The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing*

Michael Greenstone
MIT
and NBER
mgreenst@mit.edu

John A. List
University of Chicago
and NBER
jlist@uchicago.edu

Chad Syverson
University of Chicago
Booth School of Business
and NBER
chad.syverson@chicagobooth.edu

September 2012

Abstract

The economic costs of environmental regulations have been widely debated since the U.S. began to restrict pollution emissions more than four decades ago. Using detailed production data from nearly 1.2 million plant observations drawn from the 1972-1993 Annual Survey of Manufactures, we estimate the effects of air quality regulations on manufacturing plants' total factor productivity (TFP) levels. We find that among surviving polluting plants, stricter air quality regulations are associated with a roughly 2.6 percent decline in TFP. The regulations governing ozone have particularly large negative effects on productivity, though effects are also evident among particulates and sulfur dioxide emitters. Carbon monoxide regulations, on the other hand, appear to increase measured TFP, especially among refineries. The application of corrections for the confounding of price increases and output declines and sample selection on survival produce a 4.8 percent estimated decline in TFP for polluting plants in regulated areas. This corresponds to an annual economic cost from the regulation of manufacturing plants of roughly $21 billion, which is about 8.8 percent of manufacturing sector profits in this period.

* We thank participants at several seminars, Lynn Riggs, and Frank Limehouse for helpful comments. Alex Verkhivker provided able research assistance. The research in this paper was conducted while Syverson was a Special Sworn Status researcher of the U.S. Census Bureau at the Chicago Census Research Data Center. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to insure that no confidential data are revealed. Contact information: Greenstone: MIT Department of Economics, E52-359, 50 Memorial Drive, Cambridge, MA 02142-1347; List: Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637; Syverson: University of Chicago Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637.

I. Introduction

The economic costs of environmental regulations have been widely debated since the U.S. began to restrict pollution emissions more than four decades ago through the Clean Air and Water Acts. The conventional wisdom is that stricter environmental standards raise polluting firms' costs of production, which weakens U.S. firms' position in international markets and raises the prices that consumers face. On the one hand, the decline in U.S. manufacturing employment from 18 million (25.3 percent of total US employment) in 1970 to 12 million (9.0 percent of total employment) in 2012 mirrors the introduction and expansion of U.S. environmental policy. On the other hand, Porter (1991) argues that more stringent regulations enhance productivity growth by causing firms to rationalize their operations.

In addition to being an important area of research, this issue is of considerable interest to policymakers. In the last few years, several lawmakers have argued that environmental regulations are "job killers". This contention is not new, as a few decades ago environmental considerations played a major role in the NAFTA debate; opponents predicted that NAFTA would induce U.S. and Canadian firms to move their operations to Mexico, where they could better compete on the global market due to lax local environmental regulations.[1] Parallel arguments on both sides can be found in recent European negotiations, where competitiveness issues played a prominent role in discussions of differences in environmental regulations among member states in the context of the Single European Act. Part of the reason that these debates have raged on is that there is a paucity of conclusive empirical evidence.

This paper helps to fill this gap by making uses of detailed production data from nearly 1.2 million plant observations from the 1972-1993 Annual Survey of Manufactures to investigate the economic costs of air quality regulations. Following the passage of the 1970 Clean Air Act Amendments, the Environmental Protection Agency (EPA) established separate national ambient air quality standards—a minimum level of air quality that all counties are required to meet—for four criteria pollutants: carbon monoxide (CO), tropospheric ozone (O_3), sulfur dioxide (SO_2), and total suspended particulates (TSPs). As a part of this legislation, every U.S. county receives annual nonattainment or attainment designations for each of the four pollutants. The nonattainment designation is reserved for counties whose air contains concentrations of the

[1] Ross Perot famously argued that "If NAFTA passes you'll hear a flushing sound of millions of American jobs going south" when debating Al Gore on CNN's Larry King Show. Vice President Al Gore's response was a broad appeal based on the Porter (1991) hypothesis.

relevant pollutant that exceed the federal standard. Emitters of the regulated pollutant in nonattainment counties are subject to greater regulatory oversight than emitters in attainment counties. Non-polluters are free from regulation in both categories of counties.

We find that among surviving plants in heavily polluting industries, the nonattainment designation is associated with a 2.6 percent decline in measured total factor productivity (TFP) among plants that emit the targeted pollutants. In plain English, this means that regulated plants' output declined by 2.6 percent after holding constant their inputs (i.e., labor, capital, and materials). The regulations governing ozone have particularly large negative effects on productivity, though negative effects are also evident among emitters of particulates and sulfur dioxide. Carbon monoxide nonattainment, on the other hand, appears to actually *increase* measured TFP, especially among refineries.

The estimated decline in TFP of 2.6 percent likely understates the true loss in output due to the nonattainment designations for at least three reasons. First, the estimates are based on comparisons between emitters and non-emitters. These categories are imprecise, however, and it is probable that some of the non-emitter group was also regulated by the Clean Air Act, albeit less intensively. Second, data constraints require us to use revenue in estimating TFP, but this is the product of prices, which may increase due to the costs imposed by regulations, and the conceptually correct outcome of physical output. Third, the estimates are based on the plants that continue to operate, yet it seems reasonable to assume that the plants most harmed by the regulations shut down.

The application of corrections for the confounding of price increases and output declines and sample selection on survival produces a much larger estimate: TFP declines by 4.8 percent for polluting plants in nonattainment counties. This corresponds to an annual economic cost from the regulation of manufacturing plants of roughly $21 billion in 2010 dollars. This translates into a loss of more than $450 billion over the studied period.

The paper makes several contributions. The first is that by measuring the impact on productivity—the amount of output obtained from a given set of inputs—we believe that this is the first study to obtain an estimate of the regulation's *economic costs* in the full manufacturing sector.[2] Our focus on productivity effects speaks to the *efficiency* of the manufacturing sector; that is, given the inputs being used, do environmental regulations change how effectively the sector converts these inputs into outputs. Consequently, the results have a clearer economic interpretation than the finding that the nonattainment designations are associated with reductions in employment, investment, shipments in the manufacturing sector (Henderson 1996; Becker and Henderson 2000; Greenstone 2002; and Walker 2012). Furthermore, these estimates along with estimates of the costs borne by workers in these firms (Walker 2012) can be contrasted with the growing literature on the benefits of cleaner air to conduct a cost-benefit analysis.[3]

The second contribution is that the paper uses the principal instruments of the Clean Air Act Amendments (CAAAs), the pollutant-specific, county-level attainment/nonattainment designations as its measures of air quality regulation. These four designations are the "law of the land" and capture the regional and industry variation that Congress imposed with this legislation.[4] In fact, these designations govern the writing and enforcement of the plant-specific regulations that restrict the behavior of polluters. Moreover, the simultaneous evaluation of each pollution-specific regulation is important, because many plants emitted multiple pollutants and many counties were designated in nonattainment for multiple pollutants.

Furthermore, there is spatial, temporal, and industry-based variation in the impact of the regulations associated with nonattainment designations which allows for the estimation of statistical models that control for several likely confounders. The analysis exploits variation across plants in different industries in the same county and within the same industry across counties allowing for adjustment for location-specific shocks to TFP and shocks common to plants in the same industry. Additionally the panel structure of the data means that the TFP effects are identified from plants in counties that experience a change in nonattainment status; thus, the estimates are based on comparisons of plants in periods when they face the nonattainment regulations to periods when they are free from them.

The third contribution is that this paper builds on the literature of the past two decades

[2] Gollop and Roberts (1983) estimates the economic costs of SO_2 regulation in the utility sector during the 1970s. They estimate the impact of a different feature of the Clean Air Act Amendments than nonattainment designations. Also, see Ryan (2011) for estimates of the costs of the 1990 CAAAs in the Portland cement industry.

[3] Some examples of the extensive literature on the benefits of clean air include Chay and Greenstone (2003 and 2005), Currie and Neidell (2005), Lleras-Muney (2010), and Deschenes, Greenstone, and Shapiro (2012)

[4] A few states and localities (e.g., California) have imposed regulations that are stricter than the federal ones. Any regulations over and above the federally mandated ones are unobserved variables in the subsequent analysis.

that has sought to explain differences in producers' TFP levels (see Syverson 2011 for a survey). Ultimately, productivity determines a nation's living standards, and our understanding of its micro determinants remains in its nascent stages. This paper demonstrates that government policy can play an important role in shaping TFP. To date, we have little credible empirical evidence on the magnitude of such policy effects, with our study providing among the first such estimates for environmental regulations.

The remainder of our study is organized as follows. The next section briefly lays out the conceptual framework that helps to structure the subsequent analysis. Section III reviews the CAAA attainment designation process and its implications for producers. Section IV describes the data. Section V presents the empirical specification, and Section VI reports the benchmark results. Section VII discusses potential reasons why our estimates in Section VI are likely underestimates of the true productivity losses suffered by plants in polluting industries when their county is designated nonattainment for a pollutant, and estimates correction factors to account for these issues. We further interpret our results and conclude in Section VIII.

II. Conceptual Framework

In this section we briefly frame our conceptual view of how environmental regulations might affect a manufacturer's productivity level and how we would measure such effects. The model motivates the empirical models and provides a lens to interpret the results.

We begin by assuming that a manufacturing plant has a Cobb-Douglas production function:

$$Q = A\tilde{L}^{\alpha}\tilde{K}^{1-\alpha},$$

where Q is the plant's output, A is a Hicks-neutral technology shifter, and $\tilde{L}$ and $\tilde{K}$ are, respectively, labor and capital inputs. The only nonstandard feature of our assumed production function is that the labor and capital inputs, $\tilde{L}$ and $\tilde{K}$ are *production-effective* labor and capital—that is, the quantity of each input that actually is used in the production of output—rather than the *observed* labor and capital at the plant. Production-effective and observed inputs are related to one another; we assume the former is proportional to the latter, with the factor of proportionality allowed to vary between the two inputs. Thus,

$$\tilde{L} = \lambda_L L \text{ and } \tilde{K} = \lambda_K K,$$

where L and K are the observed labor and capital inputs at the plant and λ_L and λ_K are the factors of proportionality that link observed to production-effective inputs.

The introduction of production-effective inputs captures the notion that plants that fall under more stringent environmental regulations need to employ inputs that are necessary to meet regulatory requirements but that are potentially not useful for producing the plants' commercial outputs. Indeed, the plant-specific requirements under the Clean Air Act in this period were generally of the "command and control" variety that involved the EPA dictating the installation of particular pollution abatement technologies (rather than imposing emissions limits and allowing plants to achieve them in whatever way they found most efficient). Thus, for example, in order to meet federal air quality standards a manufacturing plant may need to install scrubbing or gas reclamation equipment. This equipment is part of the plant's measured capital stock, but in itself is neither necessary nor useful for producing the plant's commercial output. A labor-input example of the same concept is the hiring of an environmental compliance officer for the plant.

As reflected in our proportionality assumption, we assume these additional regulatory-compliance inputs scale up with the size of the plant. While the strict proportionality assumption is made for analytical convenience and is unnecessary to make our conceptual point, we believe it is quite plausible that larger plants require more regulatory-compliance inputs. For example, larger plants need more emissions-cleaning capital and multiple employees focusing on environmental compliance because it is the larger plants that face greater regulatory scrutiny.

In this framework, the introduction of more stringent regulatory requirements—namely, a nonattainment designation for an emitting plant's county—can be interpreted as a decrease in λ_L and/or λ_K. When more compliance-related inputs are necessary, there is a larger gap between a plant's observed and production-effective inputs.[5] The productivity effects of the regulation can be seen by substituting the expressions for production-effective labor and capital into the production function:

$$Q = A(\lambda_L L)^{\alpha}(\lambda_K K)^{1-\alpha} = A{\lambda_L}^{\alpha}{\lambda_K}^{1-\alpha}L^{\alpha}K^{1-\alpha}.$$

The plant's measured TFP is its output divided by its observed inputs (weighted appropriately):

[5] Another possible channel through which air quality regulation can affect TFP besides the compliance-related input mechanism discussed here is New Source Review (NSR). Under NSR, if an existing plant makes "significant" changes to its operations, the entire plant falls under regulations for new plants. The way this channel would have TFP effects is by forcing plants to use suboptimal input mixes in order to avoid NSR. Our conceptual framework embeds this mechanism; one can interpret λ_L and λ_K in this case as measuring the difference between the observed and the optimal input levels.

$$TFP = \frac{Q}{L^{\alpha}K^{1-\alpha}} = \frac{A\lambda_L^{\alpha}\lambda_K^{1-\alpha}L^{\alpha}K^{1-\alpha}}{L^{\alpha}K^{1-\alpha}} = A\lambda_L^{\alpha}\lambda_K^{1-\alpha}.$$

This expression makes obvious the effect of regulation on TFP. Decreases in λ_L and λ_K driven by an increased need for compliance-related inputs are inward shifters of the plant's production function. In other words, the amount of output the plant obtains per unit of observed input—its TFP level—decreases. The greater the amount of compliance-related inputs (the larger the drop in λ_L and λ_K), the larger the observed decline in plant TFP.

