Why is Pollution from U.S. Manufacturing Declining? The Roles of Trade, Regulation, Productivity, and Preferences*

Joseph S. Shapiro
Yale University
and NBER
joseph.shapiro@yale.edu

Reed Walker
University of California, Berkeley
IZA and NBER
rwalker@berkeley.edu

January 2015

Abstract

Between 1990 and 2008, emissions of the most common air pollutants from U.S. manufacturing fell by 60 percent, even as real U.S. manufacturing output grew substantially. This paper develops a quantitative model to explain how changes in trade, environmental regulation, productivity, and consumer preferences have contributed to these reductions in pollution emissions. We estimate the model's key parameters using administrative data on plant-level production and pollution decisions. We then combine these estimates with detailed historical data to provide a model-driven decomposition of the causes of the observed pollution changes. Finally, we compare the model-driven decomposition to a statistical decomposition. The model and data suggest three findings. First, the fall in pollution emissions is due to decreasing pollution per unit output within narrowly defined products, rather than to changes in the types of products produced or changes to the total quantity of manufacturing output. Second, the implicit pollution tax that rationalizes firm production and abatement behavior more than doubled between 1990 and 2008. Third, environmental regulation explains 75 percent or more of the observed reduction in pollution emissions from manufacturing.
JEL: F18, F64, H23, Q56

*We thank Joe Altonji, Costas Arkolakis, Randy Becker, Brian Copeland, Arnaud Costinot, Sharat Ganapati, Michael Greenstone, Sam Kortum, Arik Levinson, Bill Nordhaus, and seminar participants at the Census Bureau, Cornell, MIT, NBER EEE, UCE3, and Yale for useful comments and discussions. We also thank Adrian Fernandez, Vivek Sampathkumar, David Silver, and Yusuf Mercan for excellent research assistance. Shapiro acknowledges support from the Yale MacMillan Center and the Yale Program on Applied Economics and Policy Research. Walker acknowledges additional support from the Robert Wood Johnson Foundation and the UC office of the President through the University of California Center for Energy and Environmental Economics. All results have been reviewed by the U.S. Census Bureau to ensure that no confidential information is disclosed.

1 Introduction

Between 1990 and 2008, emissions of the most common air pollutants from U.S. manufacturing fell by 60 percent, even as real U.S. manufacturing output grew substantially. Figure 1 shows just how stark these environmental improvements have been. Between 1990 and 2000, the real value of U.S. manufacturing output grew by a third even as manufacturing's emissions of major regulated air pollutants like nitrogen oxides, particulate matter, sulfur dioxide, and volatile organic compounds fell on average by 35 percent. After 2000, growth in real manufacturing output slowed, even while manufacturing pollution emissions fell another 25 percentage points relative to 1990 levels.

Research suggests at least four possible explanations of these substantial improvements in U.S. air quality. First, U.S. manufacturing trade has grown substantially (Autor, Dorn, and Hanson, 2013; Pierce and Schott, 2012; Fowlie, Reguant, and Ryan, Forthcoming). When dirty industries like steel or cement move abroad, total U.S. pollution emissions may fall. Second, federal and state agencies require firms to install increasingly stringent pollution abatement technologies. Some research, for example, directly attributes national changes in air quality to the Clean Air Act and to other environmental regulations (Henderson, 1996; Chay and Greenstone, 2005; Correia, Pope, Dockery, Wang, Ezzati, and Dominici, 2013). Third, due to non-homothetic tastes or secular change, Americans may gradually choose to spend less on heavy manufactured goods and more on services and cleaner goods (Levinson and O'Brien, 2013). Finally, if manufacturers use fewer dirty inputs each year to produce the same outputs, then annual productivity growth could improve air quality. In support of this final explanation, Figure 2 shows a clear negative relationship between plant-level pollution per unit of output and total factor productivity in the U.S. manufacturing sector; as total factor productivity rises, pollution per unit of output falls.[1]

In order to explain the observed reduction in manufacturing pollution emissions, this paper develops a quantitative model of firms endogenously choosing investments in pollution abatement. Production and pollution abatement choices in the model depend on trade costs, pollution taxes, productivity, and consumer preferences. We take this model to data with two main objectives. First, we use results from the model, combined with actual pollution abatement and emissions decisions, to back out the implicit tax per unit of pollution emissions that firms face. U.S. federal, state, and local environmental regulations take many overlapping forms: command-and-control technology standards, cap-and-trade programs, and many others. Summarizing all of these policies as an implicit tax on pollution lets us quantify aggregate changes in the stringency of all air pollution regulation for U.S. manufacturing.

The paper's second objective is to decompose the underlying forces that have caused changes in pollution emissions from U.S. manufacturing. We do this in two complementary ways. We start with a statistical decomposition that decomposes changes in total manufacturing emissions into changes that can be explained by the total scale of manufacturing output, the composition of products produced, and the pollution intensity of a given set of products. We then turn to the quantitative model to explore how shocks to trade costs, environmental regulation, productivity, and preferences have each contributed to changes in air pollution

[1] See Appendix C.1 for further details on Figure 2. Similar relationships between plant-level pollution intensity or fuel intensity and management or export status appear in other data (Bloom, Genakos, Martin, and Sadun, 2010; Holladay, 2011; Martin, 2011).

emissions from U.S. manufacturing between 1990 and 2008. We use the quantitative model to evaluate a range of counterfactuals, such as how pollution emissions would have evolved if air pollution regulation had remained unchanged after 1990. Many researchers in environmental economics and international trade use quantitative models to forecast the future—they study untested policies such as a global 10 percent decrease in all trade barriers or a national carbon tax. Unlike such work, this paper uses a model to interpret the past—we quantify how different kinds of economic shocks led to observed changes in actual pollution emissions.

The paper obtains three main results. First, changes in the scale of manufacturing output or changes to the composition of products produced cannot explain trends in pollution emissions from U.S. manufacturing between 1990 and 2008. Instead, changes in emissions over this time period were almost exclusively driven by decreased pollution per unit output for narrowly defined products. Second, the model-driven measure of the pollution tax rate that rationalizes observed pollution emissions and abatement decisions – a scalar measure of the stringency of environmental regulation – more than doubled between 1990 and 2008. We find similar patterns in regulation across all the main pollutants the Clean Air Act regulates ("criteria pollutants"). Third, we find that the increasing stringency of environmental regulation explains 75 percent or more of the 1990-to-2008 decrease in pollution emissions from U.S. manufacturing. Despite the plant-level relationship between pollution and productivity documented in Figure 2 and similar relationships found in related literature (Bloom, Genakos, Martin, and Sadun, 2010; Holladay, 2011; Martin, 2011), we find that changes in U.S. productivity have had small effects on U.S. pollution emissions at the economy-wide level.

The paper weaves together elements of workhorse models from the international trade and environment economics literatures. In our model, consumers have constant elasticity of substitution preferences across varieties of goods. Each firm draws a unique productivity level, which leads firms to differ in their pollution abatement investments and, ultimately, in their pollution emissions. As in Melitz (2003), entrepreneurs draw productivities from a Pareto distribution and may pay a fixed cost to produce goods and a separate fixed cost to export goods. As in Copeland and Taylor (2003), operating firms face pollution taxes and allocate a share of productive factors to abating pollution. The model accounts for endogenous changes in firm entry, exit, production, and export decisions in a tractable way that yields analytical solutions that allow us to analyze counterfactuals. The model fits a variety of key stylized facts—dirty industries invest more in pollution abatement than clean industries; productive plants emit less pollution per unit output than unproductive plants; productive plants produce more than unproductive plants; and only the most productive plants export.

We further assess the explanatory power of the model by comparing closed form expressions for regulatory stringency with actual changes in environmental regulations over this time period. We use difference-in-differences-in-differences to relate our model-driven estimates of the stringency of NO_x regulation to the EPA's implementation of the NO_x Budget Training Program, a cap-and-trade policy for firms in 19 regulated states, which was rolled out in 2003-2004. Our regression estimates show a close link between the structural measure of pollution taxes and the actual implementation of the NO_x cap-and-trade program. We also measure implicit tax rates for CO_2, a pollutant which largely has not been regulated. While our inferred tax rates for most pollutants more than doubled between 1990 and 2008, the implicit tax rate for CO_2 hardly changed over this time period.

To take this model to the data, we rely upon detailed plant-level records from the United States Census Bureau that we link to plant-level data from the U.S. Environmental Protection Agency (EPA). These data report pollution emissions, value of shipments, and production costs for each plant in the U.S. manufacturing sector, allowing us to estimate the key parameters of our model. We use these data to estimate the elasticity of pollution emissions with respect to pollution abatement, an elasticity that plays a central role in many models of pollution abatement (Copeland and Taylor, 2003; Levinson and Taylor, 2008; Forslid, Okubo, and Ultveit-Moe, 2011; Kreickemeier and Richter, 2014). In our model, this parameter has equivalent interpretations as the Cobb-Douglas cost share of pollution taxes in production or, alternatively, as the elasticity of pollution emissions with respect to productivity. Despite being widely incorporated into theory, to the best of our knowledge, this parameter has never been estimated empirically. We also estimate elasticities of substitution and the parameters governing the distribution of productivity separately by industry; these parameters play central roles in trade research and determine how changes to regulation and trade costs affect productivity and pollution emissions.

The model imposes strong functional form assumptions that only approximate reality. To appreciate the potential insight such strong assumptions can provide, consider an example of how more direct approaches could give misleading results. Plant-level evidence, including Figure 2, shows that more productive or better-managed firms emit less pollution per unit of output. This stylized fact motivates our assumptions about firm heterogeneity. However, the idea that improving productivity is likely to decrease pollution, which holds true in plant-level regressions, may be misleading at the economy-wide level. Suppose that a one percent increase in a firm's productivity causes a one percent decrease in pollution-per-unit-of-output. However, increasing national productivity also increases national income. At the economy-wide level, productivity growth may diminish factor demand per unit of output, but idle factors can then be used for additional output in the same or other plants. Since consumers may spend the additional income on pollution-emitting goods, this productivity shock may decrease pollution intensity but increase total output. In this simple example, partial equilibrium analysis may suggest that productivity growth benefits the environment, but in general equilibrium, productivity growth has ambiguous environmental effects. Similarly, while some research finds that opening markets to international trade reallocates market share to more productive firms, it would be premature from just these plant-level regressions to conclude that opening markets to trade decreases aggregate pollution emissions.

This paper builds on three strands of literature. First, research at the intersection of trade and the environment has used statistical decompositions and reduced-form regressions to explain changes in a country's pollution emissions (Koo, 1974; Grossman and Krueger, 1995; Antweiler, Copeland, and Taylor, 2001; Gamper-Rabindran, 2006; Ederington, Levinson, and Minier, 2008; Levinson, 2009). Most of this work relates pollution to economic objects that are endogenous outcomes of the global economy, such as trade flows or GDP. We extend this work by developing a new data resource that contains plant-product-year output for the U.S. manufacturing industry between 1990-2008. We use these data to examine whether the scale and/or composition of U.S. manufacturing output over this time period can explain the observed reductions in manufacturing emissions. In contrast with the existing literature, we also develop a quantitative model that allows us to decompose how fundamental shocks such as trade costs, regulation, productivity, and preferences affect pollution emissions through the lens of counterfactuals. A benefit of this approach is that we can investigate a wide variety of counterfactuals because the model specifies and quantifies underlying firm and consumer decisions.

A second body of research measures the stringency and consequences of air pollution regulation. Many papers focus on a single regulation (Greenstone, 2002; Becker, 2005; Walker, 2011, 2013; Deschenes, Greenstone, and Shapiro, 2013). The existing focus on a single regulation is due to parsimony and because researchers can form credible empirical comparison groups for a subset of environmental policies. However, the EPA, state, and local regulators have implemented dozens of overlapping air pollution regulations over the last 20 years, many of which have not been analyzed with policy evaluation tools. We believe this is the first paper to calculate or estimate the change in the overall regulatory burden, or shadow price of pollution, that manufacturing firms face due to local and national air pollution regulations. For example, while much research has explored the Clean Air Act's (CAA) county-level nonattainment designations, large polluters in counties that meet the CAA standards are still regulated, albeit with weaker stringency than in nonattainment counties. Berman and Bui (2001a,b) describe the entire menu of local air quality regulations facing manufacturing firms around Los Angeles, finding 11 local air quality regulations for petroleum refining and 46 for manufacturing (a count which excludes state and federal regulations).[2]

Third, this paper contributes to a recent literature applying firm-level, microfounded models of international trade to policy questions (Eaton and Kortum, 2002; Dekle, Eaton, and Kortum, 2008; Donaldson, Forthcoming). Like Shapiro (2013), we use these models to analyze environmental regulation. Unlike existing literature, we analyze a setting of monopolistic competition where profit-maximizing, heterogeneous firms endogenously choose investments in pollution abatement. Our model builds on several papers in the trade literature (Copeland and Taylor, 2003; Melitz, 2003; Chaney, 2008; Hsieh and Ossa, 2011; Eaton, Kortum, Neiman, and Romalis, 2011; Arkolakis, Costinot, and Rodriguez-Clare, 2012). We know of few general equilibrium decompositions like we provide, though Eaton, Kortum, Neiman, and Romalis (2011) and Burstein, Morales, and Vogel (2013) provide related decompositions for trade and for wage inequality, respectively.

The rest of the paper proceeds as follows. Section 2 presents a statistical decomposition in order to break down aggregate emissions trends in our data, while also highlighting the frontier of what we are able to say with the data alone. Section 3 outlines our trade-environment model. Section 4 discusses the data, and Section 5 discusses how we estimate the parameters. Section 6 presents the main results, and Section 7 discusses alternative explanations and additional robustness concerns. Section 8 concludes.

2 A Statistical Decomposition of U.S. Emissions 1990-2008

Much economic research interprets national changes in industrial air pollution via three pathways (Copeland and Taylor, 1994; Grossman and Krueger, 1995). One is a change in the scale of real output. The second

[2]Most of the manufacturing policies apply to only a few industries each. The analysis includes the years 1979 to 1993. Los Angeles has among the most stringent air quality regulations in the country. We thank Eli Berman and Linda Bui for sharing details of these regulations.

is a change in the composition of production from relatively clean products like "household furniture" to relatively dirty products like "carbon black." The third is a change in the production technique used to produce a single product, which could decrease a product's pollution emissions per unit of output.

We begin by presenting a statistical decomposition of manufacturing pollution emissions using newly developed administrative data on manufacturing plant-product-year production from 1990 to 2008. The Census of Manufacturers and the Annual Survey of Manufacturers collect sub-industry, product-level output data, at the plant-product-year level. We use this information to illustrate whether changes in the total scale of output and/or changes in the composition of products produced is able to explain the observed reductions in ambient air emissions. Our focus on products rather than industries is unique to the literature and is meant to capture the fact that even within a fairly narrow industry code (e.g., 4-digit Standard Industrial Classification (SIC) code), many products differ significantly in their emissions intensities.[3] For example, while the U.S. manufacturing sector contains 455 4-digit SIC codes, the product trailer from the Census and Annual Survey of Manufacturers allows us to perform this decomposition using 1,440 products. This granularity allows us to quantify by how much the scale of output versus the types of products produced can explain the observed reductions in manufacturing air emissions.

Consider the following representation of total manufacturing pollution, denoted Z:

$$Z = \sum_s z_s = \sum_s x_s e_s = X \sum_s \kappa_s e_s \tag{1}$$

Here total manufacturing pollution Z equals the sum of pollution from each manufacturing product s, z_s. A manufacturing product in our setting can be thought of as a sub-industry classification, where for example, SIC 3312 (blast furnaces and steel mills) is subdivided into 24 different products ranging from steel wire (33125) to cold rolled sheets and strip (excluding metallic coated and electrical) (33127).[4] Alternatively, we can write manufacturing pollution as equal to the total output of a product x_s multiplied by a product-specific emissions factor e_s. We can also represent manufacturing pollution emissions as the total output shipped by all manufacturing industries, X, multiplied by the sum of each product's share of total output, $\kappa_s \equiv x_s/X$, times an emissions coefficient reflecting pollution per dollar of output shipped of that product ($e_s \equiv z_s/x_s$). In vector notation, we have

$$Z = X\kappa'\mathbf{e}$$

where κ and $\mathbf{e}$ are $S \times 1$ vectors containing the market shares of each of the S products and their pollution intensities, respectively. Totally differentiating yields three terms representing the scale, composition, and technique effects:

$$dZ \quad = \quad \underbrace{\kappa'\mathbf{e}dX}_{\text{scale}} + \underbrace{X\mathbf{e}'d\kappa}_{\text{composition}} + \underbrace{X\kappa'd\mathbf{e}}_{\text{technique}} \tag{2}$$

[3]Previous research has explored trends in manufacturing pollution emissions using industry level data. The previous literature has acknowledged that a limitation of industry level production data the difficulty in distinguishing changes in the reallocation of production towards cleaner products from industry-level "technique" based reductions in emissions intensity (Koo, 1974; Gamper-Rabindran, 2006; Ederington, Levinson, and Minier, 2008; Levinson, 2009).

[4]Output at the five-digit SIC level is the most disaggregate data available for all plants in the Census and Annual Survey of Manufacturers.

Taking the decomposition in equation (2) to the data requires annual data on total pollution, total output, each product's contribution to output, and each product's emissions intensity. Pollution and total output come from the EPA's National Emissions Inventory (NEI) and the NBER-CES Manufacturing Industry Database, respectively. We construct product-level output shares in each year using the product trailer from the Census and Annual Survey of Manufacturers. In order to construct product-level emissions factors we match the National Emissions Inventory to the Annual Survey of Manufacturers in 1990 via name and address string matching. Appendix C.2 describes the string matching process in more detail. Since the NEI reports emissions at the plant-level and the Census product trailer reports output at the product level, we apportion plant-level emissions to a plant-product-year using product shares produced at the plant in 1990.[5,6] We take the total emissions attributable to each product in 1990 and divide by the total product shipments in 1990 to construct emissions intensities.[7] We then use these 1990 product-level emissions intensities to project the scale and composition effects forward in time, holding technology (i.e. our emissions intensities) constant at 1990 emissions rates. The decomposition allows us to observe what emissions would have looked like in 2005 if firms still produced products with 1990 emissions intensities. Appendix C.3 provides additional details about both the underlying data and the decomposition.

Figure 3 illustrates the resulting statistical decomposition for nitrogen oxide emissions (NO_x). Appendix Figure A1 presents results from other pollutants, which show similar patterns. The top solid line in Figure 3 depicts the total real value of manufacturing shipments, where each industry's output is deflated by the NBER-CES industry specific price index and then totaled. We scale total output so it equals 100 in 1990. This line summarizes what emissions would have been if we had kept the same emissions rates and the same product composition as in 1990. The middle dashed line plots NO_x emissions that would have occurred if emissions intensities had remained fixed at 1990 levels but the composition of output across manufacturing products had equaled observed, historical values. The bottom dotted line plots actual NO_x emissions from manufacturing, as reported by the NEI. The bottom line implicitly summarizes the joint result of changing the scale, composition, and technique of manufacturing production over this time period.

The statistical decomposition leads to several conclusions. First, the dotted line shows that actual NO_x emissions fell by almost 50 percent. Second, the proximity of the solid and dashed lines shows that the composition between clean and dirty manufacturing products has not changed much over time, although between 2000 and 2007, manufacturing shifted to slightly cleaner products. Third, the solid and dashed lines each show that if the pollution intensity of industries had not changed, NO_x emissions would have risen by about 30 percent. Finally, the gap between the solid line on top and dotted line at bottom shows that changes in the pollution intensity of individual products (i.e. "technique) explains why NO_x emissions fell by 50 percent rather than rising by 30 percent.

Despite the relatively clear conclusion that most reductions in emissions are driven by within-product

[5] We have also created emissions intensities by simply dividing total plant emissions by product shipments for each product produced at a plant. Conclusions are not sensitive to either method of constructing product-level emissions intensities.

[6] Previous research has used the World Bank's Industrial Pollution Projection System (IPPS) for emissions intensities. The IPPS data provides a list of emissions intensities by four-digit Standard Industrial Classification (SIC) codes (Hettige, Martin, Singh, Wheeler, and Mundial, 1995; Levinson, 2009). Levinson (2014) constructs industry-level emissions intensities using the NBER-CES productivity database combined with raw NEI data.

[7] We deflated total product output by industry-year specific price indices, from the NBER-CES database, where 2008=1.

changes in emissions intensity, the data are relatively silent on what might be causing these changes. The rest of the paper investigates, in more detail, the underlying causes of these patterns in the data. If more productive plants emit less pollution per unit output, then product-level productivity growth could explain these patterns. Alternatively, changes in trade costs like the introduction of NAFTA or China's WTO ascension may have caused a reallocation of production away from unproductive and dirty firms toward more productive and perhaps cleaner firms that produce the same product. Lastly, increases in environmental regulatory stringency may also explain these reductions.

The quantitative model, which fills the remainder of this paper, makes different and arguably stronger assumptions than this statistical decomposition. The advantage of these stronger assumptions is an ability to explain how trade, regulation, productivity, and preferences contribute to the environmental improvements documented in Figures 1 and 3. The disadvantage is that these assumptions provide only a rough approximation of reality. In explaining the model, we discuss ways in which these assumptions are restrictive or could be relaxed in future research. The reader interested in additional detail on specific assumptions of this model is referred to Costinot and Rodriguez-Clare (2014), Melitz and Redding (2014), and to Copeland and Taylor (2003).

3 Model of Heterogeneous Firms with Endogenous Pollution Abatement

We describe a model of firm entry, production, trade, and pollution abatement, which is designed to reflect a stylized description of polluting industries. In the model, firms differ in their productivity level, which leads these firms to differ in their pollution abatement investments and ultimately pollution emissions. The model accounts for endogenous changes in firm entry, exit, production, and export decisions in a tractable way that yields analytical solutions that allow us to analyze counterfactuals. The model fits a variety of key stylized facts—dirty industries invest more in pollution abatement than clean industries; productive plants emit less pollution per unit output than unproductive plants; productive plants produce more than unproductive plants; and only the most productive plants export. Like all models, this approach seeks to reflect systematic patterns across firms while recognizing that some strict assumptions which enhance tractability like monopolistic competition and constant elasticity of substitution (CES) utility are not literally accurate descriptions of firms and consumers.

The model has a straightforward economic environment. We analyze a world of two countries (US and Foreign), each with a representative agent. Each country has a single productive factor, which we call labor, and which is inelastically supplied. We present the model's main results here and intermediate derivations in Appendix B. This section explains the model's four assumptions, shows how we use the model to analyze counterfactuals, and then discusses how we obtain empirical counterparts to the key economic objects in the model.

3.1 Model Assumptions

A1. Preferences. The representative agent in destination country d has the following utility function:

$$U_d = \prod_s \left(\left[\sum_o \int_{\omega \in \Omega_{o,s}} q_{od,s}(\omega)^{\frac{\sigma_s - 1}{\sigma_s}} d\omega \right]^{\frac{\sigma_s}{\sigma_s - 1}} \right)^{\beta_{d,s}} Z_d^{-\delta} \tag{3}$$

Equation (3) describes CES utility across product varieties within a sector, Cobb-Douglas preferences across sectors, and multiplicative damages from pollution. The representative agent allocates expenditure across varieties of goods ω from the measure $\Omega_{o,s}$ of goods produced by industry s in origin county o. The parameter $\beta_{d,s}$ represents the share of country d's expenditure devoted to industry s, where $\sum_s \beta_{d,s} = 1$. We refer to $\beta_{d,s}$ as consumer preferences. The variable $q_{od,s}(\omega)$ represents the quantity of variety ω goods in industry s which are exported from origin country o to destination country d. The industry-specific parameter $\sigma_s > 1$ represents the elasticity of substitution across varieties. The parameter δ governs the disutility of pollution Z_d. We assume that pollution is a pure externality, so the representative agent ignores the term $Z_d^{-\delta}$ in making expenditure choices.[8]

The assumption of CES utility, which is common in trade and macroeconomic research, implies that consumers experience decreasing marginal utility from consuming a given variety and increasing utility in the total measure of varieties. We assume this utility function because it provides a simple way to account for different varieties within a sector while leading to parsimonious aggregate descriptions of production and trade flows across countries and industries.[9]

A2. Firms and Market Structure.

A competitive fringe of entrepreneurs may choose to pay the sunk entry cost $f^e_{o,s}$ to draw a productivity φ. The productivity is drawn from a Pareto distribution with shape parameter θ_s and location parameter $b_{o,s}$. We make the standard assumption that $\theta_s > \sigma_s - 1$ so the entrants have finite expected profits.[10] After observing the productivity draw, an entrepreneur who decides to produce must pay a separate fixed cost. Firms engage in monopolistic competition so that conditional on choosing to operate, an entrepreneur chooses prices $p_{od,s}$ and abatement investments ξ to maximize profits:

$$\pi_{o,s}(\varphi) = \sum_d \pi_{od,s}(\varphi) - w_o f^e_{o,s}, \tag{4}$$

$$\text{where} \qquad \pi_{od,s}(\varphi) = p_{od,s}(\varphi) q_{od,s}(\varphi) - w_o l_{od,s}(\varphi) \tau_{od,s} - t_{o,s} z_{od,s}(\varphi) \tau_{od,s} - w_d f_{od,s}$$

[8]Because the empirical analysis of the paper does not use information on the marginal damages of air pollution, the choice of functional form for pollution damages does not affect the empirical section of the paper. The empirical results would be numerically equivalent if pollution damages were assumed to be multiplicative, additive, or any other form.

