Dirty and perverse: regulation-induced pollution substitution

Matthew Gibson*

April 16, 2014

Abstract

Pollution regulations may induce firm responses that undermine their effectiveness. By regulating air emissions in particular counties, the Clean Air Act (CAA) gives firms incentives to substitute: 1) toward polluting other media, like landfills and waterways; and 2) toward pollution from plants in other counties. Using EPA Toxic Release Inventory data, I examine the effect of CAA regulation on these types of substitution. Regulated plants increase their ratio of water to air emissions by 42 percent. In multi-plant firms, regulation of an average plant increases air emissions at unregulated plants by 17 percent, resulting in a net emissions increase. (JEL Q53, Q52, H23)

1 Introduction

Not only does air pollution increase mortality (Chay & Greenstone 2003b), there is growing evidence that it reduces productivity (Graff Zivin & Neidell 2012) and human capital (Sanders 2012). In the US the central check on

*University of California, San Diego; magibson@ucsd.edu; 9500 Gilman Drive # 0508, La Jolla, CA 92093-0508. I would like to thank Prashant Bharadwaj, Richard Carson, Andrew Chamberlain, Julie Cullen, Joshua Graff Zivin, Michael Greenstone, Kelsey Jack, Mark Jacobsen, Lynn Russell, Jeffrey Shrader, Junjie Zhang and participants in the UC San Diego environmental seminar for invaluable assistance with this project.

air emissions is the Clean Air Act (CAA). Most welfare analyses of the CAA have concluded that its benefits vastly exceed its costs (EPA 2011). As a result, researchers have generally ignored the possibility that firm responses to the CAA might reduce its effectiveness. The CAA regulates particular pollutants in particular counties, which may create incentives for firms to substitute among different forms of pollution. This paper tests two variants of this hypothesis: 1) Do firms respond to CAA regulation by polluting other channels, like landfills and waterways? (cross-media substitution); and 2) Do multi-plant firms substitute toward pollution from other plants? (spatial leakage).

If one of a county's air pollution monitors exceeds the CAA standard, the EPA designates the county as "non-attainment." The state then issues regulations to reduce that county's air pollution, including emissions requirements for industrial plants. Simple economic theory suggests that firms will respond to such regulations by substituting toward unregulated or less-regulated forms of pollution. The EPA is aware of the potential for substitution and in some cases has taken steps to mitigate it (EPA 2001). In particular, EPA is currently developing rules to restrict water emissions from power plants (ENS 2013). The potential for pollution substitution remains, however, in a wide variety of industries.

Several previous studies have examined the possibility of cross-media substitution and found little evidence for it. Sigman (1996) tests for substitution in chlorinated solvent releases by metals and manufacturing plants. The author finds no substitution driven by the CAA, but does find substitution driven by hazardous disposal prices. Greenstone (2003) tests for CAA-induced substitution in releases from the iron and steel industry and finds no evidence for it. Gamper-Rabindran (2009) models health-weighted VOC emissions by chemical manufacturers as a function of CAA non-attainment, proxying for output changes with employment changes. She finds no increased emissions into other media. My approach differs in several key respects: 1) motivated by a simple theoretical model, I use emissions ratios to disentangle output and substitution effects; 2) by estimating in levels rather than differences, I account for the discrete nature of many abatement decisions; 3) I control for spatial heterogeneity in regulation; and 4) I pool across industries to improve statistical power.

Using EPA Toxic Release Inventory (TRI) data, this study tests the substitution hypotheses outlined above by comparing regulated ("treated") plants in particulate non-attainment counties to unregulated plants. Both cross-media substitution and spatial leakage occur and responses can be large. Regulated plants increase their ratio of water to air emissions by 42 percent. In multi-plant firms, particulate regulation of an average plant increases air emissions at unregulated plants by 17 percent, resulting in a net emissions increase. Not all substitution responses are harmful; regulated plants increase their ratio of recycling to air emissions by 47 percent.