Given the production function above, cost minimization implies the plant's marginal cost is

$$MC = \frac{1}{A\lambda_L^{\alpha}\lambda_K^{1-\alpha}}\phi w^{\alpha}r^{1-\alpha}.$$

(ϕ is a constant that depends on parameter α.) Marginal cost rises as λ_L and λ_K fall, so regulations requiring more pollution abatement-related inputs increase the plant's marginal cost. Although we do not model the output market, with market power, higher marginal cost results in an increase in the price of the plant's output. This price effect of regulation is important in our empirical work below, because we cannot directly measure the output quantity Q in the theoretical TFP measure above. We must instead measure output using revenue. Because revenue is the product of Q and price, our estimates of regulation's TFP impacts reflect not just the TFP effect elucidated above, but also the price effect. The size of the price effect—that is, the relationship between price and marginal cost—depends on the nature of demand for the plant's product. We measure this markup relationship below to derive a correction that turns our estimated revenue-based TFP measures into quantity-based TFP metrics, consistent with our model.

III. The CAAA as the Basis of a Research Design

A. Background on the CAAAs and Their Enforcement

The 1970 CAAA marked the federal government's ambitious entry into the business of restricting the emission of pollutants into the air. It required that all states meet national ambient air quality standards (NAAQS) for certain criteria air pollutants—carbon monoxide, sulfur dioxide, total suspended particulates[6], and ozone[7] (other pollutants, such as lead, have subsequently been added to the list). To do so, states with air quality exceeding the federal guidelines were required to submit a State Implementation Plan (SIP) that detailed their plans to bring violating areas into compliance. Given the amount of confusion, and the inadequate resources to carry-out these plans, many areas of the country had failed to meet the standards by the 1975 deadline.

Due to this lack of progress, Congress passed the 1977 CAAA. The 1977 CAAA stipulated that starting in 1978 every county in the U.S. was to be designated annually as being in-attainment or out-of-attainment (nonattainment) of NAAQS. A county's attainment status was to be determined with respect to each of the criteria air pollutants. If a county is not in attainment of the federal standard with respect to one of these pollutants, the state must submit periodic comprehensive plans that will lead to attainment status in the near future. If standards are not met in due time, states risk losing federal monies that help to fund state-level public goods and services (see, e.g., Becker and Henderson, 2000; Greenstone, 2002).[8]

Environmental regulations in nonattainment counties are intended to be stringent. Polluting plants entering or expanding in a nonattainment county are subject to a standard of "Lowest Achievable Emission Rate (LAER)" without consideration of cost for all investments. The resulting rules frequently involve compliance with "command and control" style regulations that requires the installation and operation of specified pollution abatement equipment. Further, emissions from new investment must be offset by emissions reductions from an existing source within the same county, and plant expansion or modification leads to the entire plant being

---

[6] In 1987 the EPA changed its focus from the regulation of all particulates (i.e., TSPs) to the smaller PM_{10}s, which have an aerodynamic diameter equal to or less than 10 micrometers. In 1997 the PM_{10} regulation was replaced with a $PM_{2.5}$ one.

[7] There are separate standards for ozone (O_3) and nitrogen dioxide (NO_2). In principle, a county could meet one of these standards, but not the other. However, O_3 is the result of a complicated chemical process that involves NO_2, and the vast majority of counties that were nonattainment for NO_2 were also nonattainment for O_3. As a result, we designated a county nonattainment for O_3 if the EPA labeled it nonattainment for either O_3 or NO_2. All future references to O_3 refer to this combined measure.

[8] While the EPA denoted each county beginning in 1978 as either in or out of attainment for each criteria air pollutant, Greenstone (2002) compiled the data back to 1972 using air quality data collected via filing a Freedom of Information Act petition.

subject to more stringent regulations (see List et al., 2004).[9]

Polluting plants locating in attainment areas, on the other hand, face a more lax regulatory standard. These plants are subject to the standard of "Prevention of Significant Deterioration (PSD)." This entails permitting and the installation of the "Best Available Control Technology (BACT)" for new plants that have the potential to emit over 100 tons of a criteria pollutant in a year. The BACT is negotiated on a case-by-case basis and the economic burden on the plant is considered in arriving at a final solution. Given that the installation of BACT in attainment areas is likely to be much less costly than the installation of LAER in nonattainment areas, new polluting plants and expansions of existing ones could face significantly lower pollution control capital construction costs in attainment areas versus nonattainment counties.

Given that SIPs require states to develop plant-specific regulations for every major source of air pollution, existing plants in nonattainment areas also face greater regulatory scrutiny than plants in attainment areas. These plant-specific regulations typically have come in the form of emissions limits. Beyond the necessary abatement investments, inspections and regulatory oversight are more persistent in nonattainment areas. Further, the size of the existing polluter importantly determines the level of regulation (see Becker and Henderson, 2000).

Both the states and the EPA are given substantial enforcement powers to ensure that the CAAA's intent is met. For instance, the EPA must approve all state regulation programs in order to limit the variance in regulatory intensity across states. On the compliance side, states run their own inspection programs and frequently fine non-compliers. The 1977 legislation also made the plant-specific regulations both federal and state law, which gave the EPA legal standing to impose penalties on states that do not aggressively enforce the regulations *and* on plants that do not adhere to the regulations. A number of studies document the effectiveness of these regulatory actions at the plant level (Nadeau 1997, Cohen 1998). Perhaps the most direct evidence that the regulations are enforced successfully is that air pollution concentrations declined more in nonattainment counties than in attainment ones during the 1970s and 1980s (Henderson 1996, Chay and Greenstone 2003 and 2005, Greenstone 2002).

B. Variation in Regulation as the Basis of a Research Design

[9] The reduction in pollution due to the offset has to be larger than the expected increase in pollution associated with the new investment. The offsets could be purchased from a different facility or generated by tighter controls on existing operations at the same site (Peirce, Vesilind, and Weiner 1998).

The structure of the CAAAs provides three sources of variation in which plants were affected by the nonattainment designations. This subsection summarizes these three dimensions of variation and highlights their importance from an evaluation perspective. It also briefly discusses some of the sources of this variation and why they may reinforce the credibility of the subsequent analysis.

The first dimension of variation is that at any point in time the pollutant-specific nonattainment designations are reserved for counties whose pollution concentrations exceed the federal standards. This cross-sectional variation allows for the separate identification of industry-specific shocks and regulatory effects. This may be especially important in the period we study, because there were dramatic shocks—oil crises, recessions, and increases in foreign competition—that affected industries differentially.

The second dimension of variation is that a county's attainment/nonattainment designations vary over time as its air quality changes. Consequently, individual plants might be subject to regulations in one period but not in a different one. This longitudinal variation allows for the inclusion of plant fixed effects in equations analyzing plant-level productivity. Consequently, the paper presents estimated regulation effects that are derived from within-plant comparisons or comparisons of the productivity of polluting plants in counties that experience a change in a nonattainment designation.

The third dimension of variation is that within nonattainment counties only plants that are major emitters of the relevant pollutant(s) are subject to the regulations. We follow Greenstone (2002) and classify sectors based on their emission levels. Using information from EPA's *Sector Notebook Project*, we label industrial sectors as "pollution-intensive" if they emit at least 10 percent of the total industrial sector's emissions of the pollutant under consideration.[10] All other industries are considered non-emitters. Table 1 lists the polluting industries and the pollutants they emit. The nonpolluting plants in nonattainment counties allow for the estimation of statistical models that control for shocks to productivity common to polluters and non-polluters in nonattainment counties.

[10] This is higher than the seven percent threshold used in Greenstone (2002). Below, we present evidence suggesting that the marginal plants experiencing productivity impacts from nonattainment designations are in industry groups accounting for between seven and ten percent of industrial sector emissions. Plants in cleaner industry groups see little measurable productivity effects; those in dirtier industry groups see impacts that are roughly the same size as those just above our 10 percent cutoff.

Some of the sources of variation in nonattainment status reinforce the credibility of an evaluation based on the CAAAs. County-level nonattainment designations are federally mandated and are therefore less likely to be related to differences in tastes, geographic attributes, or underlying economic conditions across counties. Moreover, nonattainment designations depend on whether local pollution levels exceed the federal standards. While pollution levels are not randomly assigned, scientific evidence suggests that during the years under study many counties were designated nonattainment due to pollution that was related to weather patterns—a factor which is likely to be unrelated to local manufacturing sector activity.[11]

IV. Data

Our primary data sources are a county-level panel on CAAA attainment status and production microdata for manufacturing plants from the U.S. Census Bureau. We describe each in turn.

A. CAAA Attainment Status Data

We collected information on annual CAAA nonattainment status of 3,141 U.S. counties from 1972 to 1993. We observe in each year whether or not the county is in attainment with CAAA standards for each of four pollutants: ozone (O_3), total suspended particulates (TSPs), sulfur dioxide (SO_2), and carbon monoxide (CO). We construct county-level nonattainment measures from this data.

Table 2 summarizes the variation in attainment status across counties and over time. It reveals the distribution of counties' lagged and current nonattainment status by pollutant. It also includes a pooled category of "any pollutant," which holds if the county is nonattaining in any one or more of the four pollutants we track. The numbers in the table are the respective counts by category of lagged and contemporaneous attainment status counts for the 65,961 county-years in our annual CAAA attainment data over the 1972-1993 period. For example, 52,390 counties were in attainment for all pollutants in both the current and previous years; 1,461 were in

[11] Cleveland et. al. (1976) and Cleveland and Graedel (1979) document that wind patterns often cause air pollution to travel hundreds of miles, and that the concentration of O_3 in the air *entering* the New York region in the 1970s often *exceeded* the federal standards. Figure 1b in Greenstone (2002) graphically depicts the counties that were designated nonattainment for O_3 and reveals that virtually the entire Northeast, even counties without substantial local production of O_3, is in O_3 nonattainment for at least one period. It is evident that this region's nonattainment designations partially reflect its location downwind from heavy O_3 emitters in the Ohio Valley.

nonattainment in the current year but were in attainment the previous year, and so on.

Counties fall into nonattainment and come back into attainment at roughly the same rates, indicating little change in average nonattainment rates over time. However, this hides clear patterns in the time series for all pollutants (not shown here for space reasons) that show a rapid rise in the number of counties in nonattainment in the 1970s and a slow but steady decline thereafter. Looking at specific pollutants, changes in attainment status are most common for ozone, with TSPs close behind. Sulfur dioxide and carbon monoxide attainment changes are notably less frequent than for the other two pollutants.

These within-county changes in attainment status are a basic source of identification for our estimates of the effects of nonattainment status. While the county-years that see attainment status changes are a modest share of the entire sample, they still correspond to hundreds and sometimes thousands of county-level changes. Moreover, because most of the attaining counties are in rural areas with little manufacturing activity, the counties that are always in attainment of the federal standards account for a much smaller fraction of economic activity than their shares in Table 2.[12] Hence the share of manufacturing plants and output that exists in counties that experience attainment status changes is much larger than the percentages in Table 2 suggest.

This extent of the CAAA's reach is seen in Figure 1, which reports the share of polluting industries' total output that is produced by plants in nonattainment counties over the course of the sample. Panel A shows the series for industry production in nonattainment counties in any (one or more) of the pollutants for which each industry is a heavy emitter; panel B shows the corresponding pollutant-specific series. Over 30 percent of emitting industries' output was produced in nonattainment counties at the beginning of the sample in 1972. This quickly rose to more than 65 percent by 1978 (due mostly to additional counties being declared nonattainment rather than firms moving activity to nonattainment counties), after which it slowly fell to around the 55 percent level by the end of the sample. Panel B shows that ozone-emitting industries have the greatest share of their output produced in nonattainment counties, peaking at over 70 percent in 1978-79. TSP and CO emitting industries have nonattainment county production shares on the order of 20-30 percent (both see steady declines in this share over the sample), while SO_2 emitters only produce around 10 percent of their output in nonattaining counties.

[12] See the maps in Greenstone (2002) for information on the location of nonattainment counties by pollutant.

B. Census Manufacturing Micro Data

The other primary data source is the plant-level micro data on manufacturers from the U.S. Census Bureau. This is comprised of the Annual Survey of Manufactures (ASM), the Longitudinal Business Database (LBD), and the Census of Manufactures (CM). These data files contain detailed production data for manufacturing plants with a total of nearly 1.2 million plant observations over the course of the sample. A plant—or "establishment" in Census Bureau terminology—is a physical location where economic activity takes place. In the manufacturing sector, this can be thought of as a factory. A firm can own one or many plants.