[9]Research finds that non-CES utility functions, such as the linear demand system, translog utility, and certain generalizations which can allow for endogenous markups, are all part of the same "gravity family of models. While this implies that their measures of the gains from trade are closely related,, these structures do not always obtain the kind of tractable closed-form relationships we use here (Melitz and Ottaviano, 2008; Feenstra and Weinstein, 2010; Arkolakis, Costinot, Donaldson, and Rodriguez-Clare, 2012).

[10]In this setting, the strict criterion for entrants to have finite expected profits is actually slightly weaker: $\theta_s > (\sigma_s - 1)(1 - \alpha)$, where α is described below but remains between 0 and 1.

The firm sells $q_{od,s} = (1 - \xi) l_{od,s} \varphi$ units. The profit function involves several terms. For simplicity, we drop the variety notation ω and index a firm by its productivity φ. A consumer in destination d pays price $p_{od,s}(\varphi)$ for goods from firm φ. Each firm receives revenue $p_{od,s}(\varphi)\, q_{od,s}(\varphi)$ and requires $l_{od,s}(\varphi)$ units of productive labor at wage w_o to produce goods for sending to destination d. A fraction of this labor $1 - \xi$ is used to produce output, and the remaining fraction ξ to abating pollution. Each firm pays the pollution tax $t_{o,s}$ per ton of pollution emitted on $z_{od,s}(\varphi)$ tons of pollution emitted for producing goods shipped to destination d. Firms face iceberg trade costs, so $\tau_{od,s} \geq 1$ units must be shipped for one unit to arrive (hence, the firm must produce $\tau_{od,s} q_{od,s}$ in order to sell $q_{od,s}$). A firm that chooses to enter the destination market d must pay the fixed cost $f_{od,s}$.[11] Domestic trade costs are normalized so $\tau_{oo,s} = f_{oo,s} = 1$. Pollution tax revenues are lost to rent-seeking. Here $b_{o,s}$ describes a country's productivity while θ_s describes the dispersion of productivity draws within an industry s.

We assume this market structure for several reasons. Many dirty industries like cement and steel are concentrated and have barriers to entry. By accounting for fixed entry costs and industry-specific markups, our assumptions reflect a stylized version of polluting industries. At the same time, this formulation lets us account for firm entry and exit, and for reallocation of productive factors and output across firms. Finally, the Pareto technology distribution has plausible theoretical microfoundations (Gabaix, 1999; Luttmer, 2007) and provides a good fit to the empirical firm distribution, at least in the upper tail (Axtell, 2001; Eaton, Kortum, and Kramarz, 2011). In addition, the Pareto assumption yields closed form solutions for the productivity cutoffs and other endogenous variables of the model.

A3. Pollution. Firms produce pollution emissions with the following technology:

$$z_{od,s} = (1 - \xi)^{1/\alpha_s}\, \varphi l_{od,s} \tag{5}$$

We assume pollution regulations are stringent enough that all firms engage in some abatement. Equation (5) states that pollution is an increasing function of output and a decreasing function of pollution abatement expenditures. In addition, this formulation also implies that more productive firms emit less pollution per unit output produced. Equation (5) is essentially the pollution production technology adopted in Copeland and Taylor (2003). Modeling emissions in this way is appealing because many sensible and seemingly different ways of modeling pollution turn out to be equivalent to equation (5). As we show later, α represents the elasticity of pollution emissions intensity with respect to pollution abatement intensity. Pollution emissions intensity is measured as units of pollution emitted per unit of output, and pollution abatement intensity is measured as abatement expenditures divided by total factor costs. We will also show that pollution emissions in this model can be described as another factor of production in a Cobb-Douglas production technology, and in this interpretation α is the Cobb-Douglas share for pollution emissions. Copeland and Taylor (2003) discuss other equivalent interpretations of this model of abatement.

It may be useful to explain conceptually how firm and consumer decisions determine pollution and

[11] The assumption that this fixed cost is paid in destination- and not origin-country labor simplifies counterfactual analysis. Related work makes the same assumption that fixed market entry costs are paid in destination labor (Hsieh and Ossa, 2011; Eaton, Kortum, and Kramarz, 2011). The equivalence between welfare calculations in some perfect and monopolistic competition models in Arkolakis, Costinot, and Rodriguez-Clare (2012) requires fixed exporting costs to be paid in destination country labor.

abatement in this model. Equation (5) shows that for an operating firm, pollution emissions decline when the firm reallocates productive factors to abatement investments. Those abatement investments affect profits because they decrease required pollution tax payments but increase abatement costs. The endogenous decisions in the model like firm entry, exit, production, and trade all respond to environmental regulations, and all of these forces interact to determine pollution emissions. A firm that is not operating produces no pollution; a firm that begins exporting increases its output and also its pollution emissions. So although equation (5) only literally describes pollution abatement as a decision for operating firms, the model more broadly accounts for a variety of ways in which firm and consumer behavior affects pollution emissions.

A4. Competitive Equilibrium.

Consumers maximize utility; firms maximize profits; and in each country, labor supply equals labor demand:

$$L_o = L_o^e + L_o^m + L_o^p \tag{6}$$

A country's labor supply L_o is allocated to three activities: paying the fixed entry cost $f_{o,s}^e$ for drawing a productivity (L_o^e); paying country-specific entry costs $f_{od,s}$, which can be interpreted as marketing costs (L_o^m); and engaging in production, including pollution abatement (L_o^p).

Utility maximization implies that the share $\lambda_{od,s}$ of country d's expenditure on goods from industry s in country o has the following "gravity" structure:[12]

$$\lambda_{od,s} = \frac{M_{o,s}^e \left(\frac{w_o}{b_{o,s}}\right)^{-\theta_s} (\tau_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} (f_{od,s})^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} (t_{o,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}}{\sum_i M_{i,s}^e \left(\frac{w_i}{b_{i,s}}\right)^{-\theta_s} (\tau_{id,s})^{-\frac{\theta_s}{1-\alpha_s}} (f_{id,s})^{1-\frac{\theta_s}{(1-\alpha_s)(\sigma_s-1)}} (t_{i,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}} \tag{7}$$

where $M_{o,s}$ represents the mass of entrepreneurs who attempt entry into production in country o and sector s. Ignoring some terms which cancel, the numerator of equation (7) represents the value of imports and the denominator represents total expenditure, where the denominator is the sum over countries i of total imports from each country.

The gravity equation (7) is an intermediate result of the model, but it summarizes nicely how environmental regulation affects economic activity. It shows that a country with a high pollution tax $t_{o,s}$ will have lower exports and less production, particularly in dirty industries with high values of the pollution parameter α_s. An empirical literature measures how environmental regulation affects plant openings, employment, shipments, and other outcomes (see discussion of literature in Walker (2013)). A voluminous trade literature derives gravity equations expressing trade flows in a way that approximate actual bilateral trade patterns (see discussion of literature in Anderson and van Wincoop (2003)). Equation (7) provides an appealing link between these two literatures. In addition, equation (7) allows empirical tests of the hypothesis that industry reallocates production to countries with weaker environmental regulation (known as the "pollution havens hypothesis"; see Levinson and Taylor (2008)).

[12]The "gravity" description reflects the fact that bilateral trade in this and many other models is proportional to two countries' incomes (in an analogy to gravity in physics, their mass) and inversely proportional to their trade cost (analogously to physics, their distance). Most trade models with constant elasticity of substitution preferences and iceberg trade costs produce a gravity equation. This relationship also has empirical support—Leamer and Levinsohn (1995, p. 1384), for example, describe empirical estimates of gravity equations as "some of the clearest and most robust empirical findings in economics."

3.2 Comparative Statics

One key motivation for using a model of heterogeneous firms is the finding that more productive plants emit less pollution per unit of output. This relationship is not immediately apparent from the pollution technology assumption in equation (5). However, the first-order conditions of the firm's problem show that more productive firms invest a greater share of productive factors in pollution abatement (see equation (23) in Appendix B). Substituting that first-order condition into the firm's choice of pollution emissions gives the following firm-level relationship:

$$\frac{z_{od,s}(\varphi)}{q_{od,s}(\varphi)} = \frac{1}{\varphi^{1-\alpha_s}} \left(\frac{w_o}{t_{o,s}} \frac{\alpha_s}{1-\alpha_s} \right)^{1-\alpha_s} \tag{8}$$

Because $\alpha_s \in (0,1)$ can be interpreted as the cost share of pollution taxes, this expression implies that more productive firms (greater φ) emit less pollution $z_{od,s}$ per unit of output $q_{od,s}$. It also implies that a regression of $log(z_{od,s}/q_{od,s})$ on $log(\varphi)$ will give a coefficient of $\alpha_s - 1$. Since α is likely to be a small number, it is notable that the slopes of the trendlines depicted in Figure 2 are around minus one.

3.3 Equilibrium in Changes

By combining assumptions A1 through A4 in specific ways, we can use this model to analyze counterfactuals. We first combine the model's assumptions into two equilibrium conditions that summarize firm behavior. One equilibrium condition is the labor market clearing assumption (6). This condition says that any counterfactual must have total labor demanded equal total labor supplied in each country.

The other equilibrium condition combines two facts. First, the expected profit from operating a firm must equal the fixed cost of forming a firm. Second, a cutoff productivity for each destination market determines whether the expected profit from entering that market exceeds the fixed marketing cost required to sell there (i.e., a free entry condition and zero cutoff profit). Equation (25) of Appendix B formally describes the second equilibrium condition. These are equilibrium conditions because if a set of data satisfies these conditions, then those data represent a competitive equilibrium in the sense of Assumption A4. We use these conditions to analyze how counterfactuals affect welfare.

The two equilibrium conditions include numerous variables that are difficult to measure. Rather than attempt to measure these variables, we rewrite each variable as a proportional change from a base year, as in Dekle, Eaton, and Kortum (2008). The benefit of writing the model in changes is that we do not need data on difficult-to-observe variables because many of them do not appear in changes. For example, constructing an empirical analogue to the model's equilibrium conditions in levels (equation (6) and equation (25)) would require data on productivity levels for each country and industry, the fixed costs of entering each market, and the measure of firms operating in each country and industry. But in changes these variables will not appear in the equilibrium conditions, and so we do not need to measure them.

Our formal use of this technique to analyze the model proceeds as follows. Let x denote some variable from the model, let x' denote the value under a counterfactual scenario, and let $\hat{x} \equiv x'/x$ denote the proportional change in x due to the counterfactual. Written in changes, the two equilibrium equations (6) and (25) become the following:

$$1 = \psi_o \left(\frac{\sum_s \hat{M}^e_{o,s} R_{o,s} \frac{(\sigma_s - 1)(\theta_s - \alpha_s + 1)}{\sigma_s \theta_s} + \eta'_o}{\sum_s R_{o,s} \frac{(\sigma_s - 1)(\theta_s - \alpha_s + 1)}{\sigma_s \theta_s} + \eta_o} \right) \tag{9}$$

$$\hat{w}_o = \sum_d \frac{\zeta_{od,s} \left(\frac{\hat{w}_o}{\hat{b}_{o,s}} \right)^{-\theta_s} (\hat{\tau}_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} \left(\hat{f}_{od,s} \right)^{1 - \frac{\theta_s}{(\sigma_s - 1)(1-\alpha_s)}} \left(\hat{t}_{o,s} \right)^{-\frac{\alpha_s \theta_s}{1-\alpha_s}}}{\sum_i \lambda_{id,s} \hat{M}^e_{i,s} \left(\frac{\hat{w}_o}{\hat{b}_{o,s}} \right)^{-\theta_s} (\hat{\tau}_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} \left(\hat{f}_{od,s} \right)^{1 - \frac{\theta_s}{(\sigma_s - 1)(1-\alpha_s)}} \left(\hat{t}_{o,s} \right)^{-\frac{\alpha_s \theta_s}{1-\alpha_s}}} \hat{\beta}_{d,s} \hat{w}_d \tag{10}$$

These two equations will play central roles in the rest of the paper and we will reference them repeatedly. For each counterfactual, we use these equilibrium conditions to solve for the values of wages and firm entry decisions that characterize that counterfactual. To express these equilibrium conditions succinctly, we have defined export shares $\zeta_{od,s} \equiv X_{od,s} / \sum_d X_{od,s}$ and the parameter combinations η_o and ψ_o.[13] Intuitively, equation (9) says that the change in labor demand from each industry must equal the change in a country's labor supply. We assume labor supply is fixed ($\hat{L}_o = 1$) but relax this in sensitivity analysis. Equation (10) summarizes two ideas: given a set of primitive economic attributes (i.e. changes to variable and fixed trade costs, regulation, productivity, and preferences), changes to wages and firm entry must be such that entrepreneurs earn zero expected profit from drawing a productivity; and firms with positive expected profits from exporting choose to export.

In order to measure pollution emissions associated with the counterfactual, we integrate pollution emissions from (5). The change in country o's pollution emissions between a baseline year and counterfactual is then given by

$$\hat{Z}_o = \frac{\sum_s \frac{\hat{M}^e_{o,s}}{\hat{w}_o \hat{t}_{o,s}} Z_{o,s}}{\sum_s Z_{o,s}}. \tag{11}$$

An industry's pollution emissions increase proportionally with firm entry $\hat{M}^e_{o,s}$ and decrease with regulation $\hat{t}_{o,s}$ and wages $\hat{w}_o$. Equation (11) says that the proportional change in pollution emissions is the sum over industries of pollution in a counterfactual scenario, all divided by observed baseline pollution.

3.4 Taking the Model to the Data

A goal of this paper is to construct empirical analogues to key terms in the model and then to use the resulting framework to analyze counterfactuals. To clarify how we pursue this goal, it may help to distinguish between three classes of variables in the equilibrium conditions (9) and (10): data and parameters; shocks; and endogenous variables. Data and parameters represent quantities that we assume are fixed in counterfactuals. The data in this model include baseline country expenditure shares and baseline export shares ($\lambda_{od,s}$ and $\zeta_{od,s}$). The parameters in the model consist of the elasticity of substitution across product varieties, the shape parameter of the Pareto distribution of firm productivities, and the pollution elasticity (σ_s, θ_s, and α_s, respectively). Each parameter differs by industry s. We estimate elasticities of substitution and shape

[13] Specifically, $\eta_{o,s} \equiv \sum_s \left[-\frac{\theta_s - (\sigma_s - 1)(1-\alpha_s) - \sigma_s \theta_s}{\sigma_s \theta_s} \beta_{o,s} NX_o - NX_{o,s} \frac{(\sigma_s - 1)(\theta_s - \alpha_s + 1)}{\sigma_s \theta_s} \right]$ and $\psi_o \equiv \frac{1 - \sum_s \frac{\theta_s - (\sigma_s - 1)(1-\alpha_s)}{\sigma_s \theta_s} \beta_{o,s}}{1 - \sum_s \frac{\theta_s - (\sigma_s - 1)(1-\alpha_s)}{\sigma_s \theta_s} \beta'_{o,s}}$

parameters using methods similar to those used by Hsieh and Ossa (2011). As far as we know, there are no existing empirical estimates of the pollution elasticity.

Shocks represent changes to the global economy's primitive attributes. The shocks in this model include changes to variable and fixed trade costs, regulation, productivity, and preferences ($\hat{\tau}_{od,s}$ and $\hat{f}_{od,s}$; $\hat{t}_{o,s}$; $\hat{b}_{o,s}$; and $\hat{\beta}_{o,s}$). We measure these shocks but in most cases do not explain why they occurred. For example, we will quantify where and when technology changed but will not investigate whether this was due to innovation, technology transfer, or other reasons.

Endogenous variables represent values that are determined by the equilibrium interaction between supply and demand in each counterfactual, as described in equations (9) and (10), so as to achieve a competitive equilibrium. The endogenous variables in this model include changes in nominal wages in each country and changes to the mass of entrepreneurs who attempt entry into production in each country and sector ($\hat{w}_o$ and $\hat{M}_{o,s}$).

We can also describe these groups of variables in more general terms that may further clarify their meaning. The parameters describe the dispersion of productivities across firms, the willingness of consumers to substitute between different varieties of goods, and the relationship between firm-level productivity and firm-level pollution intensity. The shocks describe actual annual changes between 1990 and 2008 in the following: the stringency of environmental regulation, overall U.S. competitiveness; overall foreign competitiveness; and the share of consumer expenditure on each sector (a measure of preferences). The endogenous variables represent wages and prices.

We now turn to putting this model to work—we first explain the data we compile to analyze the model, and then we present our results. This model has many moving parts, so it may be useful to summarize it informally as follows. We use the model to consider questions like the following. Suppose some of these shocks had taken on different values than what we actually observed for the years 1990-2008, what would the paths of NO_x and CO_2 emissions have looked like? Changing historic values of those shocks affects many endogenous decisions: firm entry, exit, production, trade, and abatement. Given outcomes of all those endogenous decisions, we obtain model-driven estimates of how changing values of those shocks affect pollution emissions. Ultimately, we can quantify the importance of each set of shocks (trade, regulation, productivity, preferences) in contributing to declines in U.S. manufacturing pollution emissions.

4 Data

The data for this paper fall into three categories: plant-level microdata for estimating the model's parameters; country-by-industry aggregates used to analyze counterfactuals; and ancillary data for sensitivity analysis.

We use plant-level microdata to estimate three parameters of the model, calculated separately for each industry: the elasticity of substitution across product varieties; the shape parameter of the Pareto distribution of firm productivities; and a pollution elasticity. Estimating the elasticity of substitution requires input costs and the value of total sales for each industry. We obtain these data from the U.S. Census Bureau's Annual Survey of Manufactures (ASM) in the first year of our sample, 1990.[14] The ASM is a probabilistic

[14] We use the year 1990 to estimate the pollution elasticity because 1990 and 2005 are the only years with both emissions and

sample of approximately 60,000 establishments per year.[15] All our calculations with the ASM use sampling weights provided by the Census Bureau so the calculations are representative of the industry as a whole. We also use the ASM data to estimate the Pareto shape parameter, the details of which are described below.

Estimating the pollution elasticity requires two additional pieces of information: pollution abatement expenditures and pollution emissions. Pollution abatement expenditures come from the Pollution Abatement Costs and Expenditures (PACE) survey, which was developed jointly by the U.S. Environmental Protection Agency and the U.S. Census Bureau.[16] We also use data on air pollution emissions from the U.S. Environmental Protection Agency's National Emissions Inventory (NEI), which provides a comprehensive and detailed report of air pollution emissions from all sources. The NEI was created to provide EPA, federal and state decision makers, the U.S. public, and foreign countries with accurate measures of U.S. pollution emissions.[17]

We compile aggregate data for the U.S. and foreign countries separately for each industry and for each of the years 1990-2008. In particular, we need production and trade data from each country, and we need a measure of pollution emissions in the United States. For production, we use data from the Structural Analysis Database of the Organization for Economic Co-operation and Development (OECD). For trade, we use data from the OECD's Structural Analysis Database. Both datasets are reported in two-digit International Standard Industrial Classification codes, third revision. We convert trade data, which are reported in foreign currencies, to nominal U.S. dollars using annual exchange rates from the OECD Statistics dataset.[18] We aggregate these data to two countries (the U.S. and Foreign) and to 17 manufacturing industries defined in Appendix Table A1. We abstract from non-manufacturing activity. Although almost no countries report intranational trade (goods produced in the same country where they are consumed), we measure it as total production minus total exports.

We measure U.S. pollution emissions with the same National Emissions Inventory (NEI) data used to

abatement data, and using the same year 1990 to estimate all parameters may enhance comparability.

[15]Between 1990 and 1996, firms with at least 250 employees or $500 million in sales were sampled with certainty. Beginning in 1998, firms with at least 500 employees or $1 billion in sales were sampled with certainty. Below these thresholds, the probability of appearing in the sample increases with a firm's size.

[16]Empirical research has used the PACE survey to show that pollution abatement expenditures respond to Clean Air Act county-specific regulations; other work has shown that PACE expenditures are correlated with state-specific foreign direct investment (Keller and Levinson, 2002; Becker, 2005). The 1990 and 2005 PACE data that we use have similar structure and are broadly comparable. The 1999 PACE data, which we do not use, was not comparable with these surveys (Becker and Shadbegian, 2005).

[17]Our measures of particulate matter pollution in the NEI include only filterable particulate matter. This category includes particulates that can be captured on a filter during sampling. It excludes condensible particulate matter, which are gaseous particles that condense to small particles after they cool. It also excludes "secondary" particulate matter, which is formed in the atmosphere through reactions involving other gases like NO_x and SO_2. Filterable particulate matter is the only type of particulate matter reported in all NEI years 1990-2008. Several recent studies have used NEI microdata to explore temporal and spatial patterns in emissions trends and to incorporate air pollution into national accounts (Levinson, 2009; Muller and Mendelsohn, 2009; Muller, Mendelsohn, and Nordhaus, 2011). The classification of "polluting" industries in other studies (Greenstone, 2002; Greenstone, List, and Syverson, 2012) relies on an EPA study that used pollution emissions data from the AIRS dataset, which were later integrated into the NEI.

[18]The OECD only reports production data in these years for 32 countries. To extrapolate to other countries, we take the ratio of GDP to gross output in these countries, which is 0.310. Dekle, Eaton, and Kortum (2008) report an extremely similar value for this ratio of 0.312. We impute gross output for other countries by dividing their GDP, observed in the World Development Indicators data, by 0.310. We allocate this gross output to sectors according to the share of exports in non-U.S. countries from each sector in the OECD data.

measure pollution parameters. The NEI is not conducted annually, and we use all available years: 1990, 1996, 1999, 2002, 2005, 2008.[19] The year 1993 had no inventory. We focus on industry-level emissions of six of the main air pollutants regulated under the Clean Air Act: carbon monoxide (CO), nitrogen oxides (NO_x), particulate matter less than 10 micrometers (PM_{10}), particulate matter less than 2.5 micrometers ($PM_{2.5}$), sulfur dioxide (SO_2), and volatile organic compounds (VOCs).

We use several sources of additional data for the sensitivity analysis and robustness exercises in Section 7. These datasets are described in more detail in the context of the specific applications in Section 7.

5 Estimation and Results: Parameters and Shocks

We first describe estimation of pollution parameters, then trade/macro parameters, then recovery of historic shocks.

5.1 Pollution Parameters

Although the model literally describes pollution as a second output which is taxed, algebra combining Assumptions A2 and A3 of the model shows that pollution intensity can be written as a function of abatement investments:

$$\frac{z}{q} = (1-\xi)^{(1-\alpha)/\alpha} \tag{12}$$

Equation (12) says that emissions intensity z/q is a decreasing function of pollution abatement, and this function decreases more rapidly when the pollution elasticity α is larger. As before, ξ represents the share of factors used for pollution abatement rather than for production. Taking logs of equation (12), taking first differences over time, and allowing for national trends η_t in emissions intensity and idiosyncratic disturbances $\epsilon_{i,t}$ to pollution intensity gives

$$\Delta \ln\left(\frac{z_{i,t}}{q_{i,t}}\right) = \frac{1-\alpha}{\alpha}\Delta \ln\left(1-\xi_{i,t}\right) + \eta_t + \epsilon_{i,t} \tag{13}$$

where Δ represents the first-difference operator. Since ξ is the abatement cost share, the less firms spend on abatement, the larger z/q will be. So we expect $(1-\alpha)/\alpha$ to be positive.

Equation (13) is designed to measure the effect of pollution abatement on pollution emissions intensity. We estimate equation (13) using data from the U.S. manufacturing sector, relying on a balanced county×industry panel for the years 1990 and 2005.[20] We use data on pollution emissions from NEI, the value of shipments and value of production costs from ASM, and pollution abatement costs from PACE (z, q, and ξ, respectively).[21] We "winsorize" the reported emissions data at the 99th percentile of the 4-digit NAICS-year emissions distribution, and we use sample weights from both the Annual Survey of Manufacturers and the Pollution Abatement Costs and Expenditure Survey to inflate survey values to be nationally

[19] The final year of the analysis is 2008 because while NEI data are available for the year 2011, production data are not.

[20] The years 1990 and 2005 are the only years of data for which the PACE survey overlaps with the NEI, which is necessary to be able to observe all the parameters in equation (13).

[21] We proxy for measures of physical output $q_{i,t}$ using plant revenue, deflated by industry-specific output price deflators, where the industry-specific output price deflators come from the NBER-CES database.

representative. Total abatement costs consist of the sum of abatement operating costs plus the rental cost of capital associated with the observed abatement capital at a plant.[22] Total expenditures consist of the sum of expenditures on salary and wages, materials, energy, and the industry-specific capital rental rates for a given level of capital stock.

There are several reasons to be concerned that changes in pollution abatement costs are endogenous in the regression model, leading to biased estimates of α. For example, if regulators require the dirtiest plants to spend large shares of their costs on pollution abatement, then reverse causality will create downward bias in estimates of $(1-\alpha)/\alpha$. Moreover, our measures of abatement costs and total factor costs are based on survey responses from PACE and the ASM, both of which may contain measurement error in reported expenditures.

To address possible endogeneity concerns, we instrument for changes in the abatement cost share $\ln(1-\xi_{i,t})$ in equation (13) using changes in local environmental regulatory stringency. The EPA sets National Ambient Air Quality Standards which describe the minimum air quality needed to protect human health. The EPA requires polluting firms in areas that exceed the EPA's air quality standards ("nonattainment" counties) to install pollution abatement technologies.[23] Existing research has found that changes to county nonattainment status increase pollution abatement expenditures for polluting firms (Becker, 2005). We revisit this work by examining how county-level nonattainment designation influences abatement expenditures and ultimately pollution intensity per unit of output.

