

$$\ln EMP_{pt} = \beta_4 NETSELL * YR95_{pt} + \beta_5 STRINGENCY * YR95_{pt} + \beta_6 STRINGENCY * PHASE1 * YR95_{pt} + a_p + t + t^2 + u_{pt} \quad (7)$$

Consistent with our main results, employment in the post-policy period is lower than in the previous period, though it is not statistically significant. Our measures for regulatory stringency and form for Phase I plants in the post-policy period are also not significant using either plant or utility-level fixed effects, or matching with or without replacement. This is not surprising given that our goal is to parse the insignificant employment effect associated with Phase I plants in the post-policy period into its component parts. Our ability to predict who is likely to be a net seller of allowances is also limited because of the poor performance in the first stage where only two of the hypothesized variables were statistically significant.³⁹

³⁹ We explored adding the number of boilers at the plant as well as the number of boilers owned by the same utility as regressors. Neither improved our ability to predict which plants were likely to be net sellers.