The ASM contains production data that include plant revenues, several labor input measures, book values of equipment and structures capital stocks, investment in equipment and structures, and expenditures on inputs. We use these data to calculate plants' total factor productivity levels as described below. The ASM data also include a unique permanent plant identifier that allows us to link plants across years and estimate models with plant fixed effects. Critically, we also observe the state and county in which a plant is located, allowing us to match plants to our county-level attainment status file.

The ASM, as its name indicates, is taken annually. The ASM sample is comprised of rotating five-year survey panels that begin in years ending with "4" or "9" and end in years ending with "8" or "3" respectively. The panels are selected to be representative of the manufacturing sector. Large plants (those with over 250 employees) are sampled with certainty; sampling probabilities increase with size for plants below this threshold. A typical ASM year contains about 60,000 plants. The ASM microdata contain plants' sample weights (the inverse of their sampling probabilities), allowing us to obtain values that are representative of the entire manufacturing sector. ASM data are available from 1972 on, but information on capital stocks is only available until 1993. Because we need capital inputs to construct total factor productivity measures, our ASM sample spans 1972-1993.

The most important purpose of the ASM production data is to construct measures of plants' total factor productivity (TFP) levels. Our empirical specifications use TFP measures based on index number methods, where a plant's TFP is its logged output minus a weighted sum of its logged labor, capital, materials, and energy inputs. That is,

$$TFP_{it} = y_{it} - \alpha_{lt} l_{it} - \alpha_{kt} k_{it} - \alpha_{mt} m_{it} - \alpha_{et} e_{it},$$

where the weights α_j are the input elasticities of input $j \in \{l, k, m, e\}$.[13] (Thus our TFP measure is the natural logarithm of a plant's ratio of output to inputs.) Output is the plant's inventory-adjusted total value of shipments deflated to 1987 dollars. Inputs are plant-specific, but we use industry-level input cost shares to measure the input elastiticies. These cost shares are computed using reported industry-level labor, materials, and energy expenditures from the NBER Productivity Database (which is itself constructed from the ASM). Capital expenditures are constructed as the reported plant's capital stocks multiplied by their respective BLS capital rental rates in the corresponding two-digit industry. Details on the construction of the TFP index are in the Appendix.

A major advantage of the ASM for our purposes is its frequency. That it is a survey rather than a census is a weakness. The fact that it is designed as a representative sample, however, and that we have the sampling weights, assuages our concerns on this point. Another potential weakness is that the ASM's rotating panel structure could pose a problem for identifying which plants exit, as in panels' final years it will not be clear whether a plant that disappears from the sample does so because it is rotated out of the sample or because it ceased operations. Identifying exiters will be important below when we try to correct for the likely possibility that those plants which take the largest productivity hits from regulatory action are also more likely to exit. Fortunately, we can supplement our ASM files with LBD data, which contains annual data on plants' activity status, and therefore identifies the actual year of exit (if exit has in fact occurred).

The CM, which is a census of the roughly 350,000 manufacturing plants operating in the U.S. in a typical year, is only taken quinquennially, in years ending in "2" and "7". We will use CM data for one particular industry (ready-mixed concrete) in a specification below where we investigate the effects of nonattainment designation on plants' prices. Plant-level price data is typically not available in producer microdata of this sort, but the CM collects separate information on plants' revenues and physical quantities for a limited number of industries. This allows us to compute plant-level average unit prices as well as measure TFP in physical terms, as

[13] For a brief discussion of TFP measurement and citations for further reference, see Syverson (2011). We prefer index numbers in this application for the ease with which they allow for flexible technologies and their avoidance of the problem of inputs being endogenous functions of TFP. Like any TFP measurement methodology, however, they require assumptions. Besides cost minimization, index numbers also assume the plant faces no adjustment costs in inputs. Fortunately, empirical results in the literature using micro-level TFP data have been typically quite robust to the specific method used to obtain the TFP measure.

opposed to the standard deflated-revenue-based real output measures which confound within-industry price variation with output variation. (We discuss these output measurement issues, and their implications for our estimates of regulation's productivity effects, in greater detail below.)

V. Empirical Specification

We seek to estimate the effect of CAAA nonattainment status on polluting plants' productive efficiencies as embodied in their TFP levels, as described in the conceptual framework. Because a plant's TFP reflects how much output it produces from a given amount of inputs, our estimates below measure the change in a plant's output due to the nonattainment regulation given a fixed set of inputs. We estimate the following specification:

(1) $$TFP_{it} = \sum_p \{\beta_p I[nocaaa_{cpt}] + \delta_p I[pollind_{ip}] + \gamma_p I[nocaaa_{cpt}] I[pollind_{ip}]\} + X_{it}\Phi + \eta_i + \varepsilon_{it},$$

where *i* indexes a plant, *t* references a year, *p* indicates a pollutant, and *c* indexes a county.

TFP_{it} is the natural logarithm of plant i's total factor productivity in year t. Importantly, this equation allows for permanent differences in plant productivity, which are captured with the vector η_i that includes a separate fixed effect for each of the roughly 189,000 plants in the sample. The vector X_{it} includes Census-geographic-division-by-year or Census-geographic-division-by-ASM-panel fixed effects. These fixed effects control for differential changes in average productivity levels across geography that are common to polluters and non-polluters.

Two indicator functions are the core of the regression's explanatory variables. $I[pollind_{ip}]$ equals one if plant *i* is in an industry that is classified as a heavy emitter of pollutant *p* as detailed in Table 1. In practice, we model industry effects more flexibly by also including 2-digit-SIC-industry-by-ASM-panel (i.e., 5 years) fixed effects and 4-digit-SIC-by-year fixed effects in other specifications.[14] These fixed effects control for the considerable variation in TFP

[14] The extent of disaggregation in our industry-by-period and geography-by-period controls is subject to computational constraints. There are 20 two-digit SIC manufacturing industries, five ASM periods, and nine Census geographic divisions. Thus there are 145 (20·5 + 9·5) time-varying fixed effects in addition to the roughly 200,000 plant fixed effects in our benchmark specification. Specifications using fixed effects based on more disaggregated industry, time, or geographic categories proved computationally unworkable given our sample size of almost 1.2 million plant-year observations. However, we explore the robustness of our benchmark results to two-digit-industry-by-year and division-by-year fixed effects (for a total of 20·22 + 9·22 = 638 fixed effects) below.

levels and shocks across industries.[15]

The second core indicator is $I[nocaaa_{cpt}]$, which equals one if the county *c* in which plant *i* is located is designated nonattainment for pollutant *p* in year *t*. These indicators control for productivity differences common to polluters and non-polluters in nonattainment counties. For example, counties in dense urban areas are more likely to be in nonattainment, while at the same time the plants in these counties might enjoy productivity benefits from agglomeration spillovers (Greenstone, Hornbeck and Moretti 2010).

The parameters of interest are the γ_p's. They capture the variation in TFP specific to plants that operate in a county designated nonattainment for a particular pollutant *and* are in industries that are heavy emitters of that same pollutant. In specifications that include the plant fixed effects, η_i, the γ_p's are identified from polluting plants in counties that experience a change in the nonattainment designation for the pollutants they emit.

The $\hat{\gamma}_p$ estimates are obtained using two approaches. In the first, we pool across the four pollutants (O_3, TSPs, SO_2, and CO) when defining $I[nocaaa_{cpt}]$ and $I[pollind_{ip}]$. In this case, $I[nocaaa_{cpt}]$ equals one if the county is in CAAA nonattainment for any one or more of the four pollutants, and $I[pollind_{ip}]$ equals one if the plant is in a heavily-emitting industry in one or more of the pollutants in which the county is in nonattainment. This specification captures the effect of nonattainment status on TFP and averages it across plants that face the nonattainment designation for just one or multiple pollutants. The second approach controls separately and simultaneously for each of the four pollutants, allowing for the estimation of pollutant-specific effects while holding the impacts of others constant. This specification allows for heterogeneity in the impact of the nonattainment designation across pollutants and is also informative about the impacts on plants in industries that are heavy emitters of multiple pollutants.

A few other estimation details are noteworthy. All versions of equation (1) are weighted by the plant's ASM representative real output. This is the product of the plant's reported real output and the plant's ASM weight (its inverse sampling probability in the ASM panel; see Section IV). Consequently, the regressions measure average TFP effects on a dollar-weighted basis, which means that the results can be interpreted as aggregate average effects. Additionally,

[15] Even though our benchmark specification includes plant fixed effects, $I[pollind_{ip}]$ is identified because some plants switch industry classifications during our sample. These industry switchers are substantial in number; if we exclude them our sample drops by one-third. We test our results for robustness to excluding industry changers below.

the tables report standard errors based on clustering at the county-by-year level to account for the likely dependence in TFP innovations across plants in the same county and year. We also estimated standard errors where the clustering is done at the county level to allow for arbitrary time-series correlation in TFP shocks within a county and those are discussed briefly below.

VI. Results

A. Benchmark Specification

Table 3 reports the values of $\hat{\gamma}_p$ and the associated standard errors obtained from estimating several versions of equation (1). While all specifications include the main effects—$I[nocaaa_{cpt}]$ and $I[pollind_{ip}]$—for parsimony we only report the coefficients of interest, $\hat{\gamma}_p$, from their interaction. The set of controls increase in detail, and thus control for more data variation, moving from left to right in the table.

Empirical results in columns 1 and 2 include a full set of four-digit-SIC-industry-by-year fixed effects. These capture all common productivity movements within a four-digit industry from year to year. The coefficient in column 1, which shows the estimate for the composite "any pollutant" variable, indicates that TFP is 2.4 percent lower for a polluting plant in a county that is nonattainment for that pollutant; that is after adjustment for annual shocks to TFP common to plants in the same 4-digit industry, polluting plants in the same industry in attainment counties have a 2.4 percent higher TFP than plants in nonattainment counties. Alternatively, it can be described as polluting plants in nonattainment counties producing 2.4 percent less output from a fixed set of inputs. The effect is precisely estimated, with a standard error of 0.3 percent.

Empirical results in column 2 indicate that this average effect pooled across industries masks considerable differences in pollution-specific TFP effects of nonattainment. Ozone emitting plants in ozone nonattainment counties experience productivity losses of 1.8 percent, relative to comparable plants in attainment counties. TSP emitting plants in TSP nonattainment counties appear to, if anything, see marginally significant TFP growth of about 1.0 percent. Sulfur dioxide emitters see no substantial change in TFP when their county is in nonattainment for SO_2, and CO emitters suffer significant productivity losses on the order of 2.1 percent. In interpreting these results, it is important to bear in mind that these pollutant-specific TFP effects are all estimated holding the others constant. Further, plants in industries that are emitters of multiple pollutants will experience average total productivity effects that are larger than any of these individual pollution-specific components if their county is nonattainment for multiple pollutants. Indeed, the fact that some plants face nonattainment determined regulations for multiple pollutants helps to explain why the estimate in column 1 is not a weighted average of the column 2 estimates.[16]

The specifications in columns 3 and 4 add controls for geographic variation in average productivity levels by including census-division-by-year fixed effects. There is little change in the estimates from columns 1 and 2. The composite nonattainment coefficient indicates a -2.3 percent change in TFP, with a standard error of just 0.3 percent. The pollution-specific estimates track those in column 2 closely.

Columns 5 and 6 report the results from specifications that include plant fixed effects, so the nonattainment productivity effects are identified from within-plant TFP variations. In other words, the identifying variation comes from polluting plants located in counties that move from nonattainment to attainment or attainment to nonattainment. The specifications also include two-digit-SIC-industry-by-ASM-panel and census-division-by-ASM-panel fixed effects to control for any differences in broad productivity changes across manufacturing industries or regions.

The estimated composite effect in column 5 is a statistically significant 2.6 percent TFP drop (s.e. = 0.6 percent), which is qualitatively identical to the estimates from the previous specifications. Among the pollutant-specific estimates in column 6, again O_3 is the largest, at -2.2 percent. Interestingly, adding plant fixed effects makes the TSP estimate negative and significant but causes the CO estimate to become positive and significant. Evidently, all else equal, plants with persistently high (low) productivity levels in TSP- (CO-) emitting industries are more likely to be in nonattainment counties. Controlling for plant fixed effects removes this correlation and produces the observed change in the coefficients. SO_2 nonattainment corresponds to a TFP decline of 1.6 percent among SO_2-emitting plants, though the p-value of this estimate is 0.11.[17]

Columns 7 and 8 are richer models that include 4-digit-SIC-by-year and census-division-

[16] In addition, the column (1) specification will capture some interactions due to the regulation of multiple pollutants that are not accounted for in the column (2) specification.