We estimate a single α using this regression approach, and we use an additional implication from the model to scale this estimate for each sector.[24] We use the fact that α_s represents pollution tax payments as a share of production costs. Algebra using Assumptions A2 and A3 from the model shows that we can write total output as a Cobb-Douglas function of pollution emissions and productive factors:

$$q = (z)^{\alpha_s} (l\varphi)^{1-\alpha_s}$$

Under Cobb-Douglas production with constant returns to scale, the output elasticity α_s is equal to the share of firm costs which represent pollution taxes. Since the U.S. does not have pollution taxes, we cannot directly observe the share of firm costs that represent pollution taxes. However, if the pollution tax rate is constant across industries, then the relative value of α_s across industries is proportional to the tons of pollution emitted per dollar of input costs in each industry. For example, if the basic metals industry emitted twice as much pollution per dollar of input costs as the textiles industry did, then we would have $\alpha_{\text{basic-metals}} = 2\alpha_{\text{textiles}}$. We use this approach to measure relative differences in α across industries. We then

[22]Capital rental rates are from unpublished data constructed by the Bureau of Labor Statistics for use in computing their Multifactor Productivity series. These data are commonly used in the productivity literature to proxy for industry-specific capital rental rates. See e.g., Syverson (2011). We only observe abatement capital stocks in 2005 (not 1990). We impute 1990 abatement capital stocks using our observed measure of depreciation expenditures in both 1990 and 2005. Specifically, we use the 2005 ratio of abatement capital stocks to abatement depreciation expenditures, and we multiply this ratio by the 1990 abatement depreciation expenditure for a plant to back out the 1990 abatement capital stock of the plant.

[23]For a much more detailed description of the Clean Air Act and the regulations associated with nonattainment designations, see Wooley and Morss (2013).

[24]In principle, we could use our instrumental variables regression approach to estimate α for each of the 17 sectors in our analysis. In practice, we have a limited sample of plants that we observe in all three of the NEI, PACE, and ASM datasets. When dividing these plants into 17 industries, the samples are too small to estimate equation (13) separately for each industry.

scale these values so the mean across all sectors equals the economy-wide elasticity of pollution emissions intensity with respect to abatement costs from our equation (13) regression estimate, $\widehat{\alpha}$.

Table 1 reports the first-stage, reduced-form, and instrumental variable regressions of equation (13) for the five pollutants in the NEI for which we have an instrumental variable for abatement expenditures. We analyze each pollutant in a separate regression, where county-level nonattainment designations imposed under the Clean Air Act serve as instrumental variables for the abatement cost shares in Panel C.[25] Columns (1) through (5) analyze each pollutant separately, and Column (6) uses total emissions of all pollutants in tons as a summary measure of emissions. All regressions report cluster-robust standard errors in parentheses, clustering separately by county and 3-digit NAICS industry.

Panel A of Table 1 presents the "first-stage" regressions which show that designating a county as nonattainment increases the proportion of firm costs devoted to pollution abatement in industries that account for a larger share of pollutant p emissions. All of these first-stage regressions have negative signs, implying that regulated firms increase the share of costs devoted to pollution abatement by 6 percent relative to the baseline share. Panel B presents evidence from the "reduced form" regressions of pollution emissions intensity on the regulation instrument.[26] The regression estimates show that polluting industries in newly regulated counties decrease their pollution per unit of output after the regulations go into place. The relationship between nonattainment and pollution emission rates is negative for all pollutants, imprecise for most pollutants, but precise for VOC emissions and for total pollution emissions. Panel C, which presents our instrumental variable regression estimates, shows that changes in pollution abatement cost shares, instrumented with changes in Clean Air Act regulations, predict changes in pollution intensity. Two of five pollutants are individually statistically significant and the overall regression in column (6) is precise. Panel D presents our estimates of α that come from a nonlinear transformation of the regression coefficient $(1-\alpha)/\alpha$. Although the exact value differs slightly across pollutants, the estimates of α range from 0.008 to 0.017. When we aggregate over pollutants in column (6), we obtain the value $\alpha = 0.011$, which is statistically significant at the one-percent level.

This value of α applies to the manufacturing industry as a whole. As described above, we obtain values of α_s for each sector by using information on pollution emissions per dollar of input costs in each sector. We rescale emissions per dollar of input cost in each sector so that the overall average across all sectors is equal to 0.011 from Table 1. Appendix Table A2 reports these rescaled values of α, separately for each pollutant and industry. The resulting pollution elasticities range from 0.001 to 0.048. Perhaps unsurprisingly, the cleanest industries are Radio, Television, Communication; and Motor Vehicles, Trailers. The dirtiest industries are

[25] Technically, the instrumental variable we use for changes in abatement expenditures is an interaction between a variable indicating the pollution intensity of pollutant p of an industry in 1990 (i.e. $\text{PolluterShare}_{jp} = \frac{\text{IndustryEmissions}_{jp}}{\text{TotalEmissions}_p}$) and whether the county switches into nonattainment for any pollutant between 1990 and 1993 (i.e. $1[\text{Nonattain}_c]=1$). Thus, the instrumental variable is $\text{PolluterShare}_{jp} \times 1[\text{Nonattain}_c]$. We allow for a county to be in nonattainment if it violates the EPA standards for any of the pollutants regulated under the Clean Air Act. In practice, nonattainment is pollutant specific. We model nonattainment in this way in order to capture cross-pollutant regulatory spillovers and to ameliorate the fact that many pollutants have little/no variation over this time period (e.g. CO). The focus on counties that switched into nonattainment between 1990 and 1993 is meant to capture all the counties that became newly regulated under the 1990 Clean Air Act Amendments. We include the lower order interaction terms in all regression models to facilitate identification of the difference-in-differences interaction term.

[26] The dependent variable in Panels B and C is log((pollution+1)/output) in order to prevent attrition for non-polluting county×industry×year cells in the sample.

Basic Metals, and Other Non-Metallic Minerals.

Are these estimates of α reasonable? The pollution parameter α represents the share of costs which firms devote to paying pollution taxes. Thus, the overall estimate of 0.011 implies that firms are behaving as if they pay one percent of their total production costs to pollution taxes. We lack a method to test this number independently, but we can compare it to two related statistics. First, the PACE data report that manufacturing pollution abatement costs are about half a percent of total manufacturing sales (U.S. Census Bureau, 2008). Second, Greenstone, List, and Syverson (2012) find that nonattainment designations decrease the total factor productivity of regulated firms by 2.6 percent. These figures are not directly comparable, but because they all characterize the economic costs of environmental regulation, it is notable that are of the same order of magnitude.

5.2 Macroeconomic Parameters

We also estimate the elasticity of substitution and shape parameter of the Pareto distribution separately for each sector. While the pollution elasticity describes how productivity covaries with pollution, the Pareto shape parameter describes the width of the dispersion of productivity. Because the Pareto shape parameter appears throughout both the second equilibrium condition and the gravity equation (equations (10) and (26)), it also determines how shocks to trade costs and regulation affect trade flows, firm entry, and wages. The elasticity of substitution plays a less important role—it affects how changes to fixed trade costs and changes to preferences affect expenditure patterns (equation (26)).

We estimate these parameters by building on the approach used in Hsieh and Ossa (2011). To estimate the elasticity of substitution across product varieties, we use the implication of the model that an industry's expenditure on labor for production is proportional to the industry's revenue:

$$w_o L^p_{o,s} = (1-\alpha_s)\,\frac{\sigma_s - 1}{\sigma_s} X_{o,s} \tag{14}$$

Here $L^p_{o,s}$ represents labor used in production and $X_{o,s}$ represents revenue.[27] We use the 1990 Annual Survey of Manufactures to calculate these elasticities separately for each of the 17 aggregated ISIC sectors.

Column 4 of Table 2 presents our estimates of σ_s for each sector.[28] The elasticity of substitution ranges from 2.89 to 8.18 across industries, with a cross-industry mean of 4.76. We expect a smaller elasticity of substitution for industries with more differentiated products. The pattern across sectors generally follows this pattern. The largest elasticity of 8.18 in absolute value is for the Coke, Refined Petroleum, and Nuclear Fuels sector, which has fairly homogeneous products. The smallest elasticity of 2.89 is for the Medical, Precision, and Optical Products sector, which has fairly differentiated products.

Next, we estimate the shape parameter of the Pareto distribution of firm productivities. We rely on the

[27] In the model, this prediction reflects only wage payments used for production. In applying this prediction empirically, we measure all factor payments in the data (not merely wages), and we treat all factor payments in the data as productive (since the data do not separately measure fixed entry and marketing costs). Firm revenues are "inventory-adjusted" total value of shipments for a plant in 1990, and firm costs consist of expenditures on labor, parts and materials, energy, and capital.

[28] The reported elasticity is calculated as $\sigma_s = (1-\alpha_s)/((1-\alpha_s) - wL_s/X_s)$, where α_s is the pollution elasticity estimated above and wL_s/X_s is factor costs divided by the value of shipments. Columns 1-3 of Table 2 present these intermediate inputs into the construction of σ_s.

fact that if the distribution of firm productivities is Pareto with shape parameter θ_s, then the distribution of firm sales is Pareto with shape parameter $\theta_s/(\sigma_s - 1)$. The Pareto tail cumulative distribution function is $\Pr\{x > X_{i,s}\} = (b_{i,s}/X_{i,s})^{\theta_s/(\sigma_s-1)}$ for $X_{i,s} \geq b_{i,s}$. Taking logs gives

$$\ln(Pr\{x > X_{i,s}\}) = \gamma_{0,s} + \gamma_{1,s} \ln(X_{i,s}) + \epsilon_{i,s} \tag{15}$$

We estimate equation (15) separately for each industry s, and the coefficient $\gamma_{1,s}$ in each regression is generally close to negative one. The Pareto shape parameter is then given by $\theta_s = \gamma_{1,s}(1 - \sigma_s)$.

We use a subset of the firm-level data to estimate equation (15). Because selection into exporting can bias these estimates (di Giovanni, Levchenko, and Ranciere, 2011), we estimate this regression using only domestic sales. Additionally, since the Pareto distribution best fits the right tail of the firm distribution, we estimate these regressions using firms above the 90th percentile of sales within each industry.[29]

Columns 5 and 6 of Table 2 present our estimates of the Pareto shape parameter θ_s for each industry. For each row of the table, we estimate the Pareto shape parameters θ_s by regressing the log of a firm's sales rank on the log of its sales using the microdata from the 1990 Annual Survey of Manufactures. The regression estimates of the Pareto shape parameter are extremely precise, which reflects the fact that power law distributions describe firm size well, at least in the upper tail (Gabaix, 2009). The shape parameter estimates are close to the elasticity of substitution estimates for the corresponding industry, and in all cases cohere with our assumption that $\theta_s > \sigma_s - 1$. This relationship stems from the fact that regression estimates of equation (15) obtain coefficients $\gamma_{i,s}$ near minus one, and the Pareto shape parameter is calculated as $\theta_s = \gamma_{1,s}(1 - \sigma_s)$. The economy-wide mean of 5.71 is similar to estimates of the trade elasticity of 4 to 10 (Arkolakis, Costinot, and Rodriguez-Clare, 2012), with more recent estimates closer to 4 (Simonvska and Waugh, 2014).

5.3 Recovering Historic Values of Shocks

The previous section explained how we use plant-level microdata to estimate the model's key parameters. Given these parameters, this section explains how we use country×industry aggregate data to recover historic values of the paper's four shocks: foreign competitiveness (a measure of trade); environmental regulation; domestic competitiveness (a measure of productivity); and consumer preferences.

Why do we need to measure historic values of these shocks? The model and estimated parameters are all we need to analyze counterfactuals. But this paper's research question of why pollution followed its historical path requires us to look at a specific counterfactual where some shocks take on their historical values and other shocks do not. Analyzing that kind of counterfactual requires measuring the actual, historic values of each shock for each year in 1990-2008.

[29] In the census microdata, we measure domestic sales as inventory-adjusted total value of shipments minus the value of export shipments. Estimating the regression using only the upper tail of firm sizes follows the literature by taking a set of firms for which the relationship between firm rank and size is approximately linear (Gabaix, 2009; di Giovanni, Levchenko, and Ranciere, 2011). To determine the percentile cutoff for these regressions, we bin the data into values of firm size that are equidistant from each other on the log scale, then collapse the rank/size data to the bin level for 10 bins. We examine the scatter plot of these points overlaid by the linear fit to these points. In general, the upper 90th percentile of the sales distribution is strongly linear with respect to firm rank.

Historic values of shocks to trade costs, regulation, and preferences could in principle be measured from data on tariffs and shipping costs, announcements of new environmental regulation, total factor productivity, and consumer surveys. But while such data may be proxies that are correlated with the shocks described in the model, they do not actually measure the shocks in the model. For example, the pollution tax variable $t_{o,s}$ in the model is meant to encompass the full set of air pollution regulations facing firms, whereas changes in regulation like Clean Air Act nonattainment designations or the implementation of cap-and-trade programs represent only partial changes in environmental regulation.

Hence, we use implications of the model to infer historic values of each shock from country×industry data on production, trade, and pollution. Our general approach is to use four equations describing a competitive equilibrium: gravity in changes (26), labor market clearing, (9), the second equilibrium condition (10), and the equation describing changes in pollution emissions (11). We then manipulate these equations to express each shock as a unique function of what we observe—data and parameters. For each shock described below, we use observed data (on the right-hand-side) to construct empirical analogues to the theoretical shocks (on the left-hand side). We let an asterisk (x^*) denote the historical value of shock x.

We define the shock to foreign competitiveness as

$$\hat{\Gamma}^*_{od,s} \equiv (1/\hat{b}_{o,s})^{-\theta_s} (\hat{\tau}_{od,s})^{-\theta_s/(1-\alpha_s)} (\hat{f}_{od,s})^{1-\theta_s/(\sigma_s-1)(1-\alpha_s)} \left(\hat{t}_{o,s}\right)^{-\alpha_s\theta_s/(1-\alpha_s)} \text{ for } o \neq U.S.$$

Foreign competitiveness combines foreign productivity, variable and fixed trade costs for foreign exports, and foreign environmental regulation ($\hat{b}_{o,s}$, $\hat{\tau}_{od,s}$, $\hat{f}_{od,s}$, and $\hat{t}_{o,s}$). We combine these variables into a single "foreign competitiveness" shock because these variables all decrease the ability of foreign firms to sell a wide variety of products to U.S. consumers at low prices, and because we lack the data to measure each component of this foreign shock separately.[30] By manipulating equation (26), we can write the shock to foreign competitiveness as

$$\hat{\Gamma}^*_{od,s} = \frac{\hat{\lambda}_{od,s}}{\hat{M}^e_{o,s} \left(\hat{w}_o\right)^{-\theta_s}} \left(\hat{P}_{d,s}\right)^{-\frac{\theta_s}{1-\alpha_s}} \left(\frac{\hat{\beta}_{d,s}}{\hat{w}_d} \frac{\hat{w}_d w_d L_d - \widehat{NX_d} NX_d}{w_d L_d - NX_d}\right)^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} \tag{16}$$

where $\hat{P}_{d,s}$, the change in country×sector price index, is defined in equation (27) of Appendix B. The right-hand-side of equation (16) shows that the change in foreign competitiveness equals the change in the share of U.S. expenditure on goods from a foreign country, divided by the change in nominal income times firm entry. The first term in parentheses summarizes changes to the destination price index, and the second term in parentheses accounts for the effects of changing preferences and changing trade imbalances NX_d on sector composition. We do not need to measure the terms in parentheses on the right-hand side of equation (16) because they are specific to destination d, and they cancel out in the gravity equation (7) and in the

[30] Separately measuring productivity and trade costs would require foreign producer price index data, which are not available for most countries, sectors, and years. We have tried using OECD STAT and Thompson Reuters Datastream to obtain these price data. Analyzing counterfactuals with those data suggests that the effects of productivity and trade costs separately are similar to the effects of the foreign shocks we describe here. However, in several cases it is computationally difficult to separately analyze trade costs and productivity in counterfactuals. Separately measuring the effect of foreign environmental regulation requires data on global air pollution emissions, which are not available.

second equilibrium condition (10), because they are in both the numerator and denominator of these two equations. These two equations are the only places this shock appears in analyzing counterfactuals.

We measure shocks to environmental regulation by integrating pollution emissions in equation (5) over the mass of operating firms, and then dividing the result by baseline pollution, giving

$$\hat{t}_{o,s} = \frac{\hat{w}_o \hat{M}^e_{o,s}}{\hat{Z}_{o,s}} \tag{17}$$

The change in environmental regulation equals the change in the mass of entering firms, divided by the change in pollution emissions, scaled by the change in factor prices.

We define shocks to U.S. competitiveness as

$$\hat{\Gamma}^*_{od,s} \equiv \left(1/\hat{b}_{o,s}\right)^{-\theta_s} \left(\hat{\tau}_{od,s}\right)^{-\theta_s/(1-\alpha_s)} (\hat{f}_{od,s})^{1-\theta_s/(\sigma_s-1)(1-\alpha_s)} \text{ for } o = U.S.$$

We measure this shock as

$$\hat{\Gamma}^*_{od,s} = \left(\hat{t}_{o,s}\right)^{\frac{\alpha_s\theta_s}{1-\alpha_s}} \frac{\hat{\lambda}_{od,s}}{\hat{M}^e_{o,s}\left(\hat{w}_o\right)^{-\theta_s}} \left(\hat{P}_{d,s}\right)^{-\frac{\theta_s}{1-\alpha_s}} \left(\frac{\hat{\beta}_{d,s}}{\hat{w}_d} \frac{\hat{w}_d w_d L_d - \widehat{NX}_d NX_d}{w_d L_d - NX_d}\right)^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} \tag{18}$$

We measure the change to U.S. competitiveness similarly to the change to foreign competitiveness. However, because we have pollution emissions data for the U.S. but not foreign countries, we can separate the environmental regulation term $\hat{t}_{o,s}$ from other components of U.S. competitiveness.

We measure shocks to consumer preferences as changes in the share of expenditure allocated to each sector:

$$\hat{\beta}^*_{d,s} = \frac{\sum_o X'_{od,s}/\sum_{o,s} X'_{od,s}}{\sum_o X_{od,s}/\sum_{o,s} X_{od,s}} \tag{19}$$

The numerator in this equation describes the share of a country's expenditure on sector s in some year, while the denominator describes the share of the country's expenditure on sector s in 1990.[31]

In all of these shocks, because labor is the only factor of production, the change in nominal wages equals $\hat{w}_o = \sum_{d,s} X'_{od,s}/\sum_{d,s} X_{od,s}$. Combining equations (7) and (10) shows that the change in the mass of firm entry in a sector equals the share of output growth from that sector: $\hat{M}^e_{o,s} = \left(\hat{w}_o^{-1}\right) \sum_d X'_{od,s}/\sum_d X_{od,s}$.

To reconstruct historic economic variables exactly, we also need to account for a fifth shock, namely shocks to trade imbalances. We need to incorporate this shock because one purpose of the decomposition is to recreate historic values of pollution emissions exactly when all shocks take their historic values, allowing us to investigate how pollution emissions change when some shocks do not take their historic values. In a dynamic model, trade imbalances would represent intertemporal concerns like saving or consumption smoothing, but in the comparative statics we examine here, trade imbalances appear as transfers from one country to another. In the full decomposition, we let changes in trade imbalances follow their historical path. We measure trade imbalances or net exports, NX_d, as a country's total exports minus its total imports.

[31] Our assumption of CES preferences implies that we abstract from consumer tastes changing among varieties within a sector. Our analysis of preferences only accounts for changes in the share of consumer expenditure going to each sector.

We emphasize that economy-wide variables like the mean productivity of operating firms or the value of trade flows depend on several of the shocks, including regulation. The productivity term $\hat{b}_{o,s}$ describes the location parameter of the Pareto distribution of productivities from which entrepreneurs draw. But the mean productivity of operating firms depends on which entrepreneurs choose to form a firm, and on the market share of unproductive versus productive firms. Ultimately, the mean productivity of operating firms depends not only on the shock to U.S. competitiveness but also on shocks to U.S. environmental regulation and to foreign competitiveness. Similarly, while shocks to trade frictions $\hat{\tau}_{od,s}$ and $\hat{f}_{od,s}$ are a focus of our discussion, the volume of trade flows in each industry ultimately depends on all the shocks.

Figure A2 shows the time path of the historical shocks in the paper.[32] Although we recover the value of each shock for each country×industry, it is cumbersome to describe values for 17 different industries. Instead, we plot shocks separately for "clean" and "dirty" industries. The six "dirty" industries represent the top third of α_s values, and the eleven "clean" industries represent the lowest two-thirds of α_s values.[33]

Figure A2a shows that foreign preferences for dirty versus clean goods changed relatively little over this period. In contrast, Figure A2b shows that U.S. preferences for products in dirty industries increased after 2004. This rapid increase in U.S. preferences for dirty goods may be driven by the increasing expenditure in the Coke, Refined Petroleum, and Nuclear Fuels sector, reflecting increases in global oil prices.[34] Finally, Figure A2c suggests that the stringency of environmental regulation grew rapidly over this period. The implied pollution tax for the U.S. manufacturing sector roughly doubled between 1990 and 2008. We initially discuss results for NO_x regulation, both since NO_x emissions are measured with higher-quality methods than most other pollutants are, and because we have detailed data on one major regulation, the NO_x Budget Trading Program. However, we also present results for other pollutants.[35]

Is this a realistic change in the stringency of environmental regulation? We emphasize that the U.S. does not actually have a pollution tax on NO_x. A way to think about the meaning of this tax is as follows: if all U.S. environmental regulation relevant to NO_x emissions from the manufacturing sector were replaced with a pollution tax, what change in that tax rate would lead to the changes in firm behavior that we actually observe? Given dramatic expansion of NO_x regulation over these 18 years, a doubling in the implicit tax on pollution seems plausible. An extremely incomplete list of actual changes in NO_x regulations includes:

[32] The model and counterfactuals account for competitiveness shocks to each country. As discussed earlier, although the price index $\hat{P}_{d,s}$ appears in our measure of competitiveness shocks, we dont need these price data to analyze counterfactuals. This is because destination price indices appear in both numerator and denominator of the second equilibrium condition and cancel. As a result, the historical shocks to U.S. and foreign competeviness outside of a particular counterfactual are not informative, and we omit competitiveness shocks from Figure A2.

[33] The dirty industries are: Wood Products; Paper and Publishing; Coke, Refined Petroleum, and Fuels; Chemicals; Other Non-metallic Minerals; and Basic Metals. The clean industries are: Food, Beverages, and Tobacco; Textiles, Apparel, Fur, and Leather; Rubber and Plastics; Fabricated Metals; Machinery and Equipment; Office, Computing, and Electrical; Radio, Television, and Communication; Medical, Precision, and Optical; Motor Vehicles and Trailers; Other Transport Equipment; and Furniture, Other, and Recycling.

[34] With Cobb-Douglas consumer utility, changes in the share of expenditure in a given sector translate to a change in preferences. This stylized fact that the share of U.S. expenditure on energy products nearly doubled between 2004 and 2008 appears in other data. For example, the Energy Information Agency Energy Information Administration (2011) records that consumer expenditure on all petroleum products grew in nominal terms from $470 billion in the year 2004 to $871 billion in the year 2008.

[35] According to the 2008 NEI, which reports monitoring method for almost all plants, just over half of manufacturing NO_x emissions are reported based on continuous emissions monitoring systems or other direct measures. We considered focusing on SO_2, but according to plant-level data we obtained from the EPA Clean Air Markets Division, the Acid Rain Program which created a cap-and-trade system for SO_2 in most years included only one or two manufacturing plants.

a nearly doubling of the number of counties in ozone nonattainment between 2003 and 2004, which may be the largest expansion of nonattainment areas since the Clean Air Act began;[36] the 1990 Clean Air Act Amendments, which required large NO_x emitters in ozone nonattainment areas to install stringent pollution controls by 1995;[37] the RECLAIM cap-and-trade for Los Angeles, which began in 1993; the Ozone Transport Commission cap-and-trade for New England, which began in 1999; and the NO_x Budget Trading Program for the Eastern U.S., which began in 2003.[38]

Our measures of these historic shocks depend on the changes in nominal wages in each country and changes in firm entry that occurred in each country and sector. Appendix Figure A3 plots these values. U.S. wages stagnated in the 1990s as U.S. output grew more slowly than global output did. U.S. wages grew slightly in the late 1990s and early 2000s, as U.S. output growth modestly outpaced global output growth. Finally, nominal wages declined in the 2000s as growth from foreign countries, especially China, accelerated. Foreign wages display the opposite pattern: growth in the early 1990s and late 2000s but a slight decrease in the intervening years.

Appendix Figure A3 also shows patterns in firm entry. Foreign firm entry expanded more quickly in dirty industries than in clean industries, as indicated by the slightly higher solid line in Figure A3c. U.S. firm entry, however, grew more rapidly in clean industries during the 1990s and then accelerated in dirty industries in the late 2000s. This increase in U.S. firm entry to dirty industries in the late 2000s reflects rising energy prices—greater value of output in dirty industries increases the expected profit from entry, attracting more firms to these industries.

6 Counterfactuals

We begin by discussing the mechanics of how we analyze counterfactuals. Then we use the model, parameter estimates, and recovered historic shocks to decompose how trade, environmental regulation, productivity, and preferences contributed to observed declines in pollution emissions.

6.1 Counterfactual Algorithm

To analyze counterfactuals, we use country×industry data from the year 1990 on production, trade, and U.S. pollution emissions ($X_{od,s}$ and $Z_{o,s}$), and the parameter vectors for each industry: the pollution elasticity, elasticity of substitution, and Pareto shape parameter (α_s, σ_s, and θ_s). With the full set of data and parameters, we then use the following algorithm to solve for a specific counterfactual:

1. Characterize the counterfactual scenario by choosing values for shocks to foreign and U.S. competitiveness, U.S. environmental regulation, and preferences in each of the years 1990-2008 $\{\hat{\Gamma}_{od,s}$, $\hat{t}_{o,s}$, and $\hat{\beta}_{o,s}\}$. These values can be hypothetical or they can describe the actual, historical values of these shocks.