Additionally, this study follows the implications of Auffhammer et al (2009), which finds that the average monitor-level effect of CAA non-attainment is zero, but that the effect on non-attainment monitors is -11 to -14 percent. This suggests that regulators respond to non-attainment by focusing on problematic plants, rather than requiring uniform changes across all plants in a county. I demonstrate that only plants near non-attainment monitors are treated under the CAA. To my knowledge, previous work has not shown this pattern at the plant level. My analysis of substitution accounts for this and so avoids averaging changes at treated plants with null responses from untreated plants in nonattainment counties.

These findings are important not only for air pollution regulation, but for pollution control policy generally. If firms substitute among various forms of pollution, an optimal policy must consider not just a *plant's* emissions into a particular medium, but rather a *firm's* emissions across all media, in all locations. Such policy would set a firm's emissions price for each medium and location equal to the marginal damage from emissions (leaving no medium or location unpriced). While the optimal policy might not be feasible, a second-best policy might nonetheless benefit from considering possible substitution responses.

This analysis contributes generally to the literature on regulation in the presence of mispriced substitutes (e.g. Campbell 1991). It is the first work to document regulation-induced cross-media pollution substitution. It also contributes to the literature on pollution leakage. To date this literature has focused on international trade leakage (Davis & Kahn 2010, Levinson & Taylor 2008) and simulated carbon leakage (Fowlie 2009, Bushnell & Mansur 2011). Becker and Henderson (2000) find the CAA makes firms more likely to enter attainment counties, which might be considered a form of leakage. To the best of my knowledge, mine is the first study to find evidence of domestic emissions leakage across existing plants. The finding that the treatment effect of the CAA varies over space within a county raises questions about spatially heterogeneous welfare impacts that might form the subject of future research. Lastly, this study adds to the literature on the costs of the CAA (Walker 2011, 2013).

The rest of the paper is organized as follows. Section 2 provides background on abatement technology and regulations important to cross-media substitution. Section 3 discusses a simple theoretical model that informs my estimation. Section 4 describes the data, Section 5 presents estimating equations, Section 6 presents results, and Section 7 explores their robustness. Section 8 concludes.

2 Background

Under the Clean Air Act, the EPA sets air quality standards for six criteria pollutants: carbon monoxide (CO), nitrogen dioxide (NO2), particulate matter (PM), lead (Pb), sulfur dioxide (SO2), and volatile organic compounds (VOCs). For detailed information on particulate standards, which are the focus of this paper, see appendix table A2. The agency assigns a "non-attainment" designation to a county for a particular pollutant if at least one monitor exceeds the CAA standard in a given year. In what follows, I refer to a monitor that exceeds the annual standard as a non-attainment monitor. In my data (described in section 4), PM non-attainment lasts for an average of approximately 7 years. Conditional on PM non-attainment in at least one year, I observe on average 1 entry into non-attainment and .16 exits; most of these counties remain in non-attainment through 2012, the last year of my data. For 9 counties I observe two entries into non-attainment.

When a county receives a non-attainment designation, the state in which it is located must develop a State Action Plan detailing steps that will bring the county into attainment. These typically include "lowest achievable emissions rates" (LAER) equipment requirements and plant-specific emissions limits (Becker and Henderson 2000, 2001; Walker 2013). State and EPA enforcement mechanisms include fines, inspections, and withholding of federal highway funds (Becker and Henderson 2000, Chay and Greenstone 2005).

2.1 Abatement technologies and variable costs

If abatement costs were entirely fixed or if abatement were costless, plants would have no incentive to substitute in response. While abatement technologies do have large fixed costs, they also have substantial operating costs. Pollution control devices typically require substantial energy and may yield wastes that require costly disposal. Processes that employ catalysts require periodic replacement of the catalyst. These variable costs suggest that CAA non-attainment changes the relative price of air emissions for regulated plants. I catalog the most common air pollution control technologies below:

- Selective catalytic reduction (SCR): SCR converts nitrogen oxides to nitrogen and water. The reaction takes place in a titanium or platinum chamber at a 480-800 degrees Fahrenheit. For a large gas turbine, annual operation and maintenance costs are approximately 50% of capital cost (EPA undated, Farnsworth 2011).