[17] We also estimated the column 5 and 6 specifications while clustering standard errors by county rather than county-year to account for arbitrary time series correlation in TFP shocks in the same county. The estimated TFP effect in the pooled pollutant sample remains significant at the 5 percent level, as does the effect of ozone nonattainment in the pollution-specific specification. The coefficients on the other pollutants that were marginally significant in the benchmark results become insignificant, however.

by-year fixed effects. In general, the estimated coefficients are larger in magnitude than those from the coarser fixed effects structure in columns 5 and 6. The composite pollutant interaction coefficient in column 7 is now -4.4 percent and is estimated precisely. The O_3 nonattainment designation is associated with a 5.7 percent decline in TFP for ozone emitters, and SO_2 nonattainment is tied to a 2.1 percent productivity drop. TSP effects are negative but insignificant, and again CO nonattainment is correlated with productivity increases.

A few conclusions emerge from these specifications. First, a nonattainment designation for any pollutant reduces the TFP of plants that are heavy emitters of that pollutant by around 2 to 4 percent, with the larger estimates coming from the more robust specifications. Second, the largest declines in TFP are associated with O_3 nonattainment, which incidentally is one of the most commonly emitted pollutants among our industries. Third, the estimated effects for the other pollutants are more sensitive to specification. In the more reliable specifications with plant fixed effects, the TSPs and SO_2 nonattainment effects are typically in the range of 1 to 2 percent declines, while CO nonattainment, on the other hand, is associated with productivity increases. In this manner, there is some evidence consonant with the Porter (1991) effect.

While the results are robust to the inclusion of the full array of controls in columns 7 and 8 (and in fact imply larger magnitude effects than most of the other specifications), this specification was unfortunately unworkable as a practical matter. Many regression runs with this specification failed due to insufficient hardware resources (despite the formidable capabilities of the Census Bureau servers) and successful runs were considered unfair uses of Census computer resources.[18] Consequently, the benchmark specification throughout the remainder of the paper will be the model estimated in columns 5 and 6 of Table 3, which balances computational feasibility with the ability to control for a broad set of potential confounders. Additionally, this specification's estimates of the productivity effects of CAAA nonattainment are toward the middle of those presented in Table 3, offering a somewhat conservative measure of total TFP impacts.

Before proceeding, it is useful to place the magnitude of the estimates from the benchmark specification in context. The average yearly output of plants in any polluting

[18] Furthermore, we were completely unsuccessful in running specifications with even finer fixed effect structures (e.g., state-by-year fixed effects).

industry in nonattainment counties from 1972-1993 was roughly $412.5 billion in 2010 dollars.[19] Taking the column (5) estimate at face value, the estimated average annual cost of CAAA nonattainment in lost output was therefore about $11.0 billion in 2010 dollars.[20] Under the strong assumption that the affected manufacturers are price takers, the $11 billion is the annual reduction in manufacturing sector profits. As a base of comparison, from 1972-1993, annual manufacturing sector pre-tax profits averaged $236 billion (in 2010$).[21]

B. Dynamic Effects

The specifications above assume that the nonattainment designation's productivity impact is contemporaneous. Yet it seems possible that there could be longer-lasting cumulative effects of nonattainment on plant productivity (see, e.g., List et al., 2003). For example, it is difficult for the EPA and local environmental regulators to compel compliance of all plants in high-pollution industries within the first year of nonattainment. It may take a year or more for all plants to take the required costly abatement actions. The permitting requirements for plant expansions can involve prolonged negotiations between engineers and consultants detailing required start-up capital. The regulations affect plants' production choices (e.g., their capital stocks), and the impacts of these choices on TFP may affect productivity even many years after a county has moved from nonattainment to attainment status. Finally, as an empirical matter, nonattainment is autocorrelated; dirty counties often keep that designation for several years.

To explore the possibility of cumulative productivity effects, we estimate a version of our benchmark specification where we also include indicators of lagged nonattainment status. The total estimated impact of being subject to the nonattainment related regulations for one year is

[19] We obtain this and other values in the paper expressed in 2010 dollars by inflating measures in the data, which contain real 1987 values, using the Bureau of Labor Statistics' annual producer price index for "Total Manufacturing Industries." The index equals 100.9 for 1987 and 175.4 for 2010, implying nominal price growth in the sector of 73.8 percent over the period.

[20] This is computed as the difference between the counterfactual output of $423.5 billion (= $412.5B/(1 – 0.026)) and the observed output of $412.5B. Similar calculations for the pollutant-specific results in column 6 of Table 3 indicate a smaller annual lost output cost of just over $6.5 billion (again in 2010 dollars), though recall that this value will not adequately capture the total TFP change of plants emitting multiple pollutants in counties that are nonattainment for multiple pollutants.

[21] Manufacturing profits are taken from the "Corporate Profits Before Tax by Industry" tables of the National Income and Product Accounts. These are deflated to 1993 values using the GDP price index, and adjusted to 2010 values using the manufacturing sector PPI used elsewhere in this paper.

then the summed marginal effects across the contemporaneous and lagged impacts.[22] We estimate specifications including one and two years of lagged attainment status. To keep a consistent sample across the specifications, we estimate each using plants that are in at least their third year of operation. This sample contains roughly 800,000 plant-year observations.

Columns 1 and 2 in Table 4 are repeated from columns 5 and 6 of Table 3, and are included as a basis for comparison. Columns 3 and 4 in Table 4 report the results from the same specifications fit on the smaller, two-lag sample. Here, the estimated regulation effect from the any pollutant specification is -1.7 percent and would be judged to be statistically significant. This effect is modestly smaller than the effect estimated from the full sample, and it is apparent that this is related to the much smaller and insignificant effect of ozone nonattainment status on ozone emitters. The estimated impact of SO_2 nonattainment is somewhat larger than in the whole sample, while the TSP and CO nonattainment effects are qualitatively identical.

Columns 5 and 7 in Table 4 report the estimated cumulative marginal TFP effects for the pooled pollutant specification, which are the sums of the contemporaneous and lagged coefficients on the interactions of emitting-plant and nonattainment indicators. The estimated effect from the specification that includes one lag yields an estimated TFP loss of 2.0 percent, while the specification accounting for two years of nonattainment lags indicates a total productivity drop of 3.1 percent. The latter estimate suggests that the total TFP impact for a polluting plant in a county that is nonattainment for that pollutant is 80 percent larger than the contemporaneous effect in a consistent sample.

The pollutant-by-pollutant effects shown in columns 6 and 8 indicate that the cumulative impacts in the pooled pollutant specification reflect increasingly negative cumulative effects of ozone, TSP, and SO_2 nonattainment. Interestingly, the positive estimated TFP effect of CO nonattainment falls by more than half, and is insignificant in the column 8 specification.

Overall, we conclude that a year's nonattainment designation affects a plant's operations and productivity for at least three years. This finding is consistent with the conceptual framework's assumption that the nonattainment designation leads firms to install pollution abatement equipment that does not increase output.

[22] These lags are included both as main effects and interacted with the polluting-industry indicators. As with the contemporaneous effects, the interaction coefficients are our focus.

C. Robustness Checks

This subsection probes the robustness of the results to several variations in the details surrounding our measurement practices and our empirical specification.

Excluding Industry Switchers. Our sampled universe includes a substantial number of plants that switch industries during the time period. Some of these plants may switch out of (or into) high-emissions industries. If these shifts are coincident with nonattainment designations in the plants' counties, this could impact the estimated effects of nonattainment. To test if the benchmark results are sensitive to this industry switching, we fit the benchmark specification on the sample of plants that remain in the same four-digit SIC industry during their entire time in the sample. This shrinks the sample to approximately 807,000 observations and 135,000 unique plants.[23]

Empirical results from this smaller sample are in columns 1 and 2 of Table 5. They qualitatively, and to some extent quantitatively, match those observed from the benchmark model. The estimated TFP effect of nonattainment in the pooled pollutant specification is a statistically significant -1.6 percent, though smaller than the benchmark estimate. Since ozone-emitting industry plants account for a large share of all emitters, the smaller ozone effect helps to explain this smaller 'any pollutant' effect. In contrast, the TSP and SO_2 effects are significant and larger in magnitude at -2.7 and -2.2 percent, respectively. Again, plants in industries that emit large amounts of CO experience a productivity gain.

Heavy-Emissions Industry Definitions. As we discussed above, a key source of identification for the CAAA's productivity effects is in the comparison of dirty (high-emissions) and clean (low-emissions) plants in nonattainment counties. This is based on the notion that environmental regulators will focus abatement efforts on the heaviest polluters first. However, our threshold for dirty—that the plant's industry group accounts for at least 10 percent of the

[23] Excluding any four-digit SIC industry changers is quite conservative in that some of these switchers may change from one industry that is a heavy emitter of a particular pollutant to another industry that is similarly a heavy polluter of that pollutant. For example, suppose due to a change in the products it made that a plant moved from the Brick and Structural Clay Tile industry (SIC 3251) to the Structural Clay Products, Not Elsewhere Classified industry (SIC 3259). Because both of these four-digit industries are considered part of the heavy O_3/TSP/SO_2 emitting Stone, Clay, Glass, and Concrete industry group (SIC 32), the indicator variables definitions for this plant would not change. As such, neither would this or other similar cases lead to changes in the interaction coefficients. However, if as discussed above there are systematic differences in average TFP levels across four-digit industries within particular heavy-emitter industry groups, then plant fixed effects would not control for such changes for industry-shifting plants. As such, restricting attention to plants that stay in the same narrow industry will avoid this problem.

industrial sector's emissions of a pollutant—is arbitrary. In this subsection we explore the sensitivity of the results to this 10 percent emissions share cutoff.

We first estimate our benchmark specification where any industry group accounting for at least *seven* percent of industrial emissions of a pollutant is considered a heavy emitter. This is the same cutoff used in Greenstone (2002). We then obtain estimates using a more stringent cutoff, requiring the industry group to account for at least 12 percent of industrial emissions of that pollutant.[24]

Columns 3 and 4 of Table 5 report the results using the seven percent cutoff. The pooled-pollutant estimated effect is -2.1 percent, slightly smaller but similar to the estimate using only industries that meet the 10 percent cutoff. The ozone-specific estimate exhibits the same comparative pattern, being smaller than its 10 percent analog but still statistically significant at -1.7 percent. The estimated TSP and SO_2 effects are smaller in magnitude, however, and would be judged statistically insignificant at conventional levels. The CO estimate is exactly the same because there are no industry groups that account for between 7 and 10 percent of industrial sector emissions of carbon monoxide. Overall, these small estimated TFP impacts suggest that the EPA targeted its regulatory efforts at the heaviest polluters.[25]

The results using the 12 percent cutoff group are in columns 5 and 6 of Table 5. Here, the results are quite similar to those from the benchmark. The composite effect in the any-pollutant specification is -2.3 percent. The pollutant-specific estimates for ozone, TSPs, and SO_2 are almost unchanged. The only substantial change from the 10-percent-cutoff estimates, in fact, is for CO, where the formerly positive and significant effect on the order of 1.5 to 2 percent is now a small and insignificant 0.2 percent. This finding reveals that the positive TFP effects associated with nonattainment in this pollutant are concentrated in petroleum refining industry (accounting for 11.8 percent of CO emissions), which is the only industry dropped from the set of dirty plants with the higher cutoff.[26] Overall, the similarity of these and the benchmark

---

[24] Which formerly clean industry groups become dirty (in the seven percent cutoff specification) or dirty industry groups become clean (in the 12 percent specification) can be seen in Table A2 of Greenstone (2002), which lists all industry groups tracked by the EPA and their pollution-specific emissions shares.

[25] The use of a 4.5 percent cutoff produces even smaller TFP estimates, suggesting that the marginal plants to experience significant productivity impacts from nonattainment designations are those in industry groups accounting for around seven percent of industrial emissions.

[26] It is noteworthy that petroleum refining is also a heavy emitter of ozone and SO_2, so it is possible that the positive CO effect reflects interactions across the regulation of multiple pollutants.

estimates suggests that the marginal plants to receive regulatory oversight are those in industries accounting for less than 10 percent of industrial sector emissions of particular pollutants.

Industry-Specific Productivity Effects. The benchmark specification restricts the effect of nonattainment status to be constant across all industries that are heavy emitters of a particular pollutant. Here, we test this restriction by estimating industry-specific analogs to our pooled pollutant and pollutant-specific models. In the pooled pollutant analog, we interact indicators for the industry groups detailed in Table 1 with indicators for counties that are in nonattainment with any pollutant that the industry emits, while in the pollutant-specific analog, they are interacted with nonattainment indicators for specific pollutants.