[36] Ozone pollution is formed through photochemical reactions involving NO_x, VOCs, and heat and sunlight, so ozone nonattainment regulations target NO_x and VOC emissions.

[37] The EPA began requiring plants in ozone nonattainment areas to install Reasonably Available Control Technology (RACT).

[38] Some of these policies focus more on electricity generating units than on manufacturing. However, the relevant statistic here is the share of manufacturing pollution to which these policies applied.

2. Find the changes to wages and firm entry in each country×sector×year ($\hat{w}_o$ and $\hat{M}^e_{o,s}$) which make the equilibrium conditions (9) and (10) hold with equality for all countries and sectors and years by solving a system of nonlinear equations and then inputting the values chosen in step 1.[39] This system represents $N + NS - 1$ variables in $N + NS - 1$ unknowns: one unknown wage change per country, one unknown firm entry change per country×sector, and one unknown excluded as numeraire.

3. Use equation (11) to measure the change in U.S. pollution emissions, given the values from steps 1 and 2.

The historic values of shocks to foreign and domestic competitiveness, environmental regulation, and preferences are $\{\hat{\Gamma}^*_{od,s},\, \hat{t}^*_{o,s},\, \hat{\beta}^*_{o,s}\}$, calculated using equations (16) through (19). By construction, these values solve the two equilibrium conditions (9) and (10) in every country, industry, and year for the wage changes and firm entry changes ($\hat{w}^*_o$ and $\hat{M}^{e*}_{o,s}$) which actually occurred. Hence, if we take observed levels of trade, pollution emissions, and production from the initial year 1990, add the shocks $\{\hat{\Gamma}^*_{od,s},\, \hat{t}^*_{o,s},\, \hat{\beta}^*_{o,s}\}$ which actually occurred between 1990 and some future year, and calculate the new equilibrium, we recover the historic value of pollution from that year. However, we are interested in what pollution would have been if shocks had not equaled their historic values.

To decompose the change in pollution into the effects of the separate shocks, we study a specific set of counterfactuals. Consider the shock to foreign competitiveness. To measure how foreign competitiveness affected pollution, we define the shocks as follows:

$$\{\hat{\Gamma}_{od,s}, \hat{t}_{o,s}, \hat{\beta}_{o,s}\} = \begin{cases} \{\hat{\Gamma}^*_{od,s}, 1, 1\} & \text{if } o \neq U.S. \\ \{1, 1, 1\} & \text{if } o = U.S. \end{cases} \tag{20}$$

This says that the foreign competitiveness shock $\hat{\Gamma}_{od,s}$ took on its historic value $\hat{\Gamma}^*_{od,s}$, $o \neq U.S.$, but other shocks remained fixed at their 1990 values (i.e. the proportional change for every other shock equals one). Given the shocks defined in equation (20), we use steps 2 and 3 of the algorithm to recover the pollution emitted in this counterfactual. We do a similar calculation for each shock separately. For example, to measure the pollution change due to environmental regulation, we define the shocks as $\{\hat{\Gamma}_{od,s}, \hat{t}_{o,s}, \hat{\beta}_{o,s}\} = \{1, \hat{t}^*_{o,s}, 1\}$. We then follow steps 2 and 3 of the algorithm described above to measure the implied pollution under these shocks.

Three additional points may clarify this algorithm. First, setting all shocks equal to their historic values at once *exactly* recreates the historic decline in pollution. Second, although we are choosing the shocks to characterize a counterfactual, the firm-level decisions in the model — like entry, exit, abatement, production, and exports — are all adjusting freely in response to the shocks. Third, we analyze the model separately for each pollutant.

[39] To solve the system of nonlinear equations, we use a standard trust-region dogleg algorithm. However, as we discuss below and show in Appendix Table C2, other algorithms and randomly-chosen starting values give equivalent results.

6.2 Historic Decomposition

Figure 4 plots the time paths of NO_x emissions under five separate counterfactuals, indicated in the subfigure headings. In each subfigure, the solid line shows historic pollution emissions. The dashed line in each subfigure shows the model's counterfactual prediction of what would have happened if the indicated shock had followed its historic path and other shocks had remained fixed at their 1990 levels. For example, the dashed line in Figure 4a shows the pollution which the U.S. would have emitted in a counterfactual where foreign competitiveness followed its historic path but domestic competitiveness, pollution regulation, and preferences remained fixed at their 1990 levels. Each line is normalized to 100 in the year 1990. By construction, the solid line depicts the counterfactual resulting from setting the shocks equal to their historic values $\{\hat{\Gamma}^*_{od,s}, \hat{t}^*_{o,s}, \hat{\beta}^*_{o,s}\}$. The starred markers on the dashed lines show the years 1990, 1996, 1999, 2002, 2005, and 2008, when pollution emissions are observed in the NEI rather than linearly interpolated.

Figure 4a suggests that foreign competitiveness had a very limited effect on U.S. manufacturing NO_x emissions. Between 1990 and 2000, pollution in this counterfactual would have increased by a few percentage points. After 2000, when China's growth accelerated, foreign competitiveness led to modest decreases in U.S. pollution of a few percentage points. By 2008, in this counterfactual, U.S. pollution emissions were about five percent below their 1990 value. Ultimately, shocks to foreign competitiveness explain very little of the total decline in U.S. pollution.

Given the large effects of China's economic growth over this time on U.S. manufacturing employment (Autor, Dorn, and Hanson, 2013), or the scope of NAFTA for Mexico's access to U.S. markets (Cherniwchan, 2013), one might have expected foreign competitiveness to cause large decreases in pollution. Figure 4a suggests that this was not the case. Why didn't Chinese or other foreign competition substantially affect U.S. pollution over this time period? A few ideas help explain. Although China's exports are concentrated in low-skilled sectors, the graph of foreign firm entry in Appendix Figure A3c implies that they were not especially concentrated in dirty sectors. Moreover, aggregate data on U.S. manufacturing show that the effect of China's growth on manufacturing output or on value added was much smaller than its effect on employment (see, e.g., Pierce and Schott (2012)). Finally, one effect of foreign competition would be to shift U.S. production to cleaner or dirtier products goods, and the statistical decomposition from Section 2 provided no evidence that such a shift occurred.

Figure 4b also suggests that changes to U.S. competitiveness do not explain much of the change in U.S. manufacturing NO_x emissions. Between 1997 and 2003, the effect of U.S. competitiveness alone caused U.S. pollution emissions to increase by about 10 percent. By the end of the 1990-2008 period, though, the effect of this shock alone was nearly zero. How can we make sense of this finding that U.S. competitiveness explains little of the change in pollution? Figure 2 and plant-level regressions from papers discussed in the introduction suggest that more productive firms emit less pollution per unit of output. The model reflects this fact, since at the plant level, more productive firms in the model emit less pollution per unit of output. But at the economy-wide level, while productivity growth may diminish factor demand per unit of output, factors are used for other output in the same or other plants. Unless productivity growth is much larger in dirty industries, productivity growth may have limited scope to affect pollution. Our results here suggest that this differential was not enough to affect pollution substantially.

Figure 4c quantifies how changing consumer preferences affected pollution emissions, and it suggests they play little role in explaining the historic trends in pollution emissions. Between the years 1990 and 2000, preferences for clean goods decreased slightly, and this decreased U.S. pollution emissions. After 2000, by contrast, increasing expenditure on pollution-intensive goods leads to an increase in U.S. pollution emissions of roughly 20 percent.

The first three counterfactuals suggest that foreign competitiveness, domestic competitiveness, and U.S. preferences explain little of the decrease in pollution emissions. Figure 4d, by contrast, suggests that changes in environmental regulation over this time period account for most of the decrease in NO_x emissions. However, regulation by itself would have caused about 10 percent less pollution reduction than actually occurred. By the year 2008, regulation explains almost all of the change in pollution.

Figure 5 replicates these four counterfactual analyses for the other "criteria pollutants" regulated under the Clean Air Act: CO, $PM_{2.5}$, PM_{10}, SO_2, and VOCs. Each subfigure in Figure 5 shows all the counterfactuals for a single pollutant. In each subfigure, the solid line describes historic pollution emissions while the dashed lines describe counterfactual pollution emissions under individual sets of shocks. The exact numbers vary by pollutant, but a few empirical regularities are clear. For all pollutants, foreign competition had little effect in the 1990s but then decreased pollution in the 2000s, by up to 10 percent. This is consistent with the growth of China slightly decreasing U.S. manufacturing pollution emissions. Similarly, U.S. competition slightly increased pollution in the late 1990s and early 2000s, then slightly decreased it. Depending on the pollutant, changes in U.S. consumer preferences increased pollution in the 2000s by between 5 and 20 percent. Finally, while the exact number varies by year, for all five pollutants, U.S. environmental regulation explains most of the decrease in pollution.

The findings that regulation explains most of the observed changes in emissions across pollutants and that most pollutants had similar magnitude declines in emissions together imply that environmental regulation had similar changes over this time period for all the pollutants we study. It is difficult to assess this conclusion independently, but cursory reflection suggests it is plausible. Most pollutants experienced increased regulatory stringency over this time period. We have discussed in previous sections the many NO_x regulations that took place over this time period. The pollutants PM and VOC also experienced large expansions in Clean Air Act county nonattainment designations (and increasingly stringent nonattainment standards within these designations). The 1990 Clean Air Act Amendments also established new guidelines for CO that depend on the degree of local air quality violations. Areas in "moderate" or "serious" violation were required to implement programs introducing oxygenated fuels and/or enhanced emission inspection programs, among other measures. The 1990 Clean Air Act Amendments established the Acid Rain Program, which was designed to control SO_2 emissions over this time period. As mentioned above, there are additional air pollution programs at local, state, and regional levels, but the relative importance of these regulations compared to federal regulations is empirically unknown. In addition, firms and industries which emit large amounts of one pollutant often emit large amounts of other pollutants.[40] This suggests that some types of abatement for one pollutant may affect other pollutants at the plant.

[40] For example, SIC 3312 (Steel Works, Blast Furnaces, and Rolling Mills) is the largest manufacturing source of CO emissions, and it is also the third largest source of both PM_{10} and SO_2 manufacturing emissions.

7 Discussion and Sensitivity Analysis

The baseline analysis suggests the stringency of regulation for criteria air pollutants more than doubled between 1990 and 2008, and this change — rather than trade costs, productivity, or preferences — explains most of the observed decrease in pollution emissions. We now assess the robustness of this conclusion to a wide variety of sensitivity analyses: exploring the extent to which the model can predict pollution changes through means other than regulation; testing our inferred pollution taxes by comparing them to actual regulations; performing a similar decomposition for CO_2, which largely has not been regulated; measuring the importance of fuel mix in explaining the observed reductions in U.S. manufacturing emissions; and a handful of more technical changes to the model such as shutting off firm heterogeneity.[41] The last subsection discusses several other issues that are beyond the scope of this paper. Additional model sensitivity analysis can be found in Appendix C.4.

7.1 How Do Shocks Besides Regulation Affect Pollution in this Model?

This subsection explores a series of counterfactuals in order to understand whether the model could have led to other conclusions. We explore the extent to which forces other than environmental regulation can affect pollution in this model. We begin by exploring heterogeneity in counterfactuals for clean and dirty industries separately. In each of these counterfactuals, we take the baseline data in the year 1990, increase or decrease one shock, and calculate the changes in production and pollution abatement decisions which result. For brevity, we present this sensitivity analysis for NO_x emissions only. Appendix Figure A4 plots the resulting patterns of pollution for a suite of counterfactual scenarios.[42] The x-axis describes the hypothetical shock, and the y-axis records the resulting change in U.S. pollution emissions. The value 1 on the x-axis describes a counterfactual where a shock does not change, and the value 100 on the y-axis describes an outcome where U.S. pollution emissions do not change.

Figures A4a and A4b describe a range of shocks to U.S. and to foreign competitiveness, separately for dirty and clean industries. The solid line in Figure A4a shows that as foreign competitiveness in dirty industries grows, pollution emissions from U.S. manufacturing decline. This occurs because foreign competition in dirty industries abroad lowers expected profits in dirty U.S. industries, so production in dirty U.S. industries falls and U.S. productive factors shift to clean industries. At the same time, competition from dirty industries causes exit of unproductive (and dirty) U.S. firms, increasing the market share of cleaner firms within dirty U.S. industries.

[41] We also explored allowing for exogenous factor supply changes and including a role for non-manufacturing, Allowing for global factor supply increases provides a larger role for the scale effect to increase pollution and tends to make environmental regulation play an even larger role in explaining the observed decrease in pollution. Obtaining OECD data on trade and production for the non-manufacturing sector and incorporating this sector in the analysis of counterfactuals provides a modestly larger role for foreign competitiveness shocks. In counterfactuals allowing for a non-manufacturing sector, foreign competitiveness shocks explain an 11.7 percent decrease in pollution between 1990 and 2008, but in these counterfactuals environmental regulation still explains a large majority of the decrease in pollution.

[42] To create these graphs, we consider shocks ranging from 0.25 to 4.0 in increments of 0.25. For example, a shock of 0.25 in Figure A4a represents a counterfactual where foreign competitiveness in dirty industries falls to a fourth of its 1990 value but U.S. competitiveness, U.S. environmental regulation, and both U.S. and foreign preferences remain at their 1990 values. For each counterfactual, we measure the resulting change in pollution. We then plot these results for the entire range of shocks from 0.25 to 4.0.

Figure A4b shows that the opposite patterns happen with shocks to U.S. competitiveness. As U.S. competitiveness in dirty industries grows, U.S. pollution emissions increase; and as U.S. competitiveness in clean industries increases, U.S. pollution emissions fall. Finally, Figure A4c shows that hypothetical shocks to environmental regulation also have intuitive results. Holding other shocks constant, increasing the stringency of U.S. environmental regulation decreases pollution emissions.

More broadly, Figure A4 shows that U.S. pollution depends on all of the shocks we consider, and not only on environmental regulation. Ultimately, we find that changes to environmental regulation explain most of the recent declines in pollution. Figure A4 makes clear that because shocks to domestic and foreign competitiveness each can affect U.S. pollution, this model could have led to other conclusions.

7.2 Comparing the Implied Pollution Tax with Actual Regulatory Changes

This paper infers changes in pollution taxes by combining observed firm behavior with a model, rather than by using announced changes in environmental policy. To what extent does our model-driven measure of regulatory stringency correspond with known changes in environmental regulations? The paper has discussed the wide array of NO_x regulations which affected the manufacturing sector over the 1990-2008 period. In this section, we explore how our model-driven measure of pollution taxes respond to one well-known change in environmental regulation over this time period—the NO_x Budget Trading Program (NBP).

The NBP was a cap-and-trade program for NO_x emissions from power plants and large industrial plants in the Eastern U.S. that was rolled out in 2003 and 2004. The EPA distributed permits to each source and allowed free trading and banking of permits. Most sources in the NBP were electricity generating units, and most NO_x reductions came from coal-fired power plants. Nonetheless, many oil refineries, chemical plants, paper mills, and other manufacturing plants were regulated through the NBP. A range of economic research has studied the NBP, including studies that use differences-in-differences-in-differences research designs to measure the NBP's effects on pollution, health, and employment (Fowlie, 2010; Curtis, 2012; Deschenes, Greenstone, and Shapiro, 2013).

We obtain data from the EPA's Air Markets Program Data (AMPD) on facilities regulated under the NBP. The AMPD data do not have detailed industry codes in order to explore cross-industry heterogeneity in emissions responses. However, the EPA's NEI data do include this information in addition to facility name, longitude, latitude, county, and industry. We link the AMPD data to the NEI data by requiring an exact match on county and industry and a non-exact match on facility name, longitude, and latitude.[43] The link to the NEI data allows us to observe which industries participated in the NBP, and we use this information in a set of regression models described in more detail below.

Appendix Table A3 describes the share of manufacturing regulated in this cap-and-trade program. The NBP targeted large NO_x emitters. Although only a third of a percentage point of manufacturing plants in the NEI were regulated in the NBP, about 13 percent of manufacturing emissions of NO_x came from manufacturing plants that were subject to the NBP. The proportion of NO_x emissions from regulated plants

[43] The only measure of industry in the NBP data is a facility's "source category." We exclude NBP participants with cogeneration, electric utility, or small power producer as their primary source category, since these are not manufacturing plants.

ranges from 24 to 41 percent in a few industries: Paper and Publishing, Coke, Refined Petroleum, and Fuels; Rubber and Plastics; and Basic Metals.

Our empirical setup explores how our model-driven measure of pollution taxes corresponds with actual regulatory variation in a program-evaluation setting. We estimate the relationship between our measure of pollution taxes and the NBP rollout using the following difference-in-difference-in-differences regression model:

$$\ln(\hat{t}_{jst}) = \beta_1 \left(1[NBP_s] \times 1[NBPIndustry_j] \times 1[Year > 2002]\right) + \eta_{st} + \gamma_{jt} + \psi_{js} + \epsilon_{jst} \tag{21}$$

We regress our measure of implied pollution taxes, $\hat{t}$, as defined in equation (17), in industry j of NBP region s and year t, on a three-way interaction term describing the effect of being in an NBP-regulated industry in an NBP state in the years after the regulation went into place. We aggregate the data to the industry×region×year level, where a region is defined as inside/outside the NBP region, and industries are defined by the 17 manufacturing industries defined in Table A1.[44] We control for region×year fixed effects η_{st}, industry×year fixed effects γ_{jt}, and region×industry fixed effects ψ_{js}. With these sets of fixed effects, the model effectively controls for time-invariant observed or unobserved determinants of pollution taxes by industry×region, common transitory shocks to industries across regions, and transitory shocks within a region that affect all industries similarly. The identifying assumption of the model is that there exist no transitory shocks specific to regulated industries in the NBP region in the years after the NBP went into place. While this assumption is inherently untestable, the data permit some indirect tests. For example, data from years prior to the change in regulations permit the analysis of pre-trends across treatment and control groups prior to the change in policy. The coefficient of interest is β_1 which describes how the rollout of the NBP affected the evolution of pollution taxes in polluting industries of regulated states.

Appendix Table A4 presents results from several versions of equation (21). Each column represents a separate regression, and we report cluster-robust standard errors in parentheses, clustering at the industry×region level. The first column represents the baseline specification and suggests that polluting manufacturing firms in the NBP region in the years after the NBP went into place experienced a 1.195 log point increase in pollution taxes, or approximately 2.3× increase relative to the counterfactual.[45] Column (2) presents results that include industry×year fixed effects, and the results are nearly identical. Columns (3) and (4) add region×year fixed effects which slightly attenuate results, but the results remain statistically significant across all 4 specifications.

Equation (21) is somewhat restrictive in that it makes the assumption that the NBP led to a discontinuous level-shift in the regulations that occurs instantly and lasts forever. In order to investigate the transitional dynamics explicitly, we estimate an event study version of equation (21), including leads and lags in event

[44] States in the NBP region include Alabama, Connecticut, Delaware, Illinois, Indiana, Kentucky, Maryland, Massachusetts, Michigan, New Jersey, New York, North Carolina, Ohio, Pennsylvania, Rhode Island, South Carolina, Tennessee, Virginia, Washington DC, and West Virginia. All other states are defined as outside the NBP region.

[45] 2.3× is calculated as $\exp(1.195) - 1 = 2.303$.

time.[46] In the presence of industry×region fixed effects, not all event time coefficients are identified. We normalize the event-time coefficient in the year prior to the policy to 0. Figure 6 presents results. The figure suggests two main findings: First, in the years leading up to the policy, the implied taxes in the treatment and control groups are relatively similar and are not statistically different. This lends some reassurance that the research design is capturing a sharp event that affects NBP-regulated plants in the NBP region in the years after the NBP rather than some underlying trend in the data. Second, the years after the policy reveal a sharp and statistically significant increase in the implied pollution taxes for the NO_x polluting industries in the NBP region. The magnitudes of these estimates correspond closely to those from Appendix Table A4.

7.3 What Does the Model Say About Carbon Dioxide Emissions?

This paper focuses on six "criteria air pollutants" which are known to affect local air quality, since these pollutants have been a focus of U.S. environmental regulations, and because they have readily available data. The conclusions from the model are fairly similar across all the pollutants we analyzed. This section presents a type of placebo exercise by studying a pollutant that has not experienced the same set of environmental regulations over this time period, CO_2.

We analyze CO_2 emissions from manufacturing for two primary reasons. CO_2 emissions contribute to climate change, so it is important to understand the underlying forces driving changes in manufacturing CO_2 emissions. Additionally, local and federal policymakers have imposed little regulation on CO_2 emissions during this time period.[47] Therefore, using our model to analyze the extent to which environmental regulation has affected CO_2 emissions may provide a useful placebo test.

We use the Manufacturing Energy Consumption Survey (MECS) in order to calculate CO_2 emissions from manufacturing over this time period. MECS reports BTUs of eight categories of energy: coal; distillate fuel oil; residual fuel oil; electricity; liquefied petroleum gases and natural gas liquids; natural gas; coke and breeze; and other fuels. We use data at the level of 2-digit SIC code or a 3-digit NAICS code, concord them to the 17 industries codes used in this paper, and convert BTUs to tons of CO_2 emissions using standard rates from the U.S. EPA.[48] Appendix C.5 provides additional details on the MECS data used here. Finally, we treat CO_2 like the other pollutants in this study to investigate how CO_2 regulation has changed and how each of the shocks in this paper has contributed to changes in CO_2 emissions from manufacturing.

Figure 7a shows our inferred measure of the stringency of environmental regulation for each of the pollutants in this study, including CO_2. The dashed lines in that figure show that between 1990 and 2008, the stringency of regulation for most air pollutants increased by 100 to 250 percent. However, the solid line suggests the stringency of CO_2 regulation was essentially flat over this entire time period, and actually

[46] Specifically, we estimate models of the following form

$$\ln(\hat{t}_{jst}) = \sum_{\tau=1990}^{2008} \beta_k \left(1[NBP_s] \times 1[NBPRegulated_j] \times 1[Year = \tau]\right) + \eta_{st} + \gamma_{jt} + \psi_{js} + \epsilon_{jst}.$$

[47] The Northeast states began a cap-and-trade system for CO_2 emissions, the Regional Greenhouse Gas Initiative (RGGI), in 2008. Boulder, Colorado, and San Francisco implemented small carbon taxes in 2006 and 2008, respectively. States and the federal governments operate other taxes on various fuels which emit CO_2.

[48] CO_2 emissions rates are available at http://www.epa.gov/appdstar/pdf/brochure.pdf

decreased by about 4 percent. Although criteria air pollution regulation became much more stringent over this time period, CO_2 regulation hardly changed. The fact that we find such large increases in the implicit tax rate for air pollution emissions and small changes in the implicit tax rate for CO_2 emissions provides an additional piece of evidence that the model-driven measures of pollution taxes capture realistic features of the regulatory environment rather than changes in other associated economic variables.

Figure 7b shows the counterfactual decomposition for CO_2 emissions. The graph is the same as Figure 5, except that it shows results for CO_2 rather than for criteria air pollutants. The solid blue line shows that CO_2 emissions from manufacturing initially increased then decreased, but overall changed little relative to 1990. The dashed red lines show counterfactual CO_2 emissions under different sets of shocks. Overall, no one set of shocks completely explains the modest changes in CO_2 emissions, and regulation plays relatively little role in the observed CO_2 decreases in the 2000s.

7.4 Changing Fuel Inputs

Some manufacturing processes directly use fossil fuels to create heat and steam needed at various stages of production, and manufacturing plants may substitute from high-sulfur coal towards low-sulfur coal. Plants may also switch from coal to natural gas due to changes in relative prices across fuel types, or they can substitute towards non-fuel inputs due to rising global energy prices. This suggests that "fuel mix" may affect manufacturing pollution emissions, and this decision has not been explicit in the model and previous discussions.

We explore the importance of fuel mix in explaining the observed reductions in U.S. manufacturing emissions. We do not analyze whether environmental regulation or other forces like railroad deregulation, fracking, or others affect fuel mix. We merely quantify the importance of fuel mix and thereby bound the share of the change in pollution emissions which environmental regulation could explain. We address fuel mix here rather than in the paper's main model for two reasons. First, the detailed fuel quality data used here, unlike the data used for the main model, are not available separately for each manufacturing sub-industry. For this reason, this section treats manufacturing as a single industry. Second, it is intractable for the main model to address fuel mix directly, and the analysis here provides a simple way to assess its potential importance.

We begin by measuring what pollution emissions would have been if the manufacturing sector had used no abatement technologies, if all pollution had come from fuel combustion, and if fuel mix followed its historic path. We call this measure "potential pollution" and define it as:

$$\text{Potential Pollution}_y = \frac{\sum_f e_{fy} Q_{fy}}{H_y} \tag{22}$$

where e_{fy} are pollution emissions per physical unit of fuel f in year y in the absence of any abatement technology, Q_{fy} are the physical units of this fuel burned, and H_y is total heat input (measured in BTUs) of all fuel to the manufacturing sector. The emissions rate e_{fy} differs by year because it depends on the sulfur and ash content of each fuel, which change over time.

We construct empirical analogues to equation (22) using data from several sources. We measure e_{fy}, the air pollution emitted by combustion of different fuels in the absence of abatement technology, from EPA engineering estimates (Eastern Research Group, 2001); we measure Q_{fy} and H_y, the physical and BTU content of energy, from the U.S. Census Bureau's Manufacturing Energy Consumption Survey (MECS); and we measure the sulfur and ash content of petroleum and coal using EIA coal and petroleum reports. The data cover each year from 1993-2008. Additional details regarding the data construction are in Appendix C.

A few summary statistics highlight key trends in the underlying data. Appendix Figure A5a shows that the share of heat input from coal fell from 11.6 percent to 8.8 percent between 1993 and 2008. Much of this coal was replaced with steam purchases, which are included in the "other" category. The SO_2 emitted per BTU of coal declined by about ten percent over this period due to the use of lower-sulfur coal. However, Appendix Figure A5b suggests that the sulfur content of petroleum inputs hardly changed.

How did this changing fuel mix affect manufacturing pollution emissions? Appendix Figure A6 shows that if total BTUs of energy consumed for fuel had remained constant, fuel mix evolved as we observe, and nothing else had changed, pollution emissions from manufacturing would have declined slightly. Specifically, pollution emissions from manufacturing would have fallen by 11 percent for CO, 7 percent for NO_x, 8 percent for PM_{10}, 27 percent for SO_2, and 8 percent for VOCs.

Appendix Table A5 compares these statistics to the historic decreases in pollution emissions. On average across the five pollutants between 1993 and 1998, pollution emissions fell by 54 percent.[49] The change in fuel mix would predict a 12 percent decrease in pollution from fossil fuel combustion. To summarize, changing composition and quality of fuels can explain at most 23 percent of the decrease in observed pollution emissions. For SO_2, however, fuel mix is more important for explaining emissions reductions (52 percent).[50]

The paper's main analysis concludes that environmental regulation alone can account for most of the decrease in pollution emissions over the 1990-2008 period. Over the 1993-2008 time period, we find that changing the mix of fuels used in production can explain a fourth of the decrease in pollution emissions. We do not quantify the extent to which changes in fuel mix are due to environmental regulation versus other forces.[51] However, we take these results to conclude conservatively that environmental regulation can account for three-fourths or more of the decrease in pollution emissions from manufacturing.

[49]Recall that most of the paper analyzes the period 1990-2008, but this section begins in 1993, when fuel quality data become available, so this section is analyzing a smaller decline in pollution than the rest of the paper does.

[50]The engineering estimates used to measure air pollution emissions in the absence of abatement technology report different values for various methods of using a fuel. Technically, these various methods are called Source Classification Codes (SCC). The calculations in this subsection take an unweighted mean across SCCs within each industry. Because the 2002, 2005, and 2008 National Emissions Inventories report total pollution emissions separately for each SCC, we also tried doing this calculation with a weighted average across these SCCs, with weights equal to total tons of pollution emitted in NEI by each SCC. Finally, dividing an SCC's pollution emissions from NEI by its uncontrolled emissions rate from the engineering data provides one measure of how many physical units of each fuel are used for each SCC. So we also obtained results using a separate weighting system, where we weighted across SCCs by the physical tons of fuel used in each SCC. The unweighted mean across SCCs found that 23 percent of the observed decrease in pollution can be explained by fuel mix. The two alternative weighting schemes (weighting across SCCs by tons of pollution emitted or by tons of the fuel used in NEI) obtained very similar results of 25 and 26 percent.

[51]For example, increasing use of low-sulfur coal could be due to environmental regulation or to deregulation of railroads (Busse and Keohane, 2007).

7.5 Firm Heterogeneity

We now consider how a model with monopolistic competition but homogeneous firms, as in Krugman (1980), affects the decomposition. The paper develops a model built in part on Melitz (2003) where each firm draws a unique productivity from a Pareto distribution. That model is motivated by the plant-level stylized fact of Figure 2 whereby more productive firms emit less pollution per unit of output.

Mathematically, shutting off firm-level heterogeneity is relatively straightforward. Given the values of the Pareto shape parameter θ_s that we estimate, we shut off firm heterogeneity by setting $\sigma_s = \theta_s/(1-\alpha_s)$.[52] Equation (7) shows that this change removes the fixed entry cost term $f_{od,s}$ from the gravity equation entirely, and implies that all firms export.

Table C1 shows that shutting off firm heterogeneity has small effects on the decomposition results. In the main decomposition, environmental regulation alone accounts for a 47.96 percent decrease in pollution emissions. When we shut off firm heterogeneity, environmental regulation alone accounts for a nearly equivalent 47.95 percent decrease in pollution emissions. For the other shocks, shutting off firm heterogeneity has quantitatively larger effects on the decomposition, but the qualitative conclusions are unchanged.

Why does firm heterogeneity have such a small effect on our estimates, particularly given the strong relationship between plant-level pollution intensity and plant-level productivity documented in Figure 2? We emphasize two explanations. First, as discussed earlier, more productive firms may have lower pollution intensity at the plant level, but increasing plant-level productivity for a given level of output may free up productive factors which can be used in other factories to make widgets and pollution elsewhere. Second, an active body of research in trade has discussed the extent to which accounting for firm heterogeneity affects the magnitude of the gains from trade. In some settings, the effects of firm heterogeneity on the magnitude from gains from trade are modest or even zero.[53] Our setting of measuring changes in pollution emissions differs from this literature's focus on the gains from trade, but the intuition persists that adding more margins by which policy can affect welfare need not mean policy has larger effects on welfare.

7.6 Other Explanations

This paper builds a model that focuses on several key aspects of pollution emissions, but for simplicity and tractability it abstracts from others. This section informally discusses how other forces outside the model might affect its conclusions.

One abstraction of the model is industry detail. If firms have changed their focuses of production within one of our 17 sectors from more- to less-dirty industries and products, then our analysis may confound regulation with product substitution. The limit on our level of industry detail is the requirement to have

[52] Costinot and Rodriguez-Clare (2014) undertake a similar analysis of a heterogeneous firms model with no environmental components.

[53] Costinot and Rodriguez-Clare (2014) find that in a model with multiple sectors but no intermediate goods, a 40 percent global tariff would create a 1.2% global decrease in welfare (measured as the average across regions) in a world with monopolistic competition and heterogeneous firms and a 1.4% decrease in welfare in a world with monopolistic competition and homogeneous firms. More broadly, Arkolakis, Costinot, and Rodriguez-Clare (2012) show that the gains from trade are equivalent in these two frameworks for a model with one sector, no intermediate goods, and the same trade elasticity. Melitz and Redding (2014) argue that the gains from trade are strictly larger in a model with firm heterogeneity because given primitive parameters, the trade elasticity differs across models.

gross output data for foreign countries and the ability to calculate a counterfactual equilibrium. However, the statistical decomposition presented in Section 2 suggests that compositional changes in the type of goods produced within very narrow product categories are not able to explain a significant fraction of the observed emissions reductions. This evidence is indirect but suggests that industry detail may have limited power to explain changes in pollution emissions.

The model also assumes constant returns to scale in pollution abatement. A model with increasing returns to scale in abatement would have different structure (see e.g., Forslid, Okubo, and Ultveit-Moe (2011)). We considered the implications from such a model but chose not to pursue it for two reasons. First, the importance of fixed costs for abatement technologies is empirically unknown. Scale economies could be positive for capital investments like scrubbers, zero for fuel-switching like low-sulfur coal, and negative due to principal-agent issues for management innovations. Second, prices in such a model depend directly on market size, and market size appears in the equilibrium conditions in ways that make it difficult to apply the methodology we use.

The model is also silent on improvements in abatement technology, such as learning-by-doing. The value of the pollution technology α could change merely because regulation becomes more stringent. For example, increasing the market size of scrubbers creates additional incentives for research and development of more effective scrubbers and could decrease the future price of scrubbers. We conjecture that this sort of mechanism would make our model understate the effects of regulation on pollution emissions. Research on health care has highlighted the importance of induced innovation for pharmaceuticals (Finkelstein, 2004), but empirical research on induced innovation stemming from environmental regulation is in its infancy.[54] We leave this topic for future work.

We leave comparisons of monopolistic competition with other models (both single-factor perfect competition models like Eaton and Kortum (2002) and multiple-factor Heckscher-Ohlin models) for future work. However, we note that the gravity equation in this paper has very similar structure to the gravity equation in models of perfect competition. We also note that the trade and productivity shocks backed out from this model are similar to those in a model of perfect competition (Eaton, Kortum, Neiman, and Romalis, 2011). So, while we have described a model of imperfectly competitive heterogeneous firms to capture the stylized facts about polluting industries, we have shown that firm heterogeneity does not play a central role in our results, and we conjecture that imperfect competition does not either.

Finally, we consider the assumption that pollution is proportional to a firm's outputs rather than to its inputs, which is implicit in equation (5). This paper's model implies that more productive firms emit less pollution only because they invest more in pollution abatement. One could imagine a different model in which more productive firms emit less pollution because they use fewer factor inputs to produce a unit of output. We investigated this alternative by removing the productivity term φ from equation (5), giving an expression for pollution which is proportional to inputs rather than outputs. We then re-derived expressions for a firm's chosen pollution emissions under that modified assumption. This modification turns out to produce identical expressions for firm-level and economy-level pollution emissions both in observed

[54] Notable exceptions include the work by David Popp and coauthors. See e.g., Popp (2002) and Popp, Newell, and Jaffe (2010).

data and in counterfactuals. This modification does produce a different mechanism by which productivity affects pollution—by decreasing factor inputs rather than increasing abatement investments. However, the magnitude of the effect of productivity on pollution in the two models is numerically equivalent.

8 Conclusions

Public observers once worried that U.S. economic growth would lead to increasingly dangerous levels of pollution. Instead, U.S. air quality has improved dramatically. This paper focuses on the U.S. manufacturing sector and assesses four candidate explanations for why pollution emissions have fallen since 1990. The first explanation is that increasing production of pollution-intensive goods in China, Mexico, and other foreign countries has decreased U.S. pollution. Second, environmental regulation may have led to adoption of increasingly effective abatement technologies. Third, if productivity decreases pollution intensity, then rising productivity may have decreased pollution emissions. Lastly, consumer preferences may have led people to prefer goods that require less pollution to produce.

We begin with a statistical decomposition which shows that almost all of the change in pollution emissions from U.S. manufacturing is due to changes in pollution intensity within narrowly defined product-categories. To quantify the importance of trade, regulation, productivity, and preferences, we build on recent trade and environmental research to develop a model of heterogeneous firms that choose optimal investments in pollution abatement in response to environmental regulation. Although the methods we use are typically applied to research questions in international trade, we use them to address an open question in environmental economics: why are pollution emissions from U.S. manufacturing declining? While many quantitative models are used to forecast how untested future policies like carbon taxes or tariff reductions would affect pollution and welfare, we use our model to analyze the past—to recover the implied changes in environmental regulation and other shocks that firms actually faced in each year 1990-2008. We then use the implied changes to quantify how pollution would have changed under scenarios other than those that actually occurred.

The paper obtains three main conclusions. First, the fall in pollution emissions is due to decreasing pollution per unit output in narrowly defined manufacturing product categories, rather than reallocation across products or changes in the scale of real manufacturing output. Second, environmental regulation has grown increasingly stringent, and the pollution tax that explains U.S. data more than doubled between 1990 and 2008. Third, environmental regulation explains 75 percent or more of the observed reduction in pollution emissions from manufacturing. Trade costs, productivity improvements, and preferences play relatively smaller roles.

We believe there are a number of worthwhile extensions to the work presented here. First, like most models of monopolistic competition, our model assumes that prices are a constant markup over marginal cost. While theory makes predictions about how markups should respond to various competitive forces, the empirical evidence on the relationship between environmental policy and markups is limited.[55] Second, we believe the decomposition methodology developed in the paper could be applied to other settings. For example, there has been a keen interest in understanding why energy efficiency has improved across the

[55] The literature as to how markups respond to environmental regulatory stringency has only studied one industry, cement (e.g., Ryan (2012); Fowlie, Reguant, and Ryan (Forthcoming)).

United States. Has this been driven by efficiency standards? Rising energy prices? Population migration? One could adapt the tools from this paper to address this important policy question. We leave these extensions and questions for future work.

References

Anderson, J. E., and E. van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–192.

Antweiler, W., B. R. Copeland, and M. S. Taylor (2001): "Is Free Trade Good for the Environment?," *American Economic Review*, 91(4), 877–908.

Arkolakis, C., A. Costinot, D. Donaldson, and A. Rodriguez-Clare (2012): "The Elusive Pro-Competitive Effects of Trade," Unpublished Mimeograph, Yale University.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.

Autor, D. H., D. Dorn, and G. H. Hanson (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103(6), 2121–68.

Axtell, R. L. (2001): "Zipf Distribution of U.S. Firm Sizes," *Science*, 293, 1818–1820.

Becker, R. (2005): "Air pollution abatement costs under the Clean Air Act: evidence from the PACE survey," *Journal of Environmental Economics and Management*, 50(1), 144–169.

Becker, R. A., and R. J. Shadbegian (2005): "A Change of PACE: Comparing the 1994 and 1999 Pollution Abatement Costs and Expenditures Surveys," *Journal of Economic and Social Measurement*, 30(1), 63–95.

Berman, E., and L. T. Bui (2001a): "Environmental regulation and labor demand: evidence from the South Coast Air Basin," *Journal of Public Economics*, 79, 265–295.

——— (2001b): "Environmental Regulation and Productivity: Evidence from Oil Refineries," *Review of Economics and Statistics*, 83(3), 498–510.

Bernard, A. B., S. J. Redding, and P. K. Schott (2011): "Multi-Product Firms and Trade Liberalization," *Quarterly Journal of Economics*, 126(3), 1271–1318.

Bloom, N., C. Genakos, R. Martin, and R. Sadun (2010): "Modern Accounting: Good for the Environment or Just Hot Air?," *Economic Journal*, 120, 551–572.

Burstein, A., E. Morales, and J. Vogel (2013): "Accounting for Change in Between-Group Inequality," Unpublished Mimeograph, UCLA.

Busse, M. R., and N. O. Keohane (2007): "Market effects of environmental regulation: coal, railroads, and the 1990 Clean Air Act," *RAND Journal of Economics*, 38(4), 1159–1179.

Chaney, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–1721.

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

Cherniwchan, J. (2013): "Trade liberalization and the environment: Evidence from nafta and us manufacturing," Discussion paper, working paper.

Copeland, B. R., and M. S. Taylor (1994): "North-South Trade and the Environment," *Quarterly Journal of Economics*, 109(3), 755–787.

——— (2003): *Trade and the Environment: Theory and Evidence*. Princeton University Press.

Correia, A. W., C. A. Pope, D. W. Dockery, Y. Wang, M. Ezzati, and F. Dominici (2013): "Effect of Air Pollution Control on Life Expectancy in the United States," *Epidemiology*, 24(1), 23–31.

Costinot, A., and A. Rodriguez-Clare (2014): *Handbook of International Economics, Vol. 4* chap. Trade Theory with Numbers: Quantifying the Consequences of Globalization, pp. 197–262. North-Holland.

Curtis, M. (2012): "Who Loses under Power Plant Cap-and-Trade Programs? Estimating the Impact of the NOx Budget Trading Program on Manufacturing Employment," Mimeo, Georgia State.

Dekle, R., J. Eaton, and S. Kortum (2008): "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 55(3), 511–539.

Deschenes, O., M. Greenstone, and J. S. Shapiro (2013): "Defensive Investments and the Demand for Air Quality: Evidence from the NO_x Budget Program and Ozone Reductions," Mimeograph, MIT.

di Giovanni, J., A. A. Levchenko, and R. Ranciere (2011): "Power laws in firm size and openness to trade: Measurement and implications," *Journal of International Economics*, 85, 42–52.

Donaldson, D. (Forthcoming): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*.

Eastern Research Group (2001): "Uncontrolled Emission Factor Listing for Criteria Air Pollutants," Discussion paper, Emission Inventory Improvement Program.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79(5), 1453–1498.

Eaton, J., S. Kortum, B. Neiman, and J. Romalis (2011): "Trade and the Global Recession," NBER Working paper 16666.

Ederington, J., A. Levinson, and J. Minier (2008): "Trade Liberalization and Pollution Havens," *Advances in Economic Analysis & Policy*, 4.2.

Energy Information Administration (2011): "Annual Energy Review," Discussion paper, EIA.

Feenstra, R. C., and D. E. Weinstein (2010): "Globalization, markups, and the US price level," Discussion paper, National Bureau of Economic Research.

Finkelstein, A. (2004): "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry," *Quarterly Journal of Economics*, 119(2), 527–564.

Forslid, R., T. Okubo, and K.-H. Ultveit-Moe (2011): "International Trade, CO2 Emissions and Heterogeneous Firms," CEPR Discussion Paper No. 8583.

Fowlie, M. (2010): "Emissions Trading, Electricity Industry Restructuring, and Investment in Pollution Control," *American Economic Review*, 100(3).

Fowlie, M., M. Reguant, and S. P. Ryan (Forthcoming): "Market-Based Emissions Regulation and Industry Dynamics," *Journal of Political Economy*.

Gabaix, X. (1999): "Zipf's law for cities: an explanation," *The Quarterly Journal of Economics*, 114(3), 739–767.

——— (2009): "Power Laws in Economics and Finance," *Annual Review of Economics*.

Gamper-Rabindran, S. (2006): "NAFTA and the Environment: What can the Data Tell Us?," *Economic Development and Cultural Change*, 54(3), 605–633.

Greenstone, M. (2002): "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Amendments and the Census of Manufactures," *Journal of Political Economy*, 110(6), 1175–1219.

Greenstone, M., J. A. List, and C. Syverson (2012): "The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing," Unpublished Mimeograph, MIT.

Grossman, G. M., and A. B. Krueger (1995): "Economic Growth and the Environment," *Quarterly Journal of Economics*, 110(2), 353–377.

Henderson, V. J. (1996): "Effects of Air Quality Regulation," *American Economic Review*, 86(4), 789–813.

Hettige, H., P. Martin, M. Singh, D. Wheeler, and B. Mundial (1995): *The industrial pollution projection system*. World Bank.

Holladay, J. S. (2011): "Are Exporters Mother Nature's Best Friends?," Mimeograph, University of Tennessee.

Hsieh, C.-T., and R. Ossa (2011): "A Global View of Productivity Growth in China," Mimeography, University of Chicago.

Keller, W., and A. Levinson (2002): "Pollution Abatement Costs and Foreign Direct Investment Inflows to U.S. States," *Review of Economics and Statistics*, 84(4), 691–703.

Koo, A. Y. C. (1974): "Environmental Repercussions and Trade Theory," *Review of Economics and Statistics*, 56(2), 235–244.

Kreickemeier, U., and P. M. Richter (2014): "Trade and the Environment: The Role of Firm Heterogeneity," *Review of International*, 22(2), 209–225.

Krugman, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–959.

Leamer, E. E., and J. Levinsohn (1995): *Handbook of International Economics,*chap. International Trade Theory: The Evidence, pp. 1339–1394. Elsevier.

Levinson, A. (2009): "Technology, International Trade, and Pollution from US Manufacturing," *American Economic Review*, 99(5), 2177–2192.

——— (2014): "A Direct Estimate of the Technique Effect: Change in the Pollution Intensity of US Manufacturing 1990 - 2008," Unpublished Mimeo, Georgetown.

Levinson, A., and J. O'Brien (2013): "Environmental Engel Curves," Unpublished Mimeograph, Georgetown University.

Levinson, A., and M. S. Taylor (2008): "Unmasking the Pollution Haven Effect," *International Economic Review*, 49(1), 223–254.

Luttmer, E. G. J. (2007): "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122(3), 1103–1144.

Martin, L. A. (2011): "Energy efficiency gains from trade: greenhouse gas emissions and India's manufacturing sector," Mimeograph, Berkeley ARE.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

Melitz, M. J., and G. I. P. Ottaviano (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75, 295–316.

Melitz, M. J., and S. J. Redding (2014): *Handbook of International Economics, Vol. 4* chap. Heterogeneous Firms and Trade, pp. 1–54. North-Holland.

Muller, N. Z., and R. Mendelsohn (2009): "Efficient Pollution Regulation: Getting the Prices Right," *American Economic Review*, 99(5), 1714–1739.

Muller, N. Z., R. Mendelsohn, and W. Nordhaus (2011): "Environmental Accounting for Pollution in the United States Economy," *American Economic Review*, 101(5), 1649–1675.

Pierce, J. R., and P. K. Schott (2012): "The Surprisingly Swift Decline of U.S. Manufacturing Employment," NBER Working Paper No. 18655.

Popp, D. (2002): "Induced Innovation and Energy Prices," *American Economic Review*, 92(1), 160–180.

Popp, D., R. G. Newell, and A. B. Jaffe (2010): *Handbook of the Economics of Innovation*chap. Energy, the Environment, and Technological Change, pp. 873–938. Academic Press.

Ryan, S. P. (2012): "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, 80(3), 1019–1061.

Shapiro, J. S. (2013): "Trade, CO2, and the Environment," Unpublished Mimeograph, Yale.

Simonvska, I., and M. E. Waugh (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92(1), 1–210.

Syverson, C. (2011): "What Determines Productivity?," *Journal of Economic Literature*, 49(2), 326–365.

U.S. Census Bureau (2008): *Pollution Abatement Costs and Expenditures: 2005*. U.S. Government Printing Office.

Walker, R. (2013): "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce," *Quarterly Journal of Economics*, 128(4), 1787–1835.

Walker, W. R. (2011): "Environmental regulation and labor reallocation: Evidence from the Clean Air Act," *The American Economic Review*, 101(3), 442–447.

Wooley, D., and E. Morss (2013): *Clean Air Act Handbook*. Clark Boardman Callaghan.

Figures and Tables

Figure 1: Trends in Manufacturing Pollution Emissions and Real Output

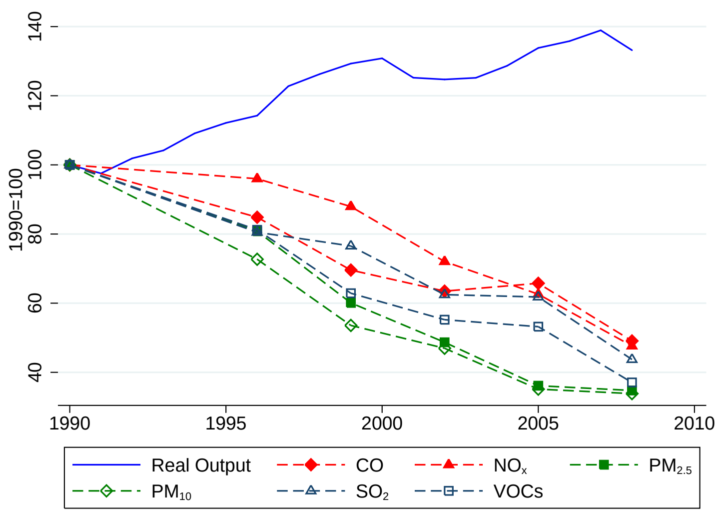

Note: This figure plots U.S. manufacturing real output and manufacturing pollution emissions in the years 1990-2008. Real output is measured annually from the NBER-CES database, where we have deflated industry level output by the NBER-CES industry-specific output price deflators. Manufacturing emissions come from the EPA's National Emissions Inventory, measured in years 1990, 1996, 1999, 2002, 2005, and 2008. Real output and pollution emissions are normalized to 100 in 1990.

Figure 2: Plant Level Pollution Intensity vs. Total Factor Productivity

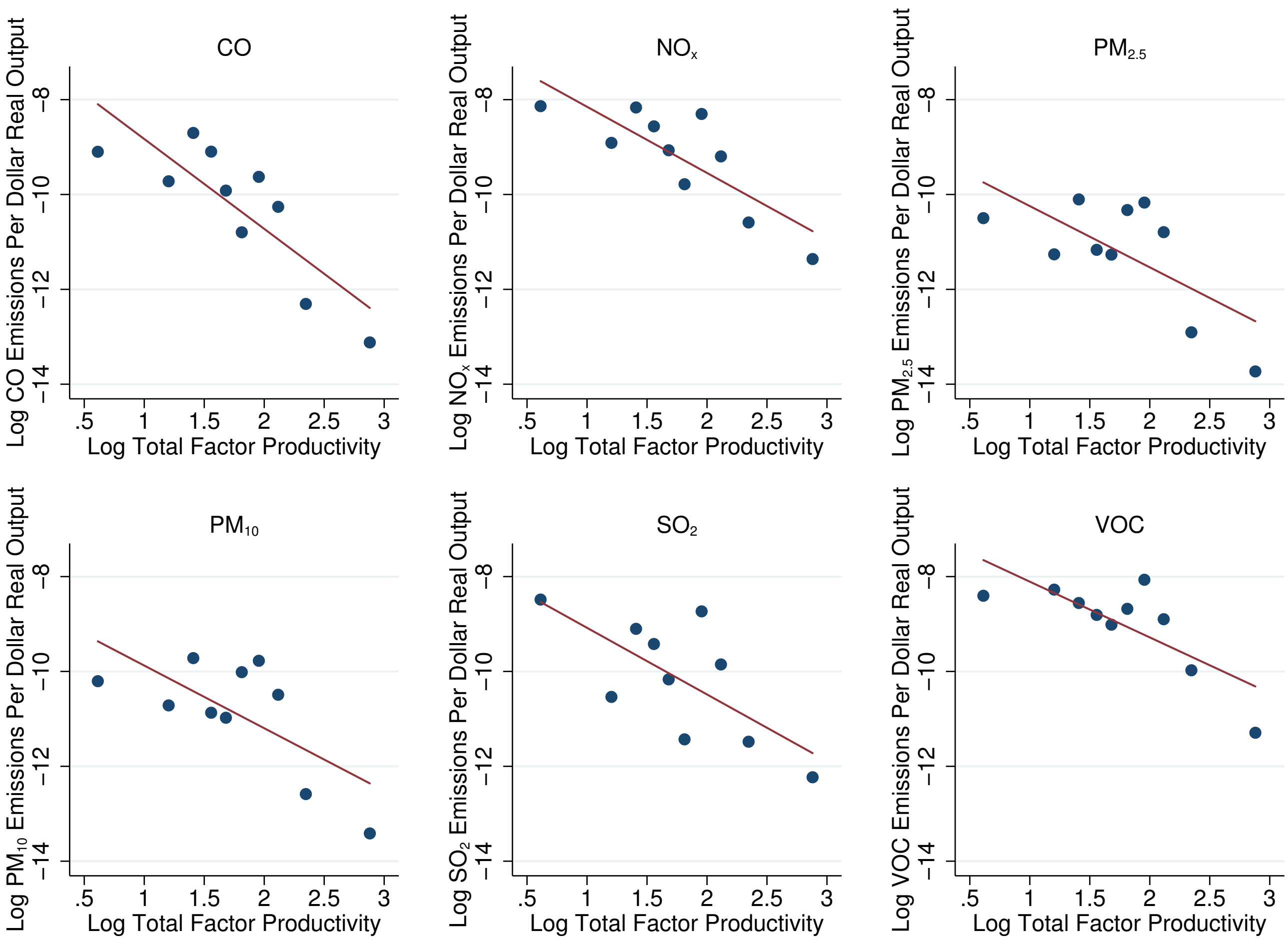

NOTE: This figure plots the relationship between plant-level total factor productivity and pollution per unit of output for the U.S. Manufacturing sector in 1990, separately for 6 different pollutants. The plant-level productivity measure is constructed from the U.S. Annual Survey of Manufacturers, using a total factor productivity index measure. We divide the sample into 10 deciles based on this plant-level productivity measure. We then compute the mean values of log productivity and log pollution per unit of real output within each decile, weighting the decile mean by plant-level inventory-adjusted, real output. Each pollutant scatter plot is accompanied by a linear fit, relating plant-specific emissions intensities to total factor productivity at the same plant. The line is fit to the entire sample, not simply the decile means. See Appendix C.1 for additional details.