The results from the pooled pollutant specification are in Table 6A. There is notable dispersion in the estimated TFP effects across industry groups with the organic chemicals industry experiencing the largest productivity decline of roughly 17 percent. When multiplied by this industry's average annual output in nonattainment counties (the table reports these values in 2010$ for each industry), their annual loss was roughly $9.2 billion. Interestingly, ozone is the only pollutant for which they are regulated. Further, the nonattainment designation is associated with TFP declines of 6.3 percent and 2.4 percent for nonferrous metals and rubber plants, respectively. Between them, this corresponds to an average lost output of about $2.3 billion per year. On the other hand, pulp and paper plants see TFP gains of 2.5 percent associated with nonattainment designations, though the relatively small size of the industry ($26.6 billion average revenue per year in 2010 dollars) imply that only about $0.6 billion of extra output was gained. The estimates for other industries are modest in magnitude and statistically insignificant. Given this dispersion, it is not surprising that an F-test of the null hypothesis that the industry-specific nonattainment effects are equal is easily rejected.

Table 6B contains the pollutant-specific results, with separate TFP effects estimated for all 15 industry-pollutant interactions. The enormous estimated productivity loss due to nonattainment for organic chemicals producers remains, here seen in the ozone nonattainment interactions (as it must because ozone is the only pollutant for which the industry is a heavy polluter). The point estimate is therefore equivalent to that in panel A. Pulp and paper's positive estimate in panel A is driven primarily by a positive effect of ozone nonattainment on TFP.

Another industry-specific result of note is that the only significantly positive effect of CO nonattainment is seen in refining. First, it counteracts significant productivity losses among industry plants associated with SO_2 nonattainment. In combination, they give the more-or-less zero result for the industry in panel A. Second, the CO coefficient is consistent with the earlier finding that the positive effect of CO went away when we used a higher threshold for dirty CO plants that excluded refiners.

Table 6B also reports test statistics and associated p-values for F-tests of the hypothesis that the pollutant-specific productivity effects are equal across industry groups. The clear rejection in the pooled pollutant specification mostly reflects heterogeneous ozone effects, as we cannot reject equality at the five percent level for TSPs, SO_2, and CO.

VII. Do the Estimated Regulation Impacts Understate the True Effect on TFP?

The estimates in the previous section likely to *understate* the true magnitude of productivity losses resulting from nonattainment designations for three main reasons. First, the estimated TFP impacts are based on a comparison of TFP among high and low emissions plants. This allows us to use the "clean" plants to control for local shocks to economic activity that may confound the nonattainment designations. However, some of the plants that we have categorized as low emitters were likely targeted by the CAAA's regulations. To the extent that this is the case, it will cause the estimated effects of regulation to be understated.

Second, like almost all users of plant- or firm-level production data, we must use revenue to measure plant output in our TFP measures. This is a result of limited data: plant- or firm-level price information is not available in the ASM, or almost any other similar data sets. The use of revenue-based TFP measures means any price differences across plants in an industry will therefore be measured as output and productivity variation. The problem in our context is that abatement actions that reduce TFP may raise marginal costs. Standard theory predicts that when these plants have market power they will increase the price they charge for their product. The consequence is that revenue-based TFP measures will conflate regulation's negative impact on technical efficiency with the positive price change, resulting in an understatement of the true technical efficiency (and output quantity) loss.

The third source of bias is due to endogenous selection of which plants survive a CAAA nonattainment designation. All else equal, plants experiencing the largest negative productivity shocks from these environmental regulations are the most likely to cease operating and shut down. Our sample of survivors experienced productivity drops that were likely to be smaller on average than those across all plants. And, the resulting estimates will not reflect the most negative productivity innovations experienced due to the application of the clean air regulations.

While there is little we can do to quantify the impact of our inability to measure any TFP effects that are common across both "clean" and "dirty" plants, the data do allow us to obtain some sense of the degree of understatement caused by the price measurement and survivor bias problems. This section explores the extent of these possibilities.

A. Regulation's Impact on Prices

Plant-level price data is unavailable for most industries. However, the Census of Manufactures does collect plant-level output data in both revenue and physical quantity terms for a few industries. We can use these quantity data to directly measure plants' physical TFP levels (the number of physical units of output they produce per unit input), which bypasses the price-as-output measurement problem altogether. Moreover, these same data allow us to compute plants' average unit prices and see how they vary with nonattainment status.

While physical quantity data are available for several industries in the CM, we focus here on ready-mixed concrete (SIC 3273). The industry has several features that lend itself to accurate measurement of quantity-based TFP. Ready-mixed concrete is a physically homogeneous product, so the output quantities (measured in thousands of cubic yards) are comparable across plants. One can imagine the conceptual difficulties of comparing quantity productivity in an industry with highly differentiated products—say in units of airplanes, where some plants make commercial jets and others make gliders for hobbyists. Also, the industry's combination of high transport costs and ubiquity mean there are many ready-mixed plants spread throughout the country. This affords a greater amount of data variation with which to measure price effects.[27] Finally, ready-mixed plants are highly specialized; on average, well over 90 percent of their revenue comes from sales of ready-mixed concrete as opposed to other products like concrete block or pipe. Specialization is important because the physical output data in the CM are collected at the product level (by seven-digit SIC classification), while inputs are only measured at the plant level. Allocating a plant's inputs to specific products, a necessary step when computing physical TFP, involves some measurement error that shrinks as the output in

[27] There are roughly 5000 plants in any given year of the CM, 3000 of which we have physical quantity data for. Adding across the CMs spanning our sample, we have over 12,000 observations of quantity-based TFP and prices.

question accounts for a larger share of plant revenues.[28]

To gauge how nonattainment designations impact plant prices and technical efficiency levels, we estimate regressions similar to our benchmark specification using data only from ready-mixed concrete plants. Rather than merely exploring plants' (revenue-based) TFP, however, we also estimate specifications with plants' quantity-based TFP levels and their (logged) unit prices as dependent variables. Estimating a specification with revenue-based TFP offers a comparison to the benchmark results. Also, because it equals the sum of the other two dependent variables by construction, we can use the three regressions to decompose the estimated regulation-induced revenue TFP change into the components driven by quantity TFP and prices.[29]

There are a few other differences between these regressions and the benchmark specifications. We now set $I[nocaaa_{cpt}]$ equal to one if county c is declared in year t to be in nonattainment with standards for either Ozone, TSPs, or SO_2—the three pollutants for which ready-mixed concrete's industry group is a heavy emitter. Thus we focus on nonattainment designations that should specifically affect ready-mixed plants. Also, because only ready-mixed concrete plants are in the sample, the industry indicator $I[pollind_{ip}]$ is not separately identified from the constant. The specification measures nonattainment effects by comparing TFP levels and prices of ready-mixed plants in nonattaining counties to those in attaining counties. The key difference with the benchmark specification is that we cannot control for shocks to TFP common to polluters and non-polluters in nonattainment counties. However, we are still able to include division-by-ASM-panel and plant-level fixed effects in the specification.[30]

The results presented in Table 7 are telling. The revenue TFP impact of nonattainment in column (1) is economically small, at -0.6 percent, and statistically insignificant. This result suggests that ready-mixed concrete plants do not appear to suffer (revenue) TFP declines from regulatory action.

---

[28] Foster, Haltiwanger, and Syverson (2008) discuss these quantity-based TFP measurement issues in greater detail. Syverson (2008) offers a general discussion of the economics of the ready-mixed concrete industry.

[29] $TFP_{rev} = TFP_q + \ln(p)$ because ln(revenue) = ln(q) + ln(p) and the input terms in both TFP measures are the same. The revenue-based productivity measure used in our full sample also includes usually minor contributions from inventory accumulations and occasionally other typically small revenue sources such as contract work. In this section, we only use revenue from the plants' ready-mixed concrete sales to ensure that the identity holds.

[30] Here, because we only have data from the quinennial CMs, division-by-ASM-period fixed effects are equivalent to division-by-CM fixed effects.

The results in columns 2 and 3 provide an explanation. Specifically, column 2 indicates that the nonattainment designation is associated with price increases: on average, a statistically significant 2.7 percent increase in price. The effect on plants' quantity-based TFP (cubic yards per unit input) is contained in column (3). This equals, as it must by construction, the revenue-based TFP effect minus the price effect. While the estimated effect of -3.3 percent is insignificant (it has a p-value of 0.111), it implies as a point estimate a considerably larger effect of nonattainment on productivity than does the revenue-based TFP estimate.

These results make clear that the nonattainment designation is associated with declines in efficiency—or equivalently, ready-mixed concrete plants suffer output losses given their input choices. However, this decline is masked in revenue-based productivity measures, because the plants are able to raise prices in response to higher regulatory costs. So at least in the case of ready-mixed concrete plants, revenue-based TFP measures are positively biased estimates of nonattainment's true impact on plants' technical efficiencies and output. This result is consistent with standard economic theory about firm's pricing behavior when they have market power.

This case study offers a vivid exhibit of the sort of understatement of true productivity effects that might plague the revenue-based TFP estimates for the entire manufacturing sector. If a similar-sized effect (i.e., an additional productivity drop of around 2.7 percent) operates in the larger sample, the true effects of nonattainment designation—which the benchmark specification estimated at -2.6 percent in revenue-based measures—could be about twice as large. Since this evidence comes from just a single industry, we are reluctant to speculate about the magnitude and instead conclude that this case study provides evidence of a strong directional bias.

By augmenting these estimates with another data source, we are able to make broader-scoped estimates of the price measurement problem. The NBER-CES Manufacturing Industry Database contains annual price indices and industry-level TFP measures (logged total industry output minus a weighted sum of logged total industry inputs) for a panel of 459 four-digit SIC manufacturing industries from 1958 to 2005.[31] We use these data to compare price growth to productivity growth at the industry level to provide a first understanding of what fraction of productivity improvements (declines) are passed on as lower (higher) prices.

[31] The data available on the NBER website span 1958 to 1997; a recent "beta" version of an extension extended this to 2005. We use the four-factor TFP measures in the NBER database to be consistent with the definition of our plant-level TFP indexes, though similar results were found using the five-factor TFP measures (which treat equipment capital and structures capital as separate inputs).

We regressed industry price growth on industry TFP growth, and included industry fixed effects to compare price and TFP growth within industries. (The results were very similar when including cross-industry variation as well.) The coefficient on the change in industry TFP is -0.349 (s.e. = 0.014). That is, for every one-percent increase (decrease) in average TFP in an industry, average industry price falls (rises) by roughly 0.35 percent.[32]

If we extend the findings from this regression to our results, it implies that for each one-percent increase in costs driven by productivity losses, firms raise prices by 0.35 percent. For our average estimated effect of nonattainment on revenue-based TFP of roughly 2.6 percent, the implied impact on true technical efficiency is then actually 2.6/(0.65) = 4.0 percent.[33] These estimates suggest that the lost output due to nonattainment designation could therefore be in the neighborhood of $17.2 billion annually (in 2010 dollars).

B. Regulation's Impact on Survivorship

It is important to note that the estimated TFP effects are conditional on survival. But plants experiencing the largest negative productivity shocks from nonattainment are the most likely to cease operations and exit the industry. This will lead such plants to disappear from our sample.[34] This implies that the paper's estimates will understate the true TFP impact of nonattainment. This subsection attempts to shed some light on the magnitude of the understatement due to this survivorship problem.

We first test whether nonattainment status predicts exit for polluting plants. Table 8 shows the results from regressing an indicator for plant exit on the same set of explanatory variables included in our benchmark specification.[35] Again, we report only the estimates of the interaction of pollution attainment and polluting industry indicators. The pooled-pollutant

---

[32] Industry-level TFP measures are not subject to the price measurement problem because industry-level deflators are available, so industry TFP changes should reflect quantity changes.

[33] Recall $TFP_{rev} = TFP_q + \ln(p)$, so $\Delta TFP_{rev} = \Delta TFP_q + \Delta\ln(p)$. If, as we estimate, $\Delta\ln(p) = -0.35*\Delta TFP_q$, then $\Delta TFP_{rev} = \Delta TFP_q - 0.35*\Delta TFP_q$, or equivalently $\Delta TFP_q = \Delta TFP_{rev}/0.65$.

[34] A negative correlation between productivity and exit probability is virtually ubiquitous in the empirical literature on the subject. It has been found in studies using samples from wide ranges of industries, time periods, and countries. See Bartelsman and Doms (2000) and Syverson (2011) for reviews of the empirical literature on plant-level productivity.