Figure 3: Nitrogen Oxides Emissions From United States Manufacturing

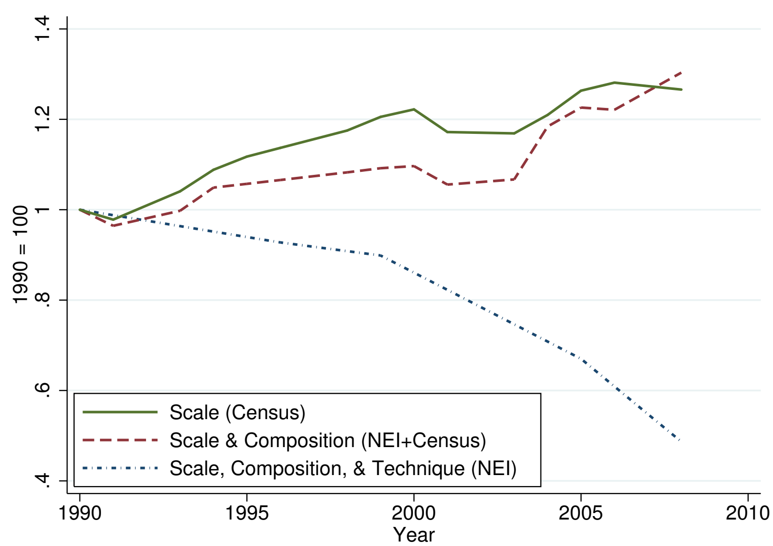

Notes: This figure plots the observed and counterfactual trends in NO_x emissions based on the statistical decomposition from equation (2). The top line plots the counterfactual for what emissions would have looked like in a world with the same composition of goods and techniques of production as was observed in the base year, 1990. The middle line represents what emissions would have looked like if we maintained the same production techniques (defined as emissions per unit of output) as in the base year, 1990. The final line represents the actual observed emissions trends, which consists of changes to both the scale, composition, and techniques associated with production since 1990. Source: NBER-CES database, ASM, and NEI.

Figure 4: Counterfactual U.S. Manufacturing Emissions of NO_x Under Subsets of Shocks, 1990-2008

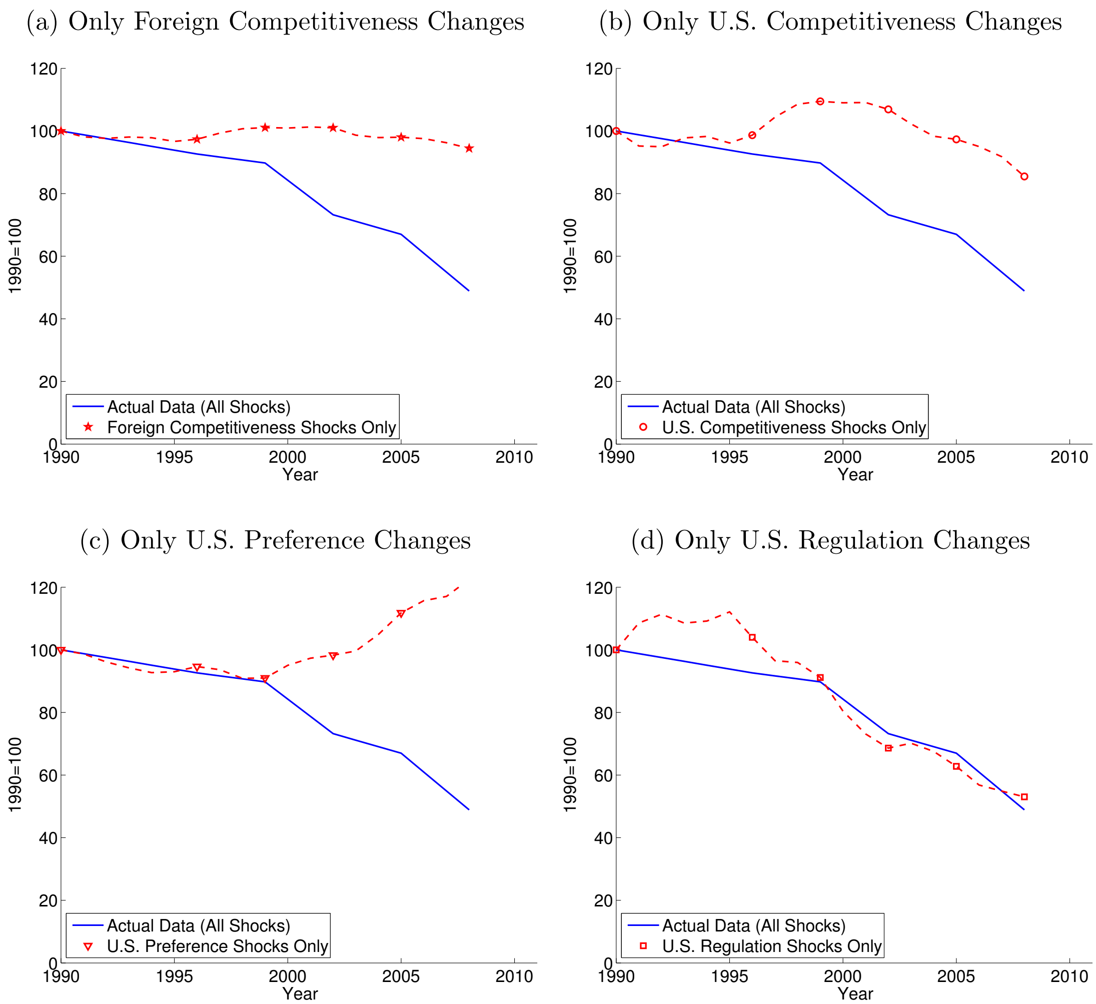

NOTES: This figure plots four separate counterfactual exercises, describing the actual and counterfactual time path of NO_x emissions in the U.S. manufacturing sector. The solid blue line displays the actual time path of emissions, and the dotted red line shows the counterfactual emissions in a scenario where only a single explanatory factor is allowed to take on its actual historical values. The scenario for each explanatory factor, or "shock", is indicated in the subfigure headings. All other explanatory factors are constrained to take their base year, 1990, values. The year 1990 NO_x emissions have been normalized to 100 in all figures. The star, circle, triangle, and square markers on the dashed lines show the years 1990, 1996, 1999, 2002, 2005, and 2008, when pollution data from NEI are observed rather than linearly interpolated.

Figure 5: Counterfactual U.S. Manufacturing Pollution Emissions Under Subsets of Shocks, 1990-2008

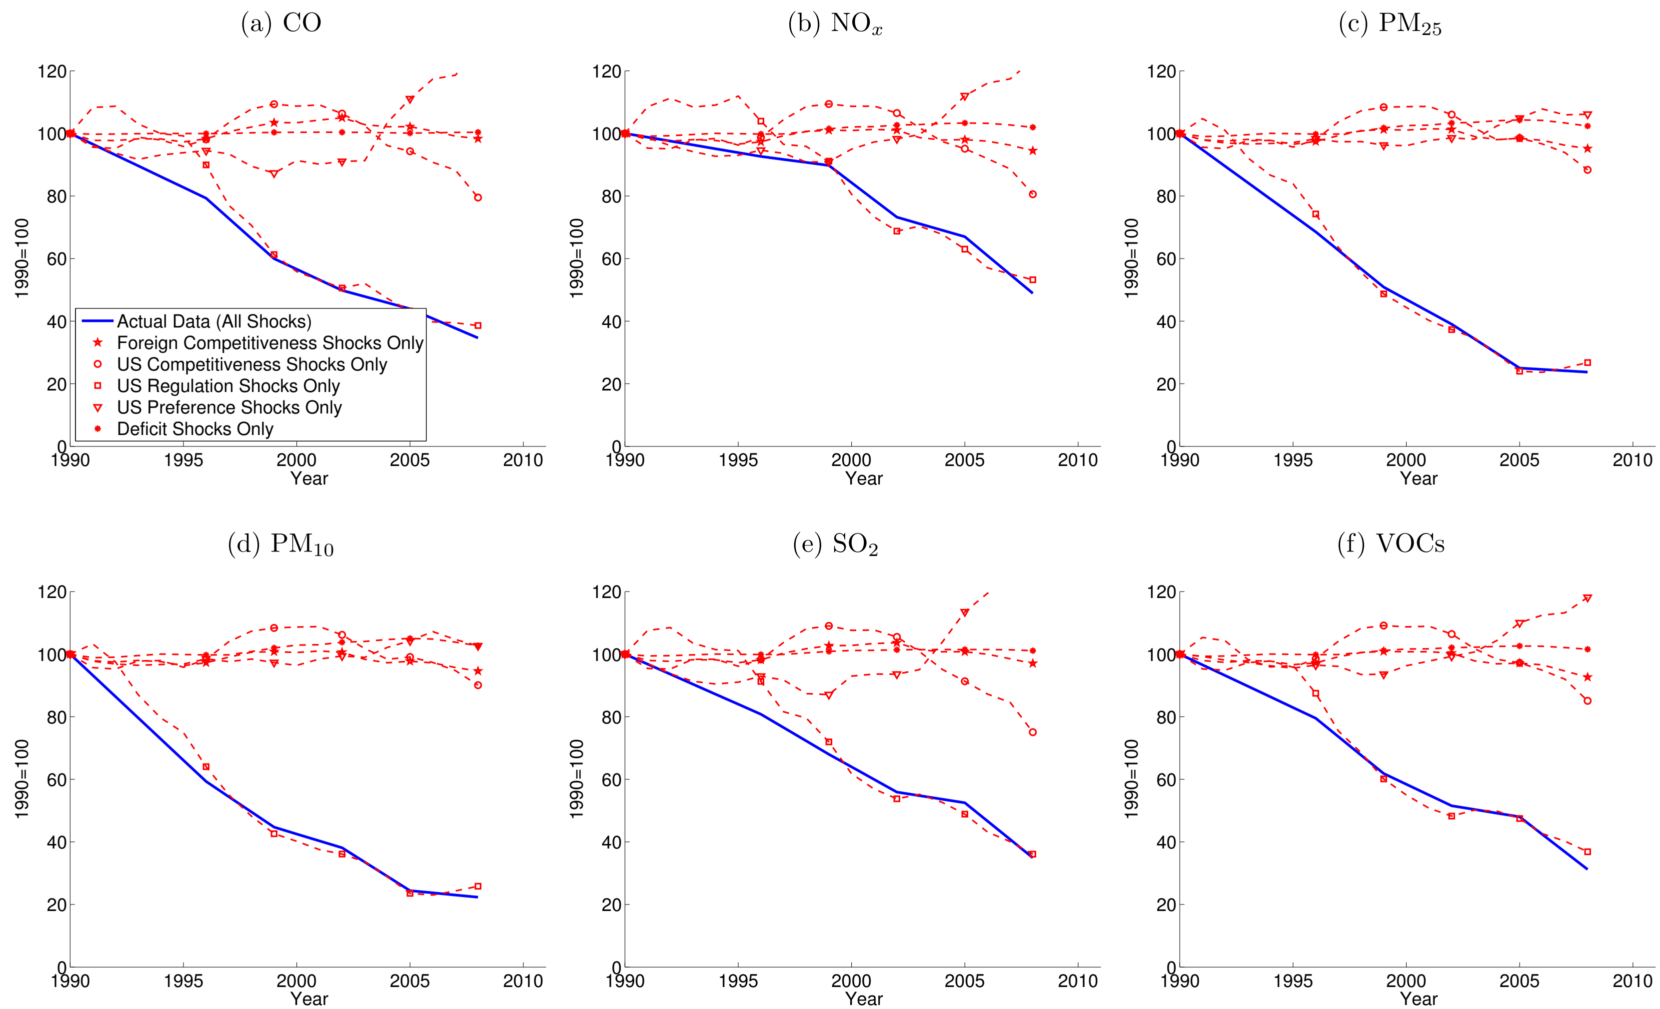

NOTES: This figure plots a separate counterfactual exercise for each pollutant and shock. Each subfigure plots the actual and counterfactual time path of the indicated pollutant emissions in the U.S. manufacturing sector. The solid blue line displays the actual time path of emissions, and the dotted red lines show the counterfactual emissions in a scenario where only a single explanatory factor is allowed to take on its actual historical values. The scenario for each explanatory factor, or "shock," is indicated in the legend in Subfigure 5a. For each counterfactual, all other explanatory factors are constrained to take their base year, 1990, values. The year 1990 values have been normalized to 100 in all figures. The star, circle, triangle, and square markers on the dashed lines show the years 1990, 1996, 1999, 2002, 2005, and 2008, when pollution data from NEI are observed rather than linearly interpolated.

Figure 6: NO_x Pollution Tax Changes as a Function of NO_x Budget Trading Program Status

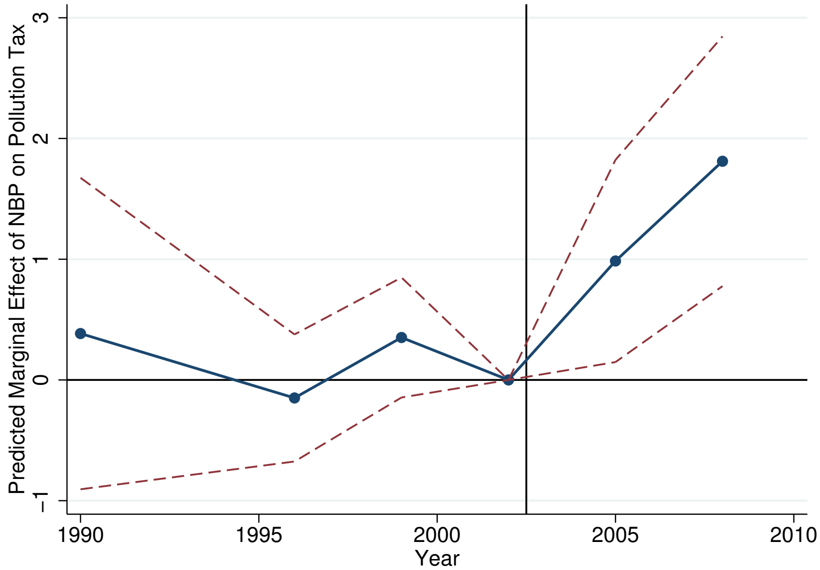

Notes: This figure reports regression coefficients from an event-study version of equation (21) in the text. The coefficients are plotted in blue and represent the time path of pollution taxes in polluting industries of NBP regions in the years just before and just after the NBP rollout, measured relative to a counterfactual. The dashed red lines represent 95 percent confidence intervals. The dependent variable is the model-driven measure of pollution taxes for a region×industry×year. The regression model includes region×year fixed effects η_{st}, industry×year fixed effects γ_{jt}, and region×industry fixed effects. Standard errors are clustered by industry×region.

Figure 7: Analysis of Carbon Dioxide Emissions

(a) Pollution Taxes for Air Pollution and for CO_2

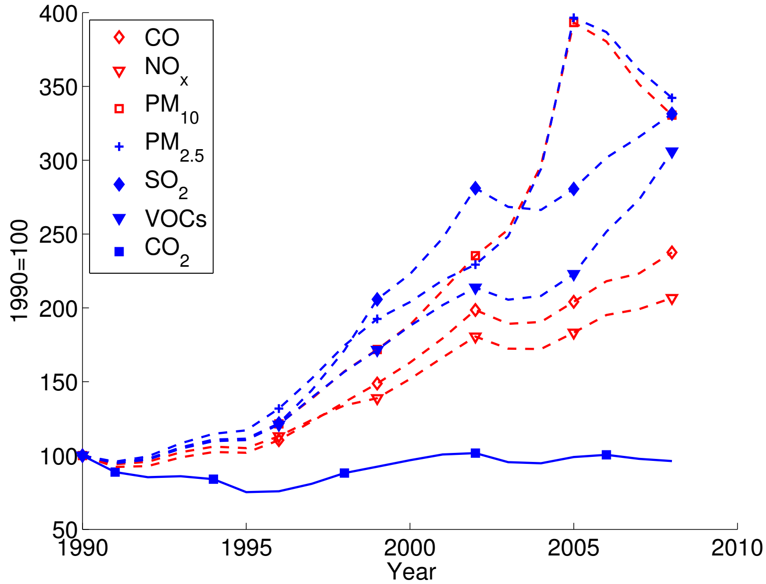

(b) Historic CO_2 Decomposition

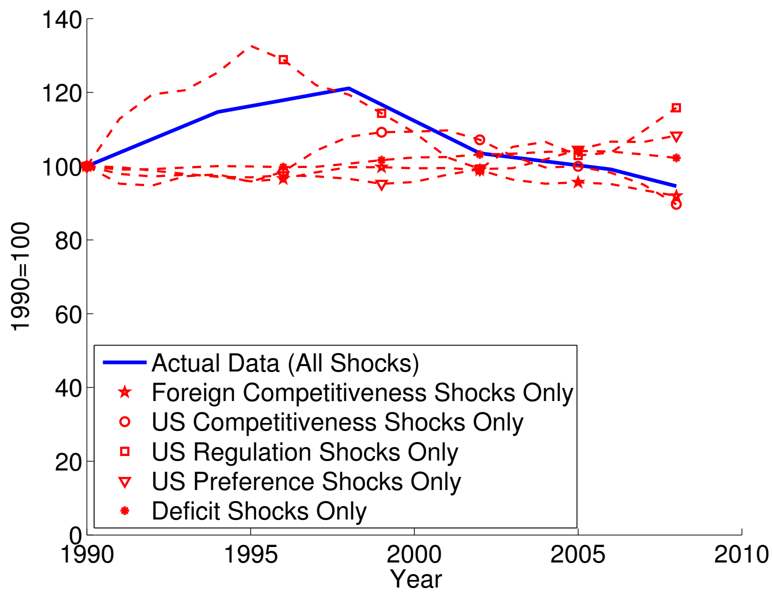

NOTES: Subfigure (a) plots implicit pollution taxes recovered for each pollutant and year, including CO_2. Subfigure (b) shows the same decomposition as Figure 5, except for CO_2 rather than for a criteria pollutant. The star, circle, triangle, and square markers on the dashed lines show the years when pollution data from NEI or MECS are observed rather than linearly interpolated.

Table 1: Pollution Elasticity: Instrumental Variables Regressions, by Pollutant

	CO	NO_x (O_3)	PM_{10}	$PM_{2.5}$	VOC (O_3)	Total (Any)
	Panel A: First Stage					
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Nonattain}_{cp} \times \text{Polluter}_p$	-0.057***	-0.061***	-0.101	-0.126*	-0.063***	-0.058***
	(0.015)	(0.011)	(0.085)	(0.068)	(0.009)	(0.009)
	Panel B: Reduced Form					
$\text{Nonattain}_{cp} \times \text{Polluter}_p$	-7.386	-5.985	-9.474	-7.399	-7.812***	-5.346**
	(5.244)	(4.782)	(6.860)	(4.427)	(1.214)	(1.979)
	Panel C: Instrumental Variables					
Abatement Expenditure Ratio	130.030**	98.592	94.118	58.551	124.907***	91.604***
	(64.278)	(72.412)	(78.483)	(46.795)	(36.827)	(25.373)
N	≈3500	≈3500	≈3500	≈3500	≈3500	≈3500
First Stage F-Stat	14	30	1.4	3.4	52	42
	Panel D: Pollution Elasticity Parameter					
Pollution Elasticity (α)	0.008**	0.010	0.011	0.017	0.008***	0.011***
	(0.004)	(0.007)	(0.009)	(0.013)	(0.002)	(0.003)
County-NAICS FE	X	X	X	X	X	X

Notes: This table presents a series of regression coefficients from 18 separate regressions, one for each column of each Panel A through C. An observation is a county×industry×year, where industry is a 6 digit NAICS code. The dependent variable in Panel A is the same in each column and represents 1 minus the abatement cost share of county×industry×year production. The regressor of interest is an interaction between two indicator variables that denote whether the industry is in a county that was newly regulated (i.e. Nonattain_{cp}=1) and whether the industry is a polluting industry (i.e. Polluter_p=1). The variable "Nonattain" changes across columns, reflecting different pollutant-specific nonattainment designations as indicated in the column headings. Parentheses in the column headings describe the type of nonattainment used as the regressor. The dependent variable in Panels B and C represent the emissions intensity, defined as pollution emissions per dollar of real output. The dependent variable in Panels B and C changes in each column, where the pollution emissions are indicated in the column headings. Panel C presents the instrumental variable estimates of pollution intensity regressed on abatement cost shares, which in practice represents the ratio of the estimates presented in Panel A and Panel B. Lastly, Panel D transforms the regression estimates in Panel C to back out a measure of α for each pollutant, where the standard errors are calculated using the delta method. Robust standard errors are in parentheses, clustering by commuting zone. Source: ASM, NEI, PACE.

Table 2: Parameter Estimates

Industry	Tons Pollution Per Dollar Costs (1)	Pollution Elasticity (α) (2)	Input Share (3)	Elasticity of Substitution (σ) (4)	Pareto Shape Parameter (θ) (5)	Shape Parameter Standard Error (6)
Food, Beverages, Tobacco	2.52	0.004	0.74	3.79	4.81	(0.13)
Textiles, Apparel, Fur, Leather	1.70	0.003	0.79	4.87	5.38	(0.10)
Wood Products	10.83	0.017	0.83	5.94	8.30	(0.17)
Paper and Publishing	10.63	0.017	0.79	4.80	4.29	(0.10)
Coke, Refined Petroleum, Fuels	17.50	0.027	0.88	8.18	17.52	(1.67)
Chemicals	10.72	0.017	0.70	3.28	4.13	(0.08)
Rubber and Plastics	2.98	0.005	0.78	4.59	5.02	(0.08)
Other Non-metallic Minerals	19.73	0.031	0.73	3.66	3.39	(0.11)
Basic Metals	30.36	0.048	0.85	6.66	9.72	(0.50)
Fabricated Metals	1.28	0.002	0.79	4.77	5.60	(0.06)
Machinery and Equipment	1.49	0.002	0.76	4.25	4.30	(0.14)
Office, Computing, Electrical	2.42	0.004	0.81	5.24	5.07	(0.15)
Radio, Television, Communication	0.51	0.001	0.79	4.66	4.13	(0.23)
Medical, Precision, and Optical	1.82	0.003	0.65	2.89	2.09	(0.06)
Motor Vehicles, Trailers	0.94	0.001	0.82	5.62	5.29	(0.18)
Other Transport Equipment	1.19	0.002	0.74	3.88	3.27	(0.13)
Furniture, Other, Recycling	2.51	0.004	0.73	3.77	4.77	(0.03)
Mean Across Industries	7.01	0.011	0.77	4.76	5.71	(0.23)

Notes: This table presents summary means and regression estimates for 17 separate industries, one per row, using data from a single year, 1990. Column 1 presents the total tons of pollution per dollar input costs for each industry, where pollution data comes from the NEI and data on input costs come from the ASM. Column 2 presents the industry specific pollution elasticity, which is calculated using the economy-wide estimate of 0.011 from Table 1, scaled across industries by the tons pollution per dollar costs from column (1). Column 3 presents the input share that is defined as the ratio of costs to revenues using data from the ASM. We deflate revenues and input expenditures using industry-specific price output and input price deflators, respectively. Column 4 displays the industry-specific elasticity of substitution, which is calculated from equation (14). Columns 5 and 6 present regression estimates and standard errors for the Pareto shape parameter, derived from equation 15. The actual parameter is a non-linear transformation of the regression coefficient, where the reported standard errors are calculated using the delta method, clustering by four digit NAICS code.

Appendix A Figures and Tables

Figure A1: Emissions Statistical Decomposition From United States Manufacturing

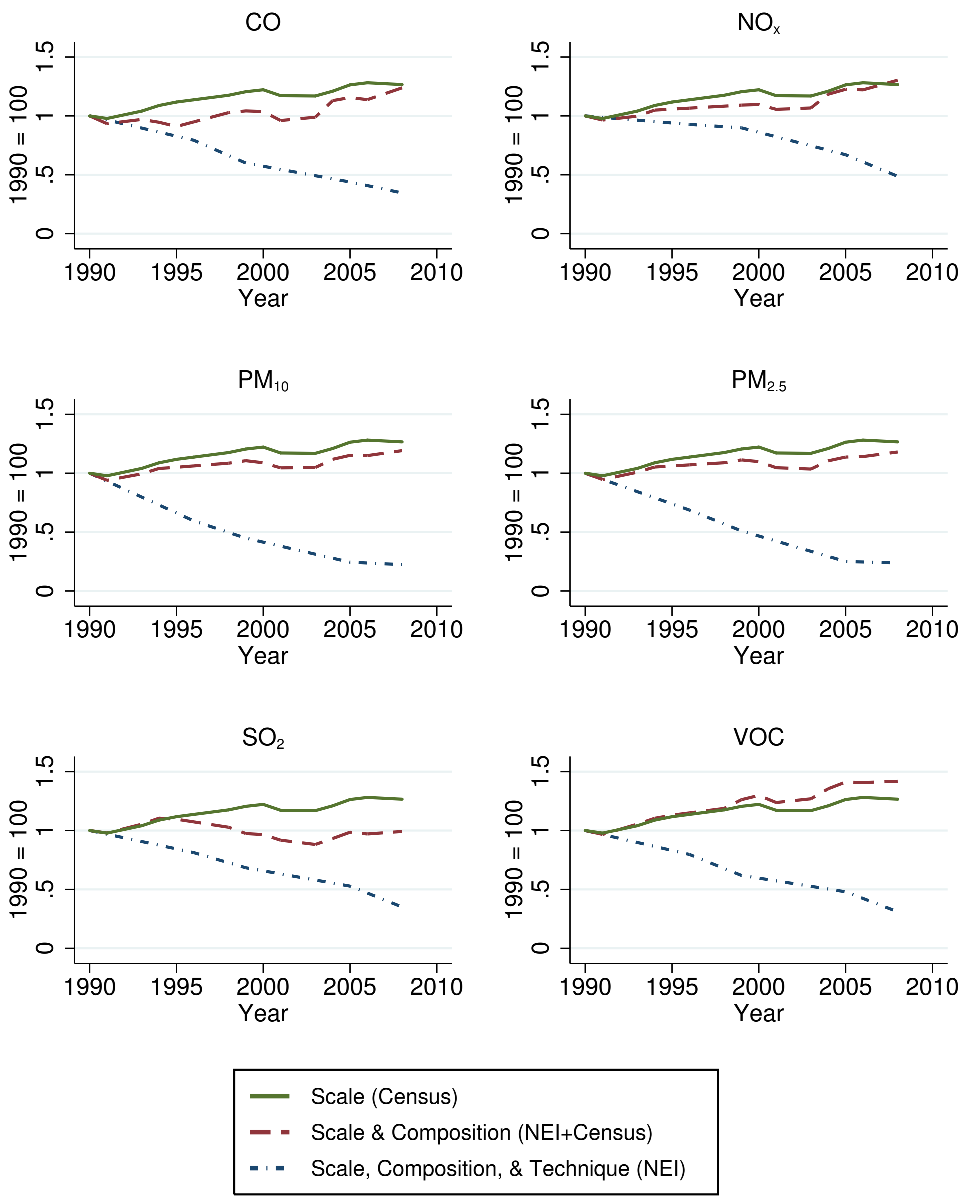

NOTES: This figure plots the observed and counterfactual trends in emissions for 6 separate pollutants based on the statistical decomposition from equation (2). The solid line of each panel plots the counterfactual for what emissions would have looked like in a world with the same composition of goods and techniques of production as was observed in the base year, 1990. The dashed line represents what emissions would have looked like if we maintained the same production techniques (defined as emissions per unit of output) as in the base year, 1990. The dashed-dotted line represents the actual observed emissions trends, which consists of changes to both the scale, composition, and techniques associated with production since 1990. Source: NBER-CES database, ASM, and NEI.

Figure A2: Historic Values of Shocks, 1990-2008.

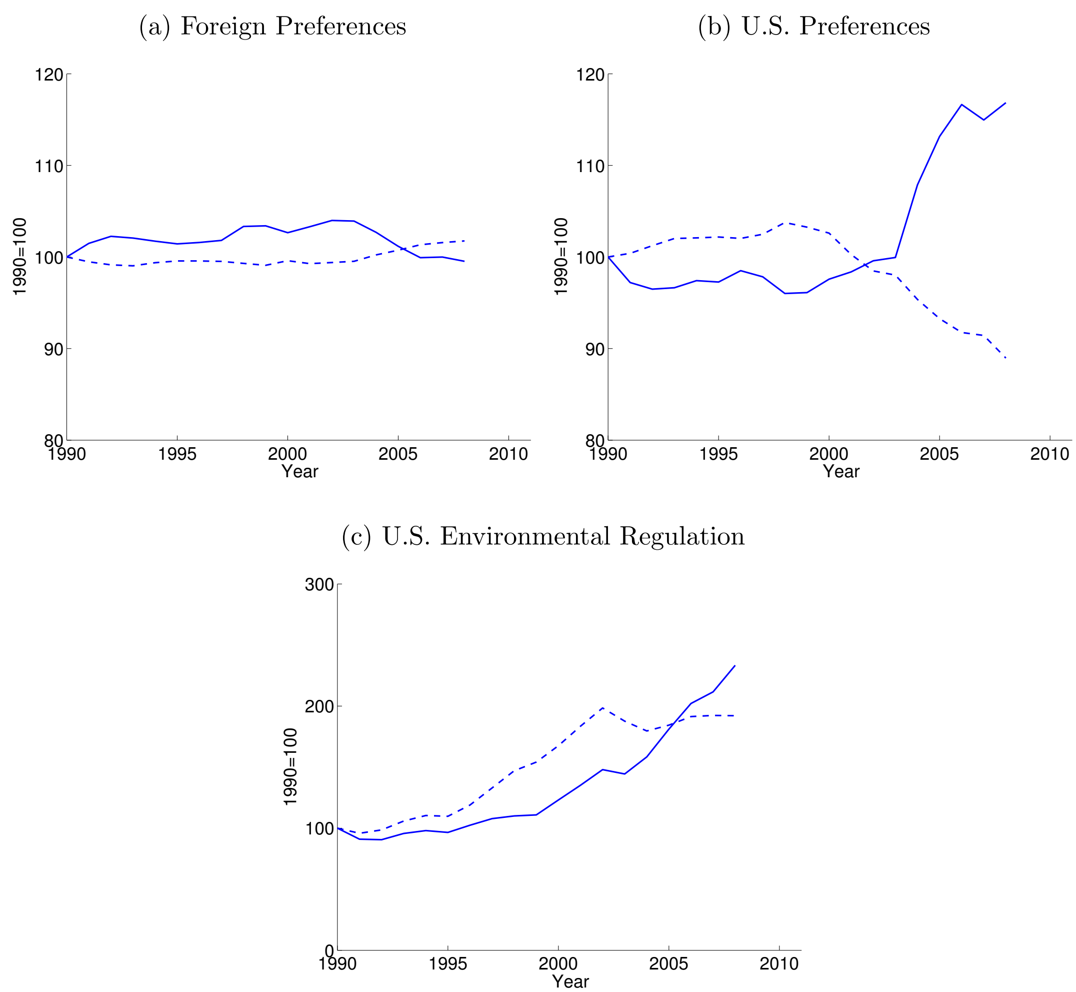

Notes: This figure plots the time path of shocks that we recover from the model outlined in Section 3 and derived using equations (16)-(19). The model delivers the value of the indicated shock for each of the 17 industries in our sample in each year. Here, we summarize the results graphically by plotting the unweighted mean separately for both dirty industries (solid line) and clean industries (dotted line).

Figure A3: Historic Values of Endogenous Variables, 1990-2008.

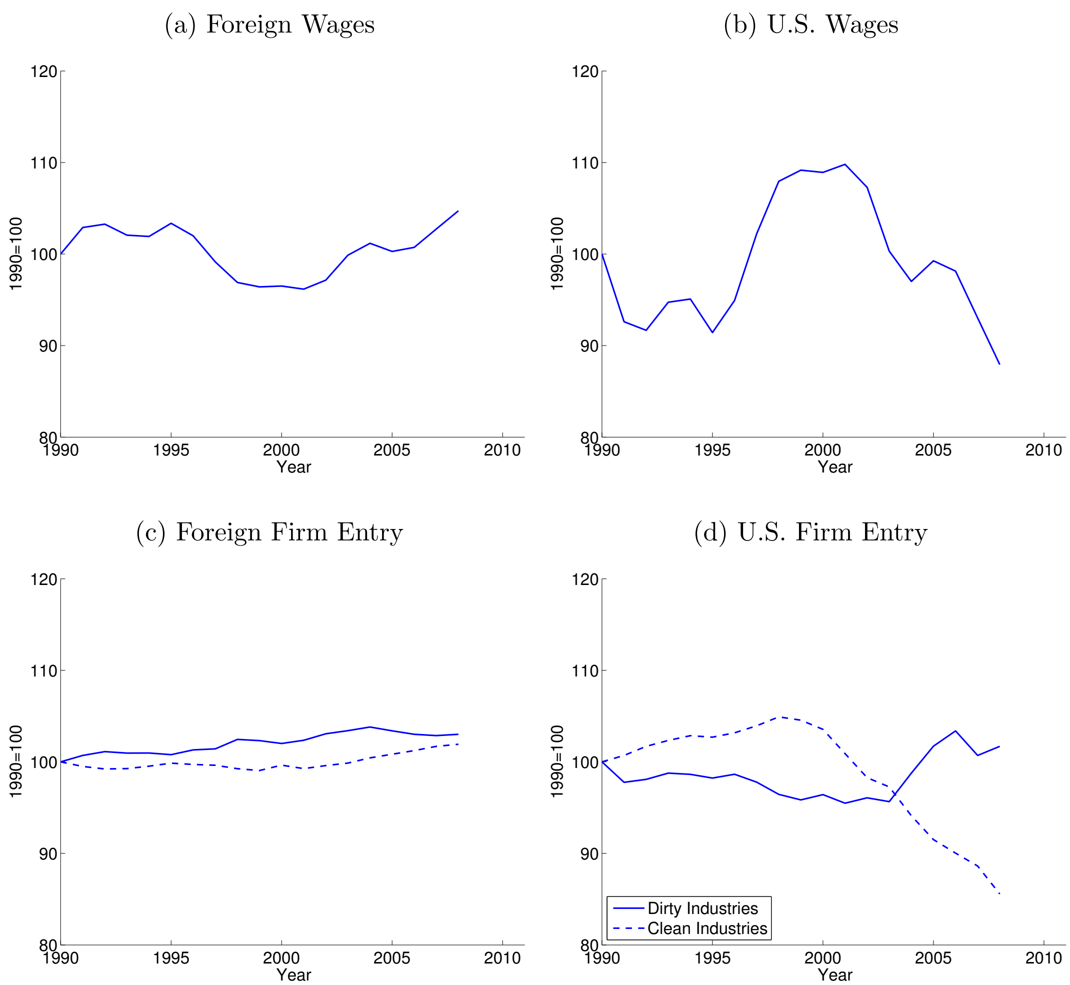

Notes: This figure plots the time path of endogenous variables that we recover from the model outlined in Section 3. The model delivers the value of firm entry changes for each of the 17 industries in our sample in each year. Here, we summarize the results graphically by plotting the unweighted mean for the indicated country-year. In subfigures (c) and (d) we plot the unweighted mean separately for both dirty industries or clean industries.

Figure A4: Counterfactual U.S. Manufacturing Pollution Under Different Economic Environments

(a) Shocks to Competitiveness of Clean and Dirty Foreign Industries

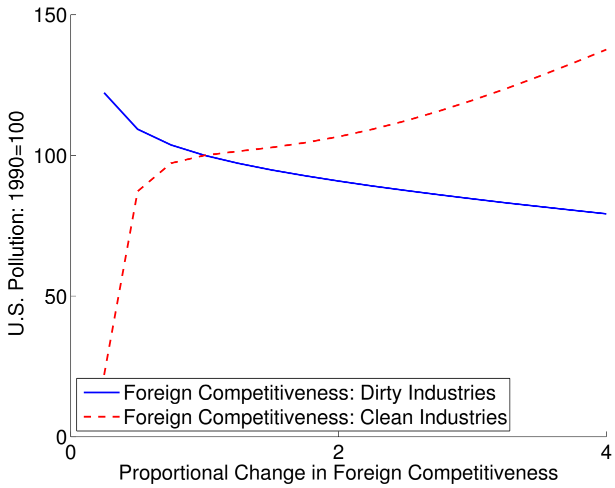

(b) Shocks to Competitiveness of Clean and Dirty U.S. Industries

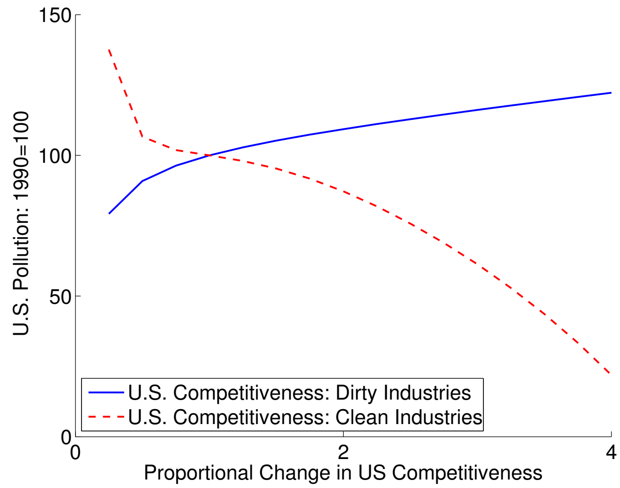

(c) Shocks to U.S. Environmental Regulation

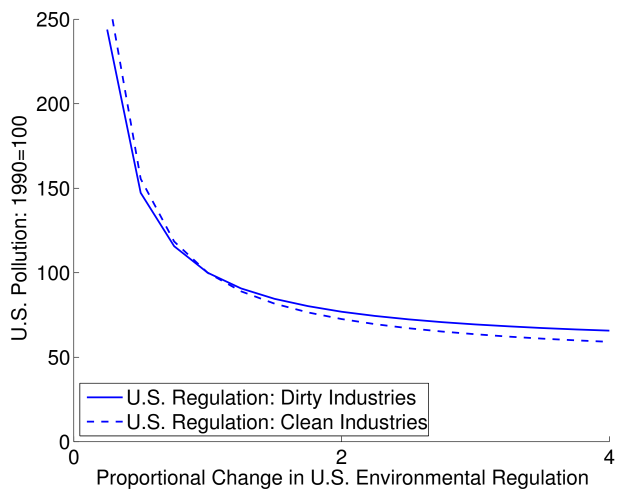

Notes: This figure plots 32 separate counterfactual exercises for each line of each subfigure. Each figure considers a counterfactual where the indicated shocks take on a value ranging from 0.25 to 8 times baseline levels (in 0.25 increments). For each counterfactual, we measure the resulting change in NO_x pollution, which is indicated on the y-axis, with the baseline value normalized to 100. The x-axis describes the counterfactual change in the indicated shock as a proportion of the baseline value. The scenario for each explanatory factors, or "shock", is indicated in the subfigure headings. Subfigures (a) and (b) conduct the counterfactual exercises separately for "clean" and "dirty" industries, as described in the text. All other explanatory factors are constrained to take their base year, 1990, values.

Figure A5: Characteristics of Fuels for Manufacturing, 1993-2008

(a) Share of Heat Input from Each Fuel (b) Potential SO_2 Emissions per BTU, by Fuel

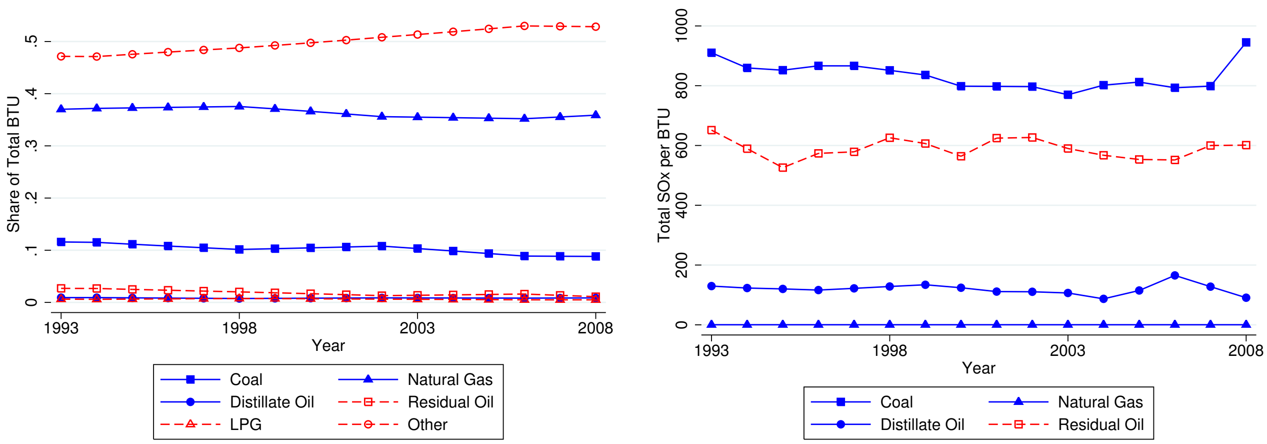

NOTES: This figure plots the time path of fuel use of different fuels in the U.S. Manufacturing industry from 1993-2008. Subfigure (a) plots BTUs of heat input from each fuel divided by total BTUs of fuel, all as used in U.S. manufacturing. The category "Other" includes electricity, steam, and anything else. Subfigure (b) describes pollution emissions that would be released in the absence of abatement technology per BTU, by fuel type. These emissions evolve over time due to changes in the sulfur content of coal and petroleum and to the physical tons of fuel per BTU. Source: EIA, MECS.

Figure A6: Potential Pollution Per BTU from Fossil Fuel Combustion, All Manufacturing, 1993-2008

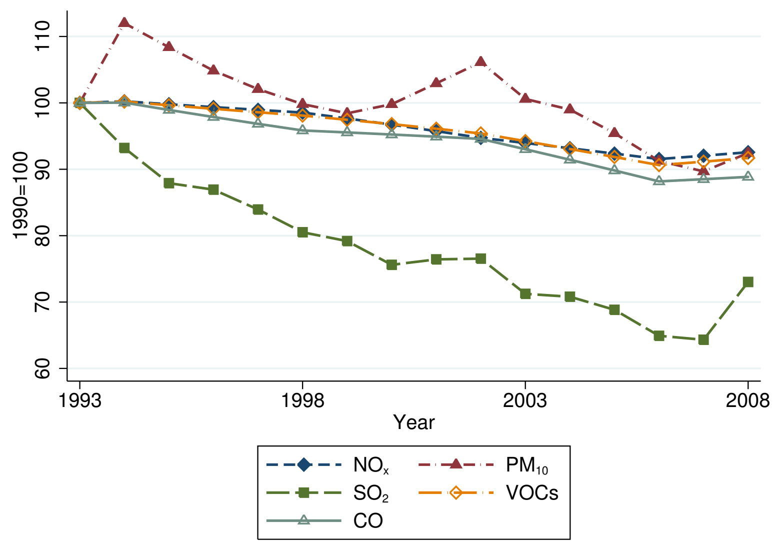

NOTES: This figure plots the time path of potential pollution per BTU in manufacturing given the mix of fossil fuels used in manufacturing in each year. The lines represent the evolution of pollution emissions for the indicated pollutant in the absence of end-of-pipe abatement. Values for the year 1990 are normalized to 100. Source: EIA, MECS.

Table A1: Industry Definitions

Code	Description	ISIC Rev. 3 Codes
1	Food, beverages, tobacco	15-16
2	Textiles, apparel, fur, leather	17-19
3	Wood products	20
4	Paper and publishing	21-22
5	Coke, refined petroleum, nuclear fuel	23
6	Chemicals	24
7	Rubber and plastics	25
8	Other non-metallic minerals	26
9	Basic metals	27
10	Fabricated metals	28
11	Machinery and equipment	29
12	Office, accounting, computing, and electrical machinery	30-31
13	Radio, television, communication equipment	32
14	Medical, precision, and optical, watches, clocks	33
15	Motor vehicles, trailers	34
16	Other transport equipment	35
17	Furniture, manufactures n.e.c., recycling	36-37

NOTES: This table presents the industry definitions used in the analysis and their corresponding two-digit International Standard Industrial Classification, third revision (ISIC Rev. 3) codes.

Table A2: Estimates of Pollution Elasticity, by Pollutant

Industry	Total (Main Estimates) (1)	CO (2)	NO_x (3)	PM_{10} (4)	$PM_{2.5}$ (5)	SO_2 (6)	VOCs (7)
Food, Beverages, Tobacco	0.0040	0.0017	0.0052	0.0060	0.0065	0.0046	0.0055
Textiles, Apparel, Fur, Leather	0.0027	0.0006	0.0029	0.0018	0.0021	0.0030	0.0081
Wood Products	0.0170	0.0238	0.0106	0.0184	0.0254	0.0086	0.0158
Paper and Publishing	0.0167	0.0152	0.0212	0.0125	0.0175	0.0209	0.0115
Coke, Refined Petroleum, Fuels	0.0275	0.0297	0.0269	0.0077	0.0105	0.0368	0.0244
Chemicals	0.0168	0.0172	0.0235	0.0067	0.0076	0.0132	0.0253
Rubber and Plastics	0.0047	0.0010	0.0040	0.0022	0.0027	0.0039	0.0186
Other Non-metallic Minerals	0.0310	0.0052	0.0536	0.0994	0.0724	0.0375	0.0091
Basic Metals	0.0477	0.0831	0.0191	0.0206	0.0272	0.0417	0.0155
Fabricated Metals	0.0020	0.0004	0.0016	0.0009	0.0011	0.0013	0.0089
Machinery and Equipment	0.0023	0.0019	0.0020	0.0019	0.0024	0.0015	0.0054
Office, Computing, Electrical	0.0038	0.0052	0.0018	0.0017	0.0021	0.0040	0.0034
Radio, Television, Communication	0.0008	0.0006	0.0006	0.0003	0.0003	0.0008	0.0021
Medical, Precision, and Optical	0.0029	0.0002	0.0081	0.0028	0.0044	0.0023	0.0052
Motor Vehicles, Trailers	0.0015	0.0004	0.0011	0.0004	0.0004	0.0012	0.0062
Other Transport Equipment	0.0019	0.0003	0.0020	0.0012	0.0012	0.0019	0.0064
Furniture, Other, Recycling	0.0039	0.0005	0.0027	0.0024	0.0030	0.0037	0.0157

Notes: This table presents estimates of the pollution elasticity for each industry and pollutant. Column (1) corresponds to column (2) of Table 2 that is calculated using the economy-wide estimate of 0.011 from Table 1, scaled across industries by the tons pollution per dollar costs from column (1) of Table 2. Columns (2)-(7) scale the economy-wide value of 0.011 according to the tons of each pollutant emitted per dollar of cost inputs.

Table A3: Share of Manufacturing Regulated by the NO_x Budget Trading Program (NBP)

	NO_x Emissions (Tons)			Number of Plants		
	NBP	Non-NBP	Share in NBP	NBP	Non-NBP	Share in NBP
Industry	(1)	(2)	(3)	(4)	(5)	(6)
All Manufacturing	163,099	1,091,048	0.13	267	79,016	0.003
Food, Beverages, Tobacco	3,901	79,824	0.047	23	3,446	0.007
Textiles, Apparel, Fur, Leather	1,068	12,594	0.078	5	5,704	0.001
Wood Products	2,294	39,445	0.055	15	4,039	0.004
Paper and Publishing	62,559	193,838	0.244	45	3,903	0.011
Coke, Refined Petroleum, Fuels	41,434	143,041	0.225	50	1,345	0.036
Chemicals	15,912	157,352	0.092	48	5,174	0.009
Rubber and Plastics	6,443	9,282	0.41	18	4,258	0.004
Other Non-metallic Minerals	9,976	324,837	0.03	11	8,177	0.001
Basic Metals	15,371	49,278	0.238	17	3,506	0.005
Fabricated Metals	401	7,726	0.049	2	8,016	0.000
Machinery and Equipment	517	9,749	0.050	6	5,782	0.001
Office, Computing, Electrical	822	12,059	0.064	6	5,287	0.001
Radio, Television, Communication	348	672	0.341	4	1,471	0.003
Medical, Precision, and Optical	348	36,320	0.009	4	3,725	0.001
Motor Vehicles, Trailers	87	6,073	0.014	2	3,536	0.001
Other Transport Equipment	465	4,449	0.095	5	4,858	0.001
Furniture, Other, Recycling	1,154	4,508	0.204	6	6,789	0.001

NOTES: This table presents descriptive statistics, by industry, pertaining to the fraction of manufacturing emissions and plants subject to the NO_x budget program. The baseline data come from plants that appear in the NEI in 2005. A plant in NEI is identified as a unique combination of a facility ID and a county. NBP plants are identified by a plant-level match between the EPA's Air Program Markets Data and the NEI. Columns (1)-(3) calculate the amount of NO_x emissions and share of NO_x emissions under the NBP. NO_x emissions in the table include the annual emissions which NEI reports, and not only the May-September emissions which NBP regulates. Columns (4)-(6) count the number of plants in each respective category. Since some plants report SIC industry codes that map into more than one ISIC code, we count a plant separately for each of the industries to which it is linked.