[35] The sample includes all plant-year observations in our benchmark sample for years 1977 and after, as 1977 is the first year for which exit indicators can be constructed from the Longitudinal Business Database. This leaves a sample of approximately 896,000 plant-year observations.

specification in column 1 indicates that plants in heavy-emitting industries are more likely to exit when their county is declared in nonattainment for one of the pollutants they emit. The increment in exit rates is 0.4 percent. Given the mean annual exit rate in the sample of about 3.7, this is about a 10 percent increase in exit likelihood. The pollution-specific estimates indicate a marginally significant 0.4 percent increase in exit associated with ozone nonattainment, and a significant decrease in exit probability associate with CO nonattainment (which is consistent with the positive TFP effects we have estimated from CO nonattainment). The TSP and SO_2 estimates are more muddled, however, with insignificant coefficients of mixed sign.

A standard approach to address this sort of selection issue is to include a Heckman (1979) style correction term (estimated in a first-stage survival regression) in the second-stage regressions estimating the TFP effects of interest. However, our specifications include plant-level fixed effects, and the average panel length in our sample is rather short (6.2 years), raising a potentially severe incidental parameters problem in the first-stage survival equation (Heckman 1981). Additionally, we face many practical boundaries in estimating a nonlinear survival specification with over one million observations and 200,000 plant fixed effects.

We therefore take an alternate approach based on logic similar to that of the traditional approaches but that is feasible to implement. Namely, we use observed plant exit rates and the distribution of TFP innovations among surviving plants to infer the unobservable in which we are interested: the distribution of TFP innovations among exiting plants (and by implication, the average TFP innovation among all plants, survivors and exiters). The difference between the conditional (on plant survival) and unconditional average TFP innovations is the correction we need to apply to our selection-biased estimates to obtain the true TFP impact of nonattainment.

Our methodology is based on the notion that the distribution of TFP innovations among surviving plants is a truncation of the unconditional distribution. The truncation point is the critical TFP innovation such that plants suffering more negative TFP innovations (i.e., larger TFP drops) exit, as their present value from operating falls below their scrap value. While the shape of the TFP innovation distribution below the threshold is unknown, we do know the exit rate. If we assume that exiters had inferior TFP innovations compared to survivors, we can infer enough about the shape of the unconditional distribution to measure the difference between the averages of the survivorship-biased TFP innovation distribution and the unconditional distribution.

Note, though, that this threshold TFP innovation is conditional on the plant's TFP level. Profitability depends on a business's productivity level, not its innovation.[36] Indeed, there is an inverse relationship between a plant's TFP level and the TFP decline necessary to cause exit. Simply put, higher-TFP plants have further to fall before they become unprofitable.

While the same estimation issues discussed above prevent us from adjusting for computing this threshold TFP innovation plant-by-plant, we can make a discrete approximation. If we cut the plant TFP *level* distribution into contiguous sections of sufficiently small size, all plants in each section have roughly the same cutoff productivity innovation value. We can then compute the selection correction section-by-section, and average over these to find the correction for the entire sample.

We compute the selection correction within each section as follows. We assume that the TFP innovation distribution that we observe among surviving plants within each section is a truncation of an unconditional innovation distribution that also includes the (unobserved) TFP innovations of plants that exited. As long as all exiters have TFP innovations below (i.e., more negative than) the median of the unconditional distribution, we can easily back out this median using the exit rate and the observable median of the truncated distribution.[37] Specifically, for an exit rate (the fraction of plants that exit) equal to x, the median of the unconditional distribution is simply the $(50 - x)^{\text{th}}$ quantile of the truncated distribution. For example, if 6 percent of plants in the segment exit, the difference between the 50^{th} and 44^{th} percentiles of TFP innovations among surviving plants in the segment equals the difference in the medians of the truncated and unconditional distributions. This difference approximates the difference in means between the truncated and unconditional distributions, and we use it as the selection correction for TFP changes in the segment.

While this method involves a compromise in that we approximate a correction in means using medians instead, it requires only very weak assumptions about the unconditional distribution of TFP innovations. Estimating the difference in means directly would involve assuming both that the unconditional distribution has a particular shape and that all plants that exited had more negative TFP innovations than the lowest innovation observed among survivors.

---

[36] See Jovanovic (1982), Hopenhayn (1992), Melitz (2003), and Asplund and Nocke (2006) for models that predict a negative relationship between productivity and survival and a threshold productivity level that determines exit.

[37] This condition can hold as long as fewer than 25 percent of plants exit, which easily holds true in our data.

The second of these assumptions, in particular, is very strong, given the myriad factors other than TFP that potentially affect exit. Hence it is likely to lead to an overstatement of the true difference in the means of the truncated and unconditional TFP innovation distributions. Using the median-based correction lets us avoid these assumptions.[38] The appendix describes the algorithm to calculate these corrections.

Table 9 reports the results. Panel A begins by displaying the average exit rates for deciles of the within-industry TFP distributions. As expected, the lowest-productivity plants are most likely to exit, and average exit rates tend to decline as TFP levels rise. On average, 9.2 percent of the plants in the lowest decile of their industry's TFP distribution exit in a given year, almost three times the rate among the other deciles. The inverse relationship between TFP levels generally holds across the higher deciles, though not monotonically and exit rates are considerably lower than among the least productive plants.

Panel B of Table 9 shows average exit rates and the calculated selection corrections for heavy emitters of each pollutant. The implied selection corrections are notably consistent across pollutants, ranging from -0.7 to -0.8 percent. These values imply that survivorship bias causes our composite pollutant benchmark estimate to understate nonattainment designations' effect on (revenue-based) TFP by over 20 percent (i.e., the estimated conditional change is -2.6 percent, but the unconditional change is around -3.3 percent). The understatements in percentage terms are even larger for the TSP- and SO_2-specific estimates, and the selection-corrected (positive) CO effect would drop by about 40 percent.

Overall, this exercise suggests that survivorship bias is relevant in this setting. In particular, it indicates that the paper's estimated regulation effects are likely to be underestimated. In total, the application of corrections for the confounding of price changes and sample selection on survival produce a 4.8 percent estimated decline in TFP for polluting plants in regulated areas. This corresponds to an annual economic cost from the regulation of manufacturing plants of roughly $21 billion.

[38] That said, we have computed mean-based corrections based on the assumptions that the observed TFP innovations among survivors in the truncated distribution are from a normal unconditional distribution and that this all exiters have TFP innovations in the truncated portion of the distribution. We found selection correction factors that were about 3 times as large as those found using the median approach above. Therefore we use the median based corrections both because they require fewer assumptions and because they offer a more conservative estimate of what are otherwise unobservable TFP effects.

VIII. Discussion and Conclusions

This paper has produced the first large-scale estimates of the economic costs of environmental regulations and finds that they are not insubstantial. Among surviving polluting plants, the nonattainment designation is associated with a roughly 2.6 percent decline in total factor productivity. The regulations governing ozone have particularly large negative effects on productivity, though negative effects are also evident among emitters of particulates and sulfur dioxide. Carbon monoxide nonattainment, on the other hand, appears to increase measured TFP, especially among refineries. Dynamics matter somewhat; indeed, the impacts of a year's nonattainment designation on a plant's TFP are detectable two years later. These results are robust to a number of alternative samples and specifications. Overall, the productivity losses among surviving plants in nonattainment counties correspond to annual lost output on the order of $11.0 billion in 2010 dollars.

These estimates, however, are likely underestimates for the three reasons discussed earlier. If we apply our estimated corrections for two of the sources of understatement—price mismeasurement and survivorship bias—the implied losses are considerably larger. We calculate that survivorship bias results in an understatement of the revenue-based TFP loss of about 0.7 percent, so the total average effect is a roughly 3.3 percent drop. Further, our inability to account for cost-driven price increases imply that the impact on plants' quantity produced is on the order of 54 percent higher than the revenue-based TFP impact (1/0.65 = 1.538). Applying this additional correction implies a total TFP loss for dirty manufacturing plants in nonattainment counties of 4.8 percent.

Are these TFP losses big? They correspond to annual lost output in the manufacturing sector of about $20.8 billion in 2010 dollars. This is roughly 8.8 percent of average manufacturing sector profits over this period. It is noteworthy that the estimated costs are substantially larger than the costs borne by workers in polluting industries (Walker 2012).[39] At least in the case of the Clean Air Act and the manufacturing sector, it seems reasonable to conclude that the claim that environmental regulations are costless or even beneficial for the

[39] Walker (2012) finds that the workers employed in regulated plants in counties that were newly designated nonattainment due to changes in the Clean Air Act from its 1990 Amendments lost approximately $9 billion in total over the next 8 year due to foregone earnings from periods of unemployment and employment at reduced wages.

regulated is contradicted by the available data.[40]

More broadly, the paper demonstrates that the recent advances in the ability to estimate TFP in micro data sets makes it feasible to estimate transparently the economic costs of regulation that are borne by businesses. In principle, this approach can be applied to determine the costs of regulations that govern firm behavior in a wide range of contexts. We envision similar exercises being fruitful in areas that regulate worker and labor conditions, health and safety legislation, and more broadly any public policy that alters the composition of inputs on the expansion path.

[40] It is tempting to use these results to conduct a cost-benefit analysis of the CAAAs, however there are several missing pieces that prevent a complete accounting. For example, the costs of the CAAAs in other industries, notably the electricity sector (noting that the incidence may be on households and businesses), remain largely unknown. On the other side of the ledger, the CAAAs has led to tremendous improvements in air quality: this has reduced rates of mortality and morbidity, made many cities in the United States more appealing places to live, and reduced purchases of drugs that protect people from the harms of air pollution (e.g., Chay and Greenstone 2003 and 2005; Sieg et al. 2004; Currie and Neidell 2005; Tra 2010; and Deschenes, Greenstone, and Shapiro 2012). These efforts have shed light on the CAAAs' benefits, but more research is necessary for a full accounting.

Appendix

A. Construction of Total Factor Productivity Index

We describe here details on the construction of our production variables.

Output. Plant output is its inventory-adjusted total value of shipments, deflated to 1987 dollars using industry-specific price indexes from the NBER Productivity Database.

Labor Hours. Production worker hours are reported directly in the Annual Survey of Manufactures (ASM) and Census of Manufactures (CM) microdata. To get total plant hours, we multiply this value by the plant's ratio of total salaries and wages to production worker wages. This, in essence, imputes the hours of non-production workers by assuming that average non-production worker hours equal average production worker hours within plants.

Real Materials and Energy Use. Materials and energy inputs are plants' expenditures on each, as reported in the ASM or CM, divided by their respective industry-level deflators from the National Bureau of Economic Research Productivity Database.

Capital. We construct capital stocks using the perpetual inventory method. Initial-year capital stocks are constructed by deflating plants' reported book values of capital by the book-to-real value ratio for the corresponding three-digit industry (the industry-level equipment and structures ratios are from published Bureau of Economic Analysis data). Thereafter, plants' reported investments (deflated by industry-specific investment goods price indices) are added to the capital stock, and depreciation (using industry-specific rates from the BEA) is subtracted. This process is repeated for every year the plant is in existence to get a capital stock series. For the 1988-1993 ASM panel, initial capital stocks could not be constructed for new plants because Census stopped collecting book values in the ASM. We instead use reported stocks from the 1992 CM, which did collect plants' book values of capital stocks, and work backwards using reported investments and imputed depreciation as described above. For our sample of ready-mixed concrete plants from the CM, we simply use the deflated values of plants' reported book values of capital. Capital is unavailable after 1993.

Total Factor Productivity. Plant TFP is its logged output minus a weighted sum of its logged labor, capital, materials, and energy inputs. That is,

$$TFP_{it} = y_{it} - \alpha_{lt} l_{it} - \alpha_{kt} k_{it} - \alpha_{mt} m_{it} - \alpha_{et} e_{it}$$

where the weights α_j are the input elasticities of input $j \in \{l, k, m, e\}$. Output is the plant's inventory-adjusted total value of shipments deflated to 1987 dollars. Inputs are plant-specific but the input elastiticies are measured using industry-level input cost shares. These cost shares are computed using reported industry-level labor, materials, and energy expenditures from the NBER Productivity Database. Capital expenditures are the reported plant's capital stocks multiplied by their respective BLS capital rental rates in the corresponding two-digit industry.[41]

[41] Capital rental rates are from unpublished data constructed by the Bureau of Labor Statistics for use in computing their Multifactor Productivity series. Formulas, related methodology, and data sources are described in U.S. Bureau of Labor Statistics (1983) and Harper, Berndt, and Wood (1989).

B. Estimating the Selection Corrections

We compute our estimated selection corrections described in Section VII.B using the following algorithm.

1. Estimate the TFP innovations for all plants in our sample that survive to the next year by regressing leads of these plants' TFP levels on our array of fixed effects from the benchmark specification. Because this array includes plant fixed effects, the predicted values from this regression are the surviving plants' TFP innovations.