Table A4: Relationship Between Implied Manufacturing Pollution Tax and NO_x Budget Program

	(1)	(2)	(3)	(4)
$1[\text{NBP}_s] \times 1[\text{NBPRegulated}_j] \times 1[\text{Year}_t > 2002]$	1.195***	1.195***	1.186***	1.186***
	(0.422)	(0.424)	(0.404)	(0.405)
N	1583	1583	1583	1583
Industry×region FE	X	X	X	X
Industry×year FE		X		X
Region×year FE			X	X

Notes: This table reports regression coefficients from 4 separate versions of equation (21), one per column. The dependent variable in all regressions is the model-driven measure of pollution taxes for a region×industry×year. All specifications control for the lower order interaction terms that ensure identification of the difference-in-difference-in-differences regression equation presented above. Standard errors are clustered by industry×region and are in parentheses.

Table A5: How Does Fuel Mix Contribute to Pollution Declines?

	CO (1)	NO_x (2)	PM_{10} (3)	SO_2 (4)	VOCs (5)	Mean (6)
Actual Decrease in Pollution by Year 2008, as % of 1993 Level	46.90	51.40	60.80	51.60	59.10	53.96
Implied Decrease in Pollution by Year 2008, as % of 1993 Level	11.20	7.40	7.60	27.00	8.30	12.3
Implied Decreases Divided by Actual Decrease	0.24	0.14	0.12	0.52	0.14	0.23

Notes: This table presents estimates for the predicted fraction of manufacturing emissions reductions that may be explained by fuel switching. The first row calculates the actual decrease in manufacturing emissions observed in the NEI. The second row uses fuel mix and emission factor data from the EIA and MECS to identify the amount of emissions reductions that could be explained by changes in the type of fuel used in manufacturing. The implied decrease in pollution refers to potential pollution per BTU of fuel input, as defined in equation (22). The last line divides the second row by the first, to try to identify the fraction of the observed emissions reductions between 1993 and 2008 that could be explained by the types of fuels used in manufacturing production.

Appendix B Theory

This appendix presents intermediate results of the model in more detail.

As in most models with constant elasticity of substitution preferences, consumer demand for variety ω in destination country d is

$$q_{od,s}(\omega) = \frac{(p_{od,s}(\omega))^{-\sigma_s}}{(P_{d,s})^{1-\sigma_s}} E_{d,s}$$

where the price index is

$$P_{d,s} = \left[\sum_i \int_{\omega \in \Omega_{o,s}} p_{id,s}(\omega)^{1-\sigma_s}\right]^{\frac{1}{1-\sigma_s}}$$

Firms engage in monopolistic competition. They choose prices $p_{od,s}$ and abatement investments ξ to maximize profits. The first-order condition for pollution abatement is

$$1 - \xi = \left(\frac{w_o}{\varphi t_{o,s}} \frac{\alpha_s}{1-\alpha_s}\right)^{\alpha_s} \tag{23}$$

This first-order condition shows that more productive firms (higher φ) invest more in pollution abatement, leading them to emit less pollution. Combining this with the first-order condition for output implies that prices equal a constant markup $\frac{\sigma_s}{\sigma_s-1}$ over marginal costs:

$$p_{od,s} = \frac{\sigma_s}{\sigma_s - 1} \frac{c_{o,s}\tau_{od,s}}{\varphi^{1-\gamma\alpha_s}}$$

where

$$c_{o,s} \equiv \frac{(t_{o,s})^{\alpha_s}(w_o)^{1-\alpha_s}}{(\alpha_s)^{\alpha_s}(1-\alpha_s)^{1-\alpha_s}}$$

Two conditions are important to the equilibrium. The zero cutoff profit $\varphi_{od,s}$ describes the productivity level which makes a firm earn zero profits from exporting to destination d, and therefore makes a firm indifferent about whether to export to d:

$$\varphi_{od,s} = \left(\frac{\sigma_s}{\sigma_1 - 1} \frac{c_{o,s}\tau_{od,s}}{P_{d,s}} \left(\frac{\sigma_s w_d f_{od,s}}{E_{d,s}}\right)^{\frac{1}{\sigma_s-1}}\right)^{\frac{1}{1-\alpha_s}}$$

The free entry condition requires that in equilibrium, entrepreneurs must earn zero ex ante expected profit from paying the fixed cost $f^e_{o,s}$ to draw a productivity value:

$$f^e_{o,s}\frac{\theta_s - (\sigma_s - 1)(1-\alpha_s)}{(\sigma_s - 1)(1-\alpha_s)} = \sum_d \frac{(b_{o,s})^{\theta_s}}{(\varphi_{od,s})} \frac{w_d}{w_o} f_{od,s}$$

We now describe several important equations that we obtain using the zero cutoff profit and free entry conditions. The value of bilateral trade can be written as follows:

$$X_{od,s} = \frac{M^e_{o,s}\left(\frac{w_o}{b_{o,s}}\right)^{-\theta_s}(\tau_{od,s})^{-\frac{\theta_s}{1-\alpha_s}}(f_{od,s})^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}}(t_{o,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}}{(P_{d,s})^{-\frac{\theta_s}{1-\alpha_s}}} \left(\frac{E_{d,s}}{w_d}\right)^{\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} (w_d)\,\chi_s$$

where $E_{d,s}$ denotes the expenditure of country d on goods from sector s, and where we have collected parameters into the constant $\chi_s \equiv \frac{(\sigma_s)^{1-\frac{\theta_s\sigma_s}{(\sigma_s-1)(1-\alpha_s)}}}{(\sigma_s-1)^{-\frac{\theta_s}{1-\alpha_s}}}\frac{(\alpha_s)^{\frac{\alpha_s\theta_s}{1-\alpha_s}}}{(1-\alpha_s)^{-\theta_s}}\frac{\theta_s}{\theta_s-(1-\alpha_s)(\sigma_s-1)}$. The labor market clearing condition can be written as

$$L_d = \sum_s \left[M^e_{d,s} f^e_{d,s} \Big(\theta_s + 1 + \frac{\alpha_s \theta_s}{1-\alpha_s} \Big) + \frac{\theta_s - (\sigma_s - 1)(1-\alpha_s)}{\sigma_s \theta_s} \beta_{d,s} \frac{(w_d L_d - NX_d)}{w_d} + \widetilde{NX}_{d,s} \right] \tag{24}$$

where the trade imbalance term is summarized by $\widetilde{NX}_{d,s} = \beta_{d,s}\frac{NX_d}{w_d} - \frac{NX_{d,s}}{w_d}\frac{(\sigma_s-1)(\theta_s-\alpha_s+1)}{\sigma_s\theta_s}$. The price index may be written as follows:

$$(P_{d,s})^{-\frac{\theta_s}{1-\alpha_s}} = \sum_o M^e_{o,s} \left(\frac{w_o}{b_{o,s}} \right)^{-\theta_s} (\tau_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} (f_{od,s})^{1-\frac{\theta_s}{(1-\alpha_s)(\sigma_s-1)}} (t_{o,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}} \left(\frac{E_{d,s}}{w_d} \right)^{\frac{\theta_s}{(1-\alpha_s)(\sigma_s-1)}-1} \chi_s$$

Finally, we can write the two equilibrium conditions in levels. The first equilibrium condition is simply labor market clearing (6). The second equilibrium condition is as follows:

$$f^e_{o,s} \frac{\sigma_s \theta_s}{(\sigma_s - 1)(1-\alpha_s)} = \sum_d \frac{(w_o)^{-1} \left(\frac{w_o}{b_{o,s}} \right)^{-\theta_s} (\tau_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} (f_{od,s})^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} (t_{o,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}}{\sum_i M^e_{i,s} \left(\frac{w_i}{b_{i,s}} \right)^{-\theta_s} (\tau_{id,s})^{-\frac{\theta_s}{1-\alpha_s}} (f_{id,s})^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} (t_i)^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}} E_{d,s} \tag{25}$$

To analyze counterfactuals, we use the following expression for the change in expenditure shares (7):

$$\hat{\lambda}_{od,s} = \frac{\hat{M}^e_{o,s} \left(\frac{\hat{w}_o}{\hat{b}_{o,s}} \right)^{-\theta_s} (\hat{\tau}_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} \left(\hat{f}_{od,s} \right)^{1-\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}} \left(\hat{t}_{o,s} \right)^{-\frac{\alpha_s\theta_s}{1-\alpha_s}}}{\left(\hat{P}_{d,s} \right)^{-\frac{\theta_s}{1-\alpha_s}}} \left(\frac{\hat{\beta}_{d,s}}{\hat{w}_d} \frac{\hat{w}_d w_d L_d - N\hat{X}_d NX_d}{w_d L_d - NX_d} \right)^{\frac{\theta_s}{(\sigma_s-1)(1-\alpha_s)}-1} \tag{26}$$

where the change in the price of index for a sector is given by

$$\hat{P}_{o,s} = \left[\left(\frac{\hat{E}_{d,s}}{\hat{w}_d} \right)^{\frac{\theta_s}{(1-\alpha_s)(\sigma_s-1)}-1} \sum_o \lambda_{od,s} \hat{M}^e_{o,s} \left(\frac{\hat{w}_o}{\hat{b}_{o,s}} \right)^{-\theta_s} (\hat{\tau}_{od,s})^{-\frac{\theta_s}{1-\alpha_s}} (\hat{f}_{od,s})^{1-\frac{\theta_s}{(1-\alpha_s)(\sigma_s-1)}} (\hat{t}_{o,s})^{-\frac{\alpha_s\theta_s}{1-\alpha_s}} \right]^{-\frac{1-\alpha_s}{\theta_s}} \tag{27}$$

Appendix C Data Overview and Additional Empirical Details

Appendix C.1 Pollution Intensity and Total Factor Productivity: Details

Figure 2 plots the relationship between plant level pollution intensity in total factor productivity. This section provides additional details underlying this figure. We use the 1990 Annual Survey of Manufacturers (ASM) which provides information on input decisions and total output at the plant level. We match the ASM to the National Emissions Inventory (NEI) using name and address matching techniques. Details of the match can be found in Appendix C.2. We use the sampling weights in the ASM to adjust plant-level output by the inverse sampling probability of a plant in the survey.

For each plant and each pollutant we divide total emissions by inventory adjusted real output.[56] We use industry-specific price deflators from the CES-NBER Productivity database to deflate output using an SIC-4, industry-level index normalized to 1 in 2008. We then compute a plant-level index measure of total factor productivity, using a Cobb-Douglas production technology and assuming constant returns to scale.[57] Production inputs include labor, capital, and materials. We approximate the output elasticities of production inputs using industry-level cost shares from the NBER-CES productivity database. All inputs were deflated using industry, input-specific price deflators from the NBER-CES productivity database.

We divide the sample into 10 deciles based on total factor productivity. We then compute the mean values of log productivity and log pollution per unit of real output within each decile, weighting the decile mean by plant-level inventory-adjusted, real output. Figure 2 plots the results for each of the six pollutants in our sample. Each pollutant scatter plot is accompanied by a linear fit, relating plant-specific emissions intensities to total factor productivity at the same plant. The line is fit to the entire sample, not simply the decile means.

Appendix C.2 Matching the 1990 Annual Survey of Manufacturers to the 1990 National Emissions Inventory

We match the 1990 Annual Survey of Manufacturers to the 1990 National Emissions Inventory using name and address string matching techniques. The ASM does not provide name and address information for plants, but the ASM can be linked to the Census Business Register via a unique, longitudinal identifier that does. The Business Register consists of the universe of establishments in the United States on an annual basis and forms the basis for the more commonly known and used Longitudinal Business Database (LBD).

We perform a match between the 1990 NEI and the Business Register for each Business Register year between 1985 and 1996. Both the NEI and the Business Register contain establishment level name and address information that we use to perform the match: county, state, SIC code (4-digit, 3-digit, and 2-digit), facility name, street, city, and zip code. We perform exact matching on county, state, and sic codes and "fuzzy" matching on facility name, street, city, and zip code. We use the "COMPGED" feature of SAS's PROC SQL to create a "generalized edit distance" score reflecting the degree of difference between two text strings. For each variable in which we use fuzzy matching techniques, we choose a score that minimizes both false positives and false negatives by visually checking the performance of the matches.

[56] Inventory adjusted total output is defined as the total value of shipments, minus the difference between finished goods inventory between the beginning and end of the year, minus the difference between work in progress inventory at the beginning and end of the period.

[57] Plant TFP is computed as 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 $j \in \{l, k, m, e\}$. Index productivity measures are common in the literature partly because they are easy to construct and also because they are a nonparametric first-order approximation to a general production function. See e.g., Syverson (2011).

We then iterate over combinations of the match variables listed above, selecting the match with the highest score in each round, and removing the residual observations from each dataset before matching again. At the end of the matching process, we are able to match 77.4% of the 1990 NEI manufacturing observations (i.e. SIC code between 2000-3999). The match percentage also reflects the fact that the ASM is a sample and not a survey, and thus we should not expect a match rate near 100%. In addition, the unmatched plant observations are not significantly different along emissions totals, relative to the matched observations. For each pollutant, we ran a plant-level regression of emissions on an indicator for whether the plant matched the census data, controlling for 4-digit SIC fixed effects and clustering standard errors by 4-digit SIC codes. Of 6 pollutants, the coefficient on the match variable is significant at the 10% level for PM_10 and $PM_{2.5}$; for other pollutants we cannot reject the null that a significant difference exists between matched and unmatched plants. For the plants that are matched and emit particulate matter, the matched plants tend to emit slightly more particulates than the unmatched plants.

Appendix C.3 Product Level Decomposition: Details

Section 2 in the text describes the statistical decomposition using the product-level production data from the Census and Annual Survey of Manufacturers (CMF and ASM, respectively). Here we provide additional details.

Bernard, Redding, and Schott (2011) provide a detailed overview of the Manufacturing Product trailer for research purposes. In terms of descriptive statistics, the typical two-digit SIC code in Manufacturing has 24 four-digit industries and 76 five-digit products, although there is heterogeneity across industries in the amount of product detail. For example, the number of products per sector ranges from a low of 12 in Leather (SIC 31) to a high of 178 in Industrial Machinery (SIC 35) (Bernard, Redding, and Schott, 2011).[58]

Within the CMF and ASM product trailers, there are several industries which report only aggregate product codes (i.e. within a 4-digit industry, more than 95% of output is produced in a product code that ends in "-", "0", or "W"). In the case that 50% or more of the product shipments within in a 4-digit SIC industry come from one of these aggregate product codes, we aggregate to the 4-digit SIC level.

There are two primary issues that emerge when looking at changes in the composition of products and how these affect manufacturing emissions over time. The first issue is associated with the introduction of new products; we calculate emissions factors using total product-level production and emissions in 1990. If new products are introduced after 1990, they will not have an emissions factor, and thus will lead to false inferences from the decomposition. In order to address product entry, we fold all new products into the adjacent product category, as defined by 5-digit SIC product codes. This implicitly assumes that the emissions factor from the new product is the same as the emissions factor calculated for the adjacent product code.

The second issue emerges from the transition between SIC and NAICS product code definitions between 1997 and 1998. We construct a product code crosswalk between 5-digit SIC product codes and 7 digit NAICS product codes. This allows us to construct a consistent 5-digit SIC by year dataset from 1990 until 2008. We develop this product-level SIC-NAICS concordance using 3 separate but complimentary strategies:

1. For the industries that only report aggregate product shipments (i.e. at the level of 4-digit SIC codes or 6-digit NAICS codes), we use the NBER-CES crosswalk which provides a linkage between 4-digit SIC codes and 6-digit NAICS codes. In the event that a 6-digit NAICS code maps into more than one 4-digit SIC code, the NBER-CES crosswalk provides value shares in order to apportion NAICS output to the relevant SIC code.

[58] As noted by Bernard, Redding, and Schott (2011), there is also substantial variation in the precision of product classifications. For example, Passenger Cars (SIC 37111) and Combat Vehicles (SIC 37114) are examples of products in the Motor Vehicle industry (SIC 3711), while Textbook Binding and Printing (SIC 27323) and Religious Books, Binding and Printing (SIC 27323) are examples of products in the Book Printing industry (SIC 2732).

2. For products that are consistently reported at the NAICS 7-digit product level, we develop a crosswalk using the 1997 Census of Manufacturing product trailer. In 1997, Census collected both NAICS and SIC product codes which we use to build the crosswalk. For 7-digit NAICS product codes that map into more than one 5-digit SIC product code, we construct apportionment shares based on the fraction of total 1997 output that is split between the respective SIC codes.

3. Lastly, there are some 7-digit NAICS codes in years 1998+ that do not match either of the two crosswalks above. For these residual product codes, we use a crosswalk developed by the Bureau of Labor Statistics between SIC and NAICS product codes.[59] There are still some cases for which NAICS 7-digit product codes map into more than one 5-digit SIC product code. In these cases, the BLS does not provide relative output shares for the "many to 1" crosswalk that would allow us to apportion NAICS output to the relevant SIC product code. This lack of apportionment for split products means that the product series in years 1998+ will overstate the amount of output for NAICS product codes that map into multiple SIC product codes. We adjust for this structural break by multiplying the scale + composition line in years 1998+ by an adjustment factor. This adjustment factor is computed by fitting a linear trend to years 1996 and 1997 and projecting the 1998 point; the value that we multiply the observed 1998 value to recover the predicted 1998 value is the adjustment factor we use to scale all post-1998 output.

Once we have a consistent 5-digit SIC product-level dataset, we construct product shares in each year by taking the total product output produced in a given year and dividing that by total manufacturing output in that year. In non-Census years, we use the weights provided by the Census to scale up plant-level output by the inverse sampling probability of the survey. We multiply these product shares by the product-level emissions factors and sum over all products in a year. Lastly, we multiply this annual number by total manufacturing output in that year in order to recover the scale+composition line in Figures 3 and A1.

Appendix C.4 Additional Model Sensitivity: Alternative Parameter Values, and Solution Algorithms

Alternate Parameter Values

We now consider the sensitivity of the paper's main results to parameter estimates and model assumptions. Appendix Table C1 begins by investigating model sensitivity to alternative parameter specifications. The first row shows that by 2008, NO_x emissions from U.S. manufacturing were 48.90 percent of their 1990 values. The paper's main estimates imply that environmental regulation alone would have caused pollution emissions to equal 53.04 percent of their 1990 value by 2008 (column (3), row (2)). Rows 3 and 4 explore how sensitive this conclusion is to changes in the underlying Pareto shape parameter estimates. Because the Pareto distribution best approximates the size distribution for the upper tail of firms, our main estimates of these parameters use the largest 10 percent of firms in each industry. Estimating the Pareto shape parameters using the top 50 percent of firms in each industry, or using the top 25 percent of firms in each industry, hardly affects the main conclusions. These two alternatives imply that environmental regulation would have led NO_x emissions to be 53.31 or 53.26 percent of their 1990 value by 2008, which are extremely close to the main results.

Rows (5) and (6) of Appendix Table C1 explore sensitivity to changes in the pollution elasticity α_s. Row (5) assumes that the pollution elasticity is one-fourth of our estimated values, and row (6) assumes that the true values of α_s are four times the values of our main estimates. The former implies that environmental regulation alone would have led pollution emissions to be 53.49 percent of their 1990 value in 2008; the latter implies that environmental regulation alone would have led pollution emissions to be 49.41 percent of their

[59]Source: http://www.bls.gov/ppi/ppinaictosic15.htm (accessed on July, 1 2014).

1990 value by 2008. These alternative parameter values modestly affect the magnitude of how environmental regulation affects manufacturing NO_x emissions. However, across the four alternative sets of results, the qualitative conclusion persists that regulation explains most of the change in pollution.

Algorithm to Calculate Equilibrium

Appendix Table C2 explores the sensitivity of the results to different sets of starting values needed for the algorithm to solve systems of nonlinear equations (9) and (10). We randomly draw 1,000 different sets of starting values from the uniform distribution [0.75,1.25].[60] Each set of starting values represents changes in wages in each country and firm entry decisions in each country$\times$sector. The objective function appears to be fairly flat around the main set of results; different starting values obtain slightly different values of the changes in wages and firm entry decisions which are not numerically equivalent to the main results. However, column (1) shows that the differences between these equilibria are very small and appear only between the 29th and 31st decimal point. Because we only have 32 digits of calculation precision, these differences in equilibria may reflect numerical precision due to computational limits. We also report results using two alternative algorithms for solving systems of nonlinear equations—a trust-region reflexive algorithm and a Levenberg-Marquardt algorithm. Both yield very similar, though not numerically equivalent values of the objective function, and yield the same estimate of how regulation affects pollution.

Column (2) of Appendix Table C2 shows that the ratio of U.S. pollution emissions in 2008 relative to 1990 is equal to 52.53 in every set of starting values and algorithms we use. Across the thousand alternative sets of starting values, the standard deviation of this value is 9.72E-13. These results suggest that while the objective function is somewhat flat near the optimum, our quantitative conclusions are the same with other starting values or algorithms.

[60] We choose this range to cover the values of shocks observed in data and described in Appendix Figure A3. Some starting values well outside this region fail to converge.

Table C1: Sensitivity Analysis: U.S. Pollution Emissions in Counterfactual Divided by 1990 Emissions, Separately for Each Shock

	Foreign Competitiveness (1)	U.S. Competitiveness (2)	U.S. Environmental Regulation (3)	U.S. Preferences (4)	Trade Deficits (5)
1. Actual Change			48.90		
2. Main Estimate	94.46	85.48	53.04	122.56	101.95
3. No Firm Heterogeneity	94.27	91.50	53.05	122.54	101.92
4. Parameter θ: Top 50%	90.17	94.22	53.31	122.19	103.48
5. Parameter θ: Top 25%	91.41	93.27	53.26	122.29	101.09
6. Parameter α: 0.25× Main Estimates	94.44	73.11	53.49	122.57	101.95
7. Parameter α: 4× Main Estimates	94.37	108.61	49.41	122.54	101.95

NOTES: This table presents a set of sensitivity analyses for the main set of counterfactuals in the text. Row 1 presents the actual observed change in NO_x emissions between 1990 and 2008. Row 2 shows the main estimates from the model, where each column corresponds to a separate counterfactual. For example, column (1) shows that if foreign competitiveness took its true historic value and all other shocks were held fixed, then manufacturing emissions of NO_x in 2008 would have been 94.46 percent of their observed 1990 values. Row 3 shows counterfactuals in a model where parameters are chosen so all firms have the same productivity and there is no firm heterogeneity. Rows 4 and 5 explore sensitivity of these counterfactuals to changes in the Pareto shape parameters that govern the distribution of firm productivity. Rows 6 and 7 explore model sensitivity to changes in the estimated pollution elasticity. Row 8 incorporates an additional shock, a global increase in factor supplies of 62%, which equals the change in real global GDP between 1990 and 2008.

Table C2: Sensitivity to Starting Values and Algorithms

	Minimized Objective Function	2008 Change in Pollution, Regulation-only Counterfactual
Main Results	1.36E-30	53.04
Starting Values Randomly Chosen:		
Mean	1.92E-30	53.04
Minimum	3.45E-31	53.04
Maximum	8.58E-29	53.04
Standard Deviation	(5.25E-30)	(6.41E-13)
Algorithm: Trust-Region Reflective	9.61E-31	53.04
Algorithm: Levenberg-Marquardt	1.74E-30	53.04

NOTES: This table presents a set of sensitivity analyses for the counterfactuals presented in the paper. The table varies the starting values and the algorithm associated with solving the system of nonlinear equations used in our counterfactuals. The main results in the paper use a starting value equal to one and a trust-region, dogleg algorithm. This table presents results from a randomly chosen set of starting values drawn from a uniform distribution [0.75,1.25]. The table reports the mean, minimum, maximum, and standard deviation of the 1,000 different results. The last two rows of the table present results using two different minimization routines.

Appendix C.5 Additional Data Details

Manufacturing and Energy Consumption Survey

MECS is a nationally representative survey of the entire U.S. manufacturing sector which was conducted in 1991, 1994, 1998, 2002, 2006 and 2010. The survey includes around 15,000 plants in a typical year. We download publicly available tables reporting BTUs of energy used for the manufacturing sector and for several subsectors.[61] We extract data on energy used for heat, power, and electricity generation. Those data exclude energy used as feedstock (e.g., they exclude the petroleum which is a physical part of plastics). The data are reported in Table 3 or Table 4 of the MECS yearly tables. We focus on energy used as fuel because fuel combustion is likely to contribute to pollution emissions from the manufacturing sector, whereas energy used as a feedstock may not. In the 1998 MECS survey, which is in the middle of our sample period, about 71 percent of total BTUs of energy were used for fuel, and the remaining 29 percent for feedstock.

EPA Air Program Markets Data

In order to operate cap-and-trade programs like the NO_x Budget Trading Programs, the EPA maintains a public database listing each facility which participates in the program and its attributes.[62] We obtain a list of facilities which were regulated under the NO_x Budget Trading Program in each year of its operation, 2003-2007. Each facility includes identifying information such as longitude and latitude, address, and a generic industry description, which we use to link these data to the National Emissions Inventory.

[61] The data are online at http://www.eia.gov/consumption/manufacturing/index.cfm.

[62] The data are available at http://ampd.epa.gov/ampd/. These data were formerly called the Clean Air Markets Database.