2. Divide all plants in a given four-digit-industry-year cell, both those that will survive to the next year and those that will cease operating, into deciles by their TFP level. Doing this separately by industry-year controls for the many variations in the other forces that drive exit rates across industries and time. These deciles are the sections described above that ensure all plants within them have roughly the same threshold productivity innovation for exit.

3. Compute the exit rate for the decile, which is simply the fraction of plants that exit by the following year.

4. Using the estimated TFP innovations for survivors, the categorization of survivors into deciles, and the corresponding decile-specific exit rates, compute the implied differences in means and medians of the conditional and unconditional TFP innovation distributions as described in the main text. Compute these differences for every decile.

5. Compute the real-revenue-weighted average of these decile-industry-year specific TFP corrections to obtain the implied unconditional TFP innovation for the entire sample.

6. Do this entire process separately for plants subject to nonattainment ratings for each pollutant, so the adjustments are pollution-specific, just as the TFP effect estimates are.

References

Asplund, Marcus and Volker Nocke. "Firm Turnover in Imperfectly Competitive Markets." *Review of Economic Studies*, 73(2), 2006, 295-327.

Bartelsman, Eric J. and Mark Doms. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature*, *38*(3), 2000, 569-95.

Becker, Randy and Vernon Henderson. "Effects of Air Quality Regulations on Polluting Industries." *Journal of Political Economy*, 108(2), 2000, 379-421.

Berman, Eli and Linda T. M. Bui. "Environmental Regulation and Productivity: Evidence from Oil Refineries." *Review of Economics and Statistics*, 83(3), 2001, 498-510.

Biorn, Erik, Rolf Golombek, and Arvid Raknerud. "Environmental Regulations and Plant Exit: A Logit Analysis Based on Establishment Panel Data." *Environmental and Resource Economics*, 11(1), 1998, 35-59.

Chay, Kenneth Y. and Michael Greenstone. "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession." *Quarterly Journal of Economics*, 118(3), 2003, 1121-67.

Chay, Kenneth Y. and Michael Greenstone. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy*, 113(2), 2005, 376-424.

Cleveland, W.S. and T.W. Graedel. "Photochemical Air Pollution in the Northeast United States." *Science*, 204, 1979, 1273-78.

Cleveland, W.S., B. Keiner, J.E. McRae, and J.L. Warner. "Photochemical Air Pollution: Transport from the New York City Area into Connecticut and Massachusetts." *Science*, 191, 1976, 179-181.

Cohen, Mark A. "Monitoring and Enforcement of Environmental Policy." Unpublished Paper, Vanderbilt University, Owen Graduate School of Management, 1998.

Currie, Janet, and Matthew Neidell. 2005. "Air Pollution and Infant Health: What Can We Learn From California's Recent Experience?" *Quarterly Journal of Economics*, 120(3): 1003-1030.

Deschenes, Olivier, Michael Greenstone, Joseph Shapiro. 2012. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program and Ozone Reductions." Mimeograph.

Foster, Lucia, John Haltiwanger, and Chad Syverson. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1), 2008.

Gollop, Frank M. and Mark J. Roberts. “Environmental Regulations and Productivity Growth: The Case of Fossil-Fueled Electric Power Generation.” *Journal of Political Economy*, 91(4), 1983, 654-74.

Gray, Wayne B. and Ronald J. Shadbegian. “Environmental Regulation, Investment Timing, and Technology Choice.” *Journal of Industrial Economics*, 46(2), 1998, 235-56.

Gray, Wayne B. and Ronald J. Shadbegian. “Plant Vintage, Technology, and Environmental Regulation.” *Journal of Environmental Economics and Management*, 46(3), 2003, 384-402.

Greenstone, Michael. “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), 2002, 1175-1219.

Greenstone, Michael. “Did the Clean Air Act Amendments Cause the Remarkable Decline in Sulfur Dioxide Concentrations?” *Journal of Environmental Economics and Management*, 47(3), 2004, 585-611.

Greenstone, Michael, Richard Hornbeck and Enrico Moretti. “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings,” *Journal of Political Economy*, 2010, 118 (3): 536-598.

Harper, Michael, Ernst R. Berndt, and David O. Wood. “Rates of Return and Capital Aggregation using Alternative Rental Prices,” in *Technology and Capital Formation,* Dale W. Jorgenson and Ralph Landau (eds.), MIT Press, 1989.

Heckman, James. “Sample Selection Bias as a Specification Error.” *Econometrica*, 47(1), 1979, 153–61.

Heckman, James. “The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process.” In: Manski, C.F., McFadden, D. (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press, Cambridge, MA. 1981.

Henderson, J. Vernon. “Effects of Air Quality Regulation.” *American Economic Review*, 86(4), 1996, 789-813.

Hopenhayn, Hugo A. “Entry, Exit, and Firm Dynamics in Long Run Equilibrium.” *Econometrica*, 60(5), 1992, 1127-50.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. “Earnings Losses of Displaced Workers.” *American Economic Review*, 83(4), 1993, 685-709.

Jaffe, Adam B., Steven R. Peterson, Paul R. Portney, and Robert N. Stavins. “Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us?” *Journal of Economic Literature*, 33(1), 1995, 132-63.

Johansen, Leif. “Substitution versus Fixed Production Coefficients in the Theory of Economic Growth: A Synthesis.” *Econometrica*, 27(2), 1959, 157-176.

Jovanovic, Boyan. “Selection and the Evolution of Industry.” *Econometrica*, *50*(3), 1982, 649-70.

List, John A., Daniel L. Millimet, Per G. Fredriksson, and W. Warren McHone. “Effects of Environmental Regulations on Manufacturing Plant Births: Evidence from a Propensity Score Matching Estimator.” *Review of Economics and Statistics*, 85(4), 2003, 944-52.

List, John A.; Millimet, Daniel L.; and McHone, Warren (2004) "The Unintended Disincentive in the Clean Air Act," *Advances in Economic Analysis & Policy*: Vol. 4: Iss. 2, Article 2.

Lleras-Muney, Adriana. “The Needs of the Army: Using Compulsory Relocation in the Military to Estimate the Effect of Environmental Pollutants on Children’s Health.” *Journal of Human Resources*, 35(3), 2010, 549-590.

Melitz, Marc J. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71(6), 2003, 1695-1725.

Nadeau, Louis W. “EPA Effectiveness at Reducing the Duration of Plant-Level Noncompliance.” *Journal of Environmental Economics and Management*, 34(1), 1997, 54-78.

Pashigian, B. Peter. “The Effect of Environmental Regulation on Optimal Plant Size and Factor Shares.” *Journal of Law and Economics*, 27(1), 1984, 1-28.

Peirce, J. Jeffrey, P. Aarne Vesilind, and Ruth Weiner. “Environmental Pollution and Control, Fourth Edition.” Woburn, MA: Butterworth-Heinemann, 1998.

Porter, Michael E. “America’s Green Strategy.” *Scientific American*, 264(4), 1991, 168.

Ryan, Stephen P. “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica* (forthcoming).

Sieg, Holger, V. Kerry Smith, H. Spencer Banzhaf, and Randy Walsh. “Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods.” *International Economic Review*, 45(4), 2004, 1047-77.

Syverson, Chad. “Markets: Ready-Mixed Concrete.” *Journal of Economic Perspectives*, 22(1), 2008, 217-233.

Syverson, Chad. “What Determines Productivity?” *Journal of Economic Literature*, 49(2), 2011, 326-65.

Tra, Constant I. “A Discrete Choice Equilibrium Approach to Valuing Large Environmental Changes.” *Journal of Public Economics*, 94(1-2), 2010, 183-96.

U.S. Bureau of Labor Statistics. “Trends in Multifactor Productivity: 1948-81.” Bulletin 2178. Washington D.C., Government Printing Office, 1983.

Walker, W. Reed. “The Transitional Costs of Sectoral Reallocation: Evidence From the Clean Air Act and the Workforce.” Working paper, 2012.

Figure 1. Share of Polluting Industries' Output Produced in Counties Nonattaining in Pollutant

A. Any Pollutant

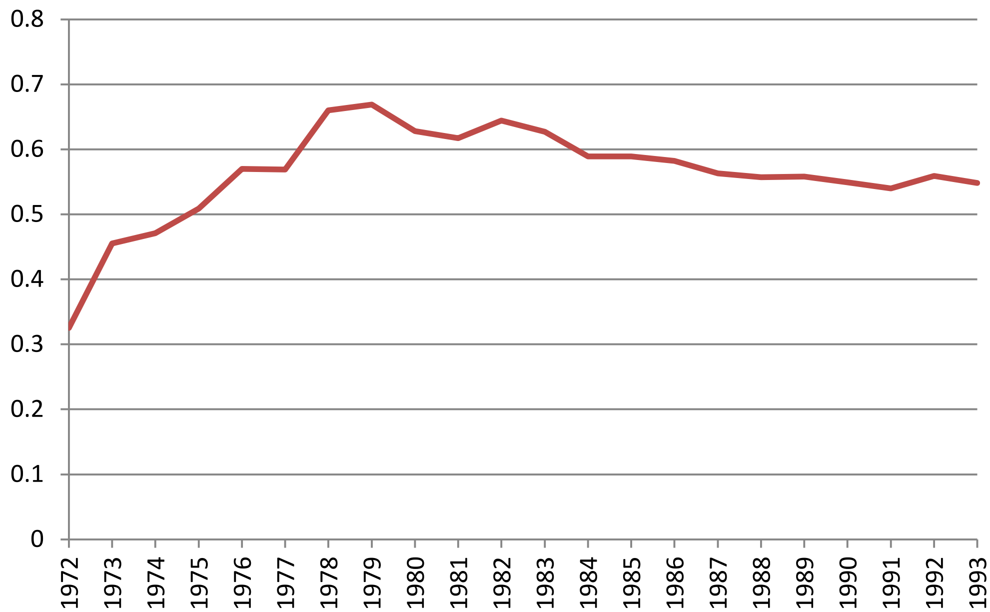

B. By Pollutant

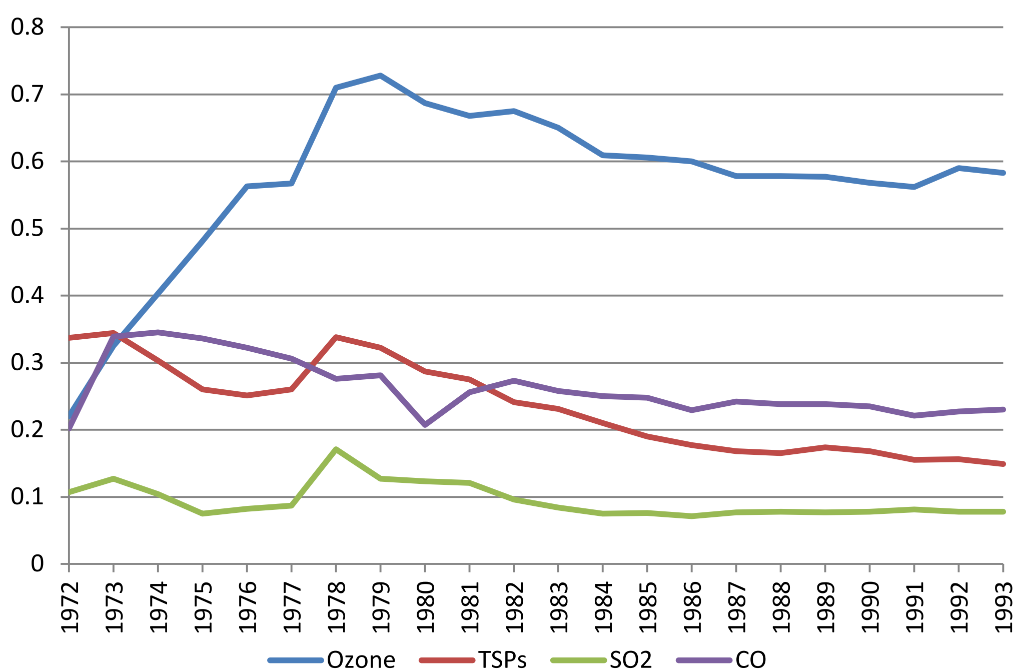

Table 1. Polluting Industry Groups and Their Pollutants

Industry (applicable SIC codes)	Pollutant
Pulp and paper (2611–31)	$CO/O_3/SO_2$/TSPs
Organic chemicals (2861–69)	O_3
Petroleum refining (2911)	$CO/O_3/SO_2$
Rubber and miscellaneous plastic products (30)	O_3
Stone, clay, glass, and concrete (32)	O_3/SO_2/TSPs
Iron and steel (3312–25, 3321–2)	CO
Nonferrous metals (333–34)	CO/SO_2

Notes: This table, based on information in Greenstone (2002), shows industries that are classified as heavy emitters of one or more of the four primary pollutants the CAAA covers, and which pollutant(s) they emit. We consider all plants in these industries to be heavy emitters and subject to CAAA abatement mandates if their county is declared in nonattainment.

Table 2. Annual Changes in Attainment Status

A. Any Pollutant

		Current Nonattainment: No	Current Nonattainment: Yes
Lagged Nonattainment	No	52,390	1,461
	Yes	1,181	10,929

B. O_3

		Current Nonattainment: No	Current Nonattainment: Yes
Lagged Nonattainment	No	55,962	1,124
	Yes	640	8,235

C. TSPs

		Current Nonattainment: No	Current Nonattainment: Yes
Lagged Nonattainment	No	61,217	715
	Yes	878	3,151

D. SO_2

		Current Nonattainment: No	Current Nonattainment: Yes
Lagged Nonattainment	No	64,663	172
	Yes	157	969

E. CO

CO		Current Nonattainment: No	Current Nonattainment: Yes
Lagged Nonattainment	No	63,034	290
	Yes	243	2,394

Notes: This table shows the distribution from 1972-1993 of counties' lagged and current nonattainment status. The reported numbers are the category-specific count of from a total number of 65,961 county-years.

Table 3: TFP Effects of Nonattainment, Core Specifications

Pollutant	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Any	-0.024**		-0.023**		-0.026**		-0.044**	
	(0.003)		(0.003)		(0.006)		(0.007)	
O_3		-0.018**		-0.018**		-0.022**		-0.057**
		(0.004)		(0.004)		(0.007)		(0.008)
TSPs		0.010**		0.009**		-0.013*		-0.011
		(0.004)		(0.004)		(0.007)		(0.008)
SO_2		0.000		-0.002		-0.016		-0.021*
		(0.006)		(0.006)		(0.010)		(0.011)
CO		-0.021**		-0.024**		0.017*		0.022**
		(0.005)		(0.005)		(0.009)		(0.010)
4-Digit SIC x Year	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Census Div x Year	No	No	Yes	Yes	No	No	Yes	Yes
2-Digit SIC x Period	No	No	No	No	Yes	Yes	No	No
Census Div x Period	No	No	No	No	Yes	Yes	No	No
Plant	No	No	No	No	Yes	Yes	Yes	Yes
R^2	0.766	0.766	0.801	0.801	0.887	0.887	0.887	0.887

Notes: This table reports the results from the estimation of alternative versions of equation (1), which involves the regression of plants' TFP levels on nonattainment indicators, polluting industry indicators, and their interaction, along with alternative sets of fixed effects that are noted in the row headings at the bottom of the table. In the row headings "period" refers to 5-year ASM panel periods and SIC refers to industry following the standard industrial classification system.. The entries in the table are the coefficients and standard errors (in parentheses) of the estimates of the interaction of pollutant nonattainment attainment and polluting industry indicators. Observations are weighted by the product of real output and the ASM weight so that the estimates are representative of effects on aggregate manufacturing activity. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. N ≈ 1,185,000 (approximate sample sizes are used to eliminate confidential data disclosure issues across samples). See the text for further details.

Table 4: TFP Effects of Nonattainment—Dynamic Specifications

Number of Lags	Zero	Zero	Zero	Zero	One	One	Two	Two
Pollutant	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Any	-0.026**		-0.017**		-0.020**		-0.031**	
	(0.006)		(0.007)		(0.008)		(0.009)	
O_3		-0.022**		-0.003		-0.008		-0.013
		(0.007)		(0.008)		(0.009)		(0.009)
TSPs		-0.013*		-0.013		-0.019*		-0.019*
		(0.007)		(0.009)		(0.010)		(0.011)
SO_2		-0.016		-0.021*		-0.024*		-0.029*
		(0.010)		(0.012)		(0.013)		(0.015)
CO		0.017*		0.018*		0.020		0.009
		(0.009)		(0.011)		(0.013)		(0.014)
Approx. N	1,185,000	1,185,000	800,000	800,000	800,000	800,000	800,000	800,000
R^2	0.887	0.887	0.898	0.898	0.898	0.898	0.898	0.898

Notes: This table shows the results of estimating specifications similar to our benchmark, except as noted in the column headers some specification include various lags of the pollutant nonattainment indicators and their interaction with polluting industry indicators. The full sample is used in columns [1]and [2] and the remaining columns use the sample of plant observations that result from imposing the restriction that plants must be operating for at least three years which is necessary for the two lag specifications. The entries in the table are the sum of the current and lagged interaction coefficients and their standard errors in parentheses. Results for specifications including zero, one, and two lags are reported (see column headers). Observations weighted by the product of real output and ASM weight. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. See the Notes to Table 3 and the text for further details.

Table 5: TFP Effects of Nonattainment—Robustness Checks

	Excluding industry switchers		7 percent cutoff for "heavy" polluters		12 percent cutoff for "heavy" polluters	
Pollutant	[1]	[2]	[3]	[4]	[5]	[6]
Any	-0.016**		-0.021**		-0.023**	
	(0.007)		(0.004)		(0.007)	
O_3		-0.011		-0.017**		-0.021**
		(0.009)		(0.006)		(0.008)
TSPs		-0.027**		0.000		0.013*
		(0.008)		(0.007)		(0.007)
SO_2		-0.022**		-0.007		-0.016
		(0.011)		(0.009)		(0.010)
CO		0.023**		0.018*		0.002
		(0.010)		(0.009)		(0.012)
Approx. N	807,000	807,000	1,185,000	1,185,000	1,185,000	1,185,000
R^2	0.910	0.910	0.887	0.888	0.887	0.887

Notes: This table shows the results of estimating specifications similar to our benchmark, except with changes in sample or variable definitions. Columns 1 and 2 exclude from the sample any plants that change four-digit-SIC industries during the sample. Columns 3 and 4 define heavy polluting plants as those in industry groups that account for at least seven percent of industrial emissions of one of our four pollutants, rather than the ten percent cutoff in our benchmark sample. Columns 5 and 6 use a more stringent 12 percent cutoff to define heavy polluters. Observations weighted by the product of real output and ASM weight. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. See the text and notes to Table 3 for further details.

Table 6: TFP Effects of Nonattainment, by Industry

A. Effect of Nonattainment in for Any Pollutant

Pollutant-Industry	Estimated Effect	Average annual output, 1972-93 (billions of 2010$)	F-test for H_0 that nonattainment effects are equal across industries
Pulp and paper	0.025** (0.009)	$26.6	
Organic chemicals	-0.168** (0.019)	54.6	
Petroleum refining	-0.005 (0.013)	152.3	
Rubber	-0.024** (0.008)	59.6	F-statistic: 15.9 p-value: 0.000
Stone, clay, glass	-0.011 (0.007)	56.3	
Iron and steel	0.005 (0.014)	51.6	
Nonferrous metals	-0.063** (0.026)	11.4	
N	1,185,000		
R^2	0.888		

Notes: This table shows the results of estimating a specification similar to the benchmark, except breaking out pollution-specific nonattainment effects by specific industry group rather than pooling all heavy emitters together. Observations are weighted by the product of real output and ASM weight. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. The table also reports the results of an F-test for equality of effects across all industry groups that emit a particular pollutant. See the text and Table 3 for further details.

Table 6 (cont.): TFP Effects of Nonattainment, by Industry

B. Effect of Nonattainment, Pollutant-by-Pollutant

Pollutant-Industry	Estimated Effect	F-test for H_0 that pollutant nonattainment effects are equal across industries
Pulp and paper/O_3	0.052**	
	0.010)	
Organic chemicals/O_3	-0.168**	
	(0.019)	
Petroleum refining/O_3	0.006	O_3 F-statistic: 30.6
	(0.013)	O_3 p-value: 0.000
Rubber/O_3	-0.024**	
	(0.008)	
Stone, clay, glass/O_3	0.013*	
	(0.007)	
Pulp and paper/TSPs	-0.024*	
	(0.012)	TSP F-statistic: 0.310
Stone, clay, glass/TSPs	-0.016**	TSP p-value: 0.579
	(0.008)	
Pulp and paper/SO_2	0.026	
	(0.020)	
Petroleum refining/SO_2	-0.027*	
	(0.014)	SO_2 F-statistic: 2.09
Stone, clay, glass/SO_2	-0.006	SO_2 p-value: 0.099
	(0.013)	
Nonferrous metals/SO_2	-0.039	
	(0.029)	
Pulp and paper/CO	0.012	
	(0.016)	
Petroleum refining/CO	0.028**	
	(0.014)	CO F-statistic: 2.25
Iron and steel/CO	0.005	CO p-value: 0.080
	(0.014)	
Nonferrous metals/CO	-0.058*	
	(0.031)	
N	1,185,000	
R^2	0.888	

Notes: This table reports the results from the estimation of a specification similar to the benchmark one, except breaking out pollution-specific nonattainment effects by specific industry group rather than pooling all heavy emitters of a particular pollutant together. Observations are weighted by the product of real output and ASM weight. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. The table also reports the results of an F-test for equality of effects across all industry groups that emit a particular pollutant. See the text and Table 3 for further details.

Table 7: TFP and Price Effects of Nonattainment, Ready-Mixed Concrete Plants

	Dependent Variable		
Pollutant	Revenue TFP	ln(price)	Physical Quantity TFP
Nonattainment for O_3, TSPs, and/or SO_2	-0.006 (0.019)	0.027** (0.010)	-0.033 (0.021)
R^2	0.635	0.660	0.649

Notes: This table reports the results from the estimation of a specification similar to the benchmark specification, but with several differences. First, the sample only includes ready-mixed concrete plants from the Census of Manufactures for which we observe revenues and output quantities in physical units (cubic yards). Since the sample is restricted to emitters, it is impossible to separately identify the industry indicator $I[pollind_{ip}]$ and the constant. Consequently, the parameter of interest is the coefficient on the nonattainment variable. Second, we estimate separate regressions (all with the same explanatory variables) for three different dependent variables: the plant's revenue-based TFP (revenue per unit input, our dependent variable in the benchmark specification), the plant's quantity-based TFP (physical units of output per unit input), and the natural log of the plant's average unit price (dollars per cubic yard). The three dependent variables are linked by the following accounting identity: revenue TFP ≡ quantity TFP + ln(price). Third, nonattainment is defined as a nonattainment designation in one or more of the three pollutants of which concrete plants are considered heavy emitters: O_3, TSPs, and SO_2. As with the benchmark specification, observations are weighted by the product of real output and ASM weight, and standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. N ≈ 12,000. See the text for further details.

Table 8: Exit and Nonattainment

Pollutant	[1]	[2]
Any	0.0042**	
	(0.0021)	
O_3		0.0039*
		(0.0023)
TSPs		-0.0050
		(0.0038)
SO_2		0.0043
		(0.0055)
CO		-0.0115**
		(0.0050)
R^2	0.413	0.413

Notes: This table reports the results from regressing an indicator for plant exit on pollution nonattainment indicators, polluting industry indicators, and their interaction, along with fixed effects at the two-digit-SIC-by-ASM-panel, Census-division-by-ASM-panel, and plant levels. The values in the table are the coefficients and standard errors of the estimates of the interaction of pollution attainment and polluting industry indicators. Standard errors are clustered by county-year. An asterisk denotes significance at the ten percent level; two asterisks denote significance at the five percent level. The sample includes all plant-year observations in our benchmark sample for years 1977 and after (1977 is the first year for which exit indicators can be constructed from the Longitudinal Business Database). N ≈ 896,000 (approximate sample sizes are used to eliminate confidential data disclosure issues across samples).

Table 9: Selection Corrections

A. Exit Rates by TFP Decile for Plants Facing Nonattainment Designation

Decile	Average Exit Rate (percent)
1st (lowest)	9.15
2nd	3.85
3rd	3.25
4th	2.93
5th	2.63
6th	3.11
7th	2.77
8th	2.91
9th	2.63
10th	3.67

Notes: This panel shows the average exit rates (the percent of plants that cease operations in the following year) by the plant's decile within its four-digit-SIC-industry's TFP distribution.

B. Selection Correction Factors (I.e., Additional TFP Change from Nonattainment Designation)

Pollutant	Average exit rate	Selection Correction for TFP Change
Any	3.68	-0.007
O_3	3.65	-0.007
TSPs	4.05	-0.008
SO_2	3.73	-0.007
CO	3.60	-0.008

Notes: This table shows average exit rates and our calculated survivorship-bias corrections to our benchmark estimates of nonattainment's TFP effects by pollutant (including the pooled, "any pollutant" specification). These corrections, shown in the rightmost column, are the estimated additional changes in TFP in heavy-emitter plants associated with nonattainment. The full effect of nonattainment is obtained by adding these selection corrections to the estimated TFP effects from the specifications above, which had only surviving plants in the sample. See the text for details of the calculations.