- Electrostatic precipitation (ESP): Typically used to abate PM and lead, electrostatic precipitators create an electrical field through which particles pass, acquiring a charge. The particles then adhere to collecting electrodes. Particles are periodically removed from the collecting electrodes using vibration or water spray. Wastes are discharged to a landfill or reused (EPA undated, Farnsworth 2011).

- Fabric filters (a.k.a. "baghouses"): Typically used for PM and lead, fabric filters trap particles in tightly woven fabric. Filters require periodic replacement (Farnsworth 2011).

- Scrubbers (a.k.a. "flue gas desulfurization"): Liquid or solid absorbing materials are sprayed into flue gas. The residue must be collected by a downstream device such as an electrostatic precipitator (Farnsworth 2011).

- Coal washing: Coal is crushed and washed to reduce its sulfur content (DOE undated).

- Fuel switching: A plant may be able to switch to a fuel that results in lower air emissions, e.g. a power plant can switch to low-sulfur coal (DOE undated). Economic intuition suggests the new fuel must be weakly more expensive than the old, or the plant would have been using it prior to abatement.

- Incineration: Incineration is commonly used for VOC abatement. Unlike most other technologies, incineration is a final disposal method that typically outputs only carbon dioxide and water. Incinerators require the use of additional fuel (e.g. natural gas) to reach operating temperature, usually 1200 to 2000 degrees Fahrenheit. Sometimes platinum or metal oxide catalysts are used to facilitate combustion at lower temperatures, from 300 to 900 degrees Fahrenheit. EPA estimates that a typical thermal incinerator costs approximately $483,000 and a catalytic incinerator $889,000. Estimated annual operating costs are $422,000 and $316,000, respectively (Vatavuk et al 2000).

Many of these options are capital-intensive and will produce abrupt reductions in emissions once installed. This is one of the reasons I estimate in levels rather than differences, as described in Section 5. SCR and incineration decrease toxic air emissions but increase carbon emissions. Most of these technologies, particularly those used to remedy particulate air pollution, produce secondary waste streams that require disposal.

2.2 Pricing of water and land emissions

CAA-induced substitution will reduce welfare only if substitute emissions are unpriced or underpriced. Such is the case for many TRI pollutants and many emissions channels. The Safe Drinking Water Act (SDWA) and the Pollutant Priority List (PPL) for the Clean Water Act do not cover many TRI chemicals (Gamper-Rabindran 2009). For example, my TRI data contain 580 chemicals. The PPL lists 126 chemicals (EPA 2013). In addition, two recent Supreme Court decisions have limited the scope of the CWA. *Solid Waste Authority of Northern Cook County v. U.S. Army Corps of Engineers* removed CWA protection from "isolated" water bodies, including many wetland areas. *Rapanos v. United States* removed CWA protection from waterways that are not navigable year-round (EPA 2008). (Note that the Mississippi River would arguably have met such a definition in 2012.)

The Resource Conservation and Recovery Act (RCRA) governs many forms of toxic disposal on land. Coal combustion residuals are currently exempt from many provisions of the RCRA, though the EPA is attempting to regulate them (EPA 2010). Some mining and petrochemical wastes are also exempt (EPA 1999). Regulation of TRI-listed air pollutants that do not fall into one of the six CAA criteria categories varies by industry. Under the 1990 CAA Amendments, EPA develops industry-specific regulations governing the air release of 187 toxic chemicals ("air toxics"). EPA "...does not prescribe a specific control technology, but sets a performance level based on a technology or other practices already used by the better-controlled and lower emitting sources in an industry" (EPA undated). While the incomplete regulations governing water and land emissions suggest cross-media substitution may reduce welfare, a full welfare analysis is beyond the scope of this paper.

3 Theory

The following simple model informs my estimating equations for cross-media substitution. Suppose there are two pollution inputs. Let A be air emissions and W be water emissions. The CAA may be viewed as shift in relative prices $\frac{p_A}{p_W}$, with the increased price of air emissions having two components: 1) pecuniary cost, like the variable abatement cost described in section 2.1; and 2) non-pecuniary cost, for example the cost of incurring the displeasure of a regulator. The object of policy interest is unconditional factor demand W^*, incorporating firms' possible output response to regulation. Suppose a CES production function, so the firm problem becomes:

$$\max_{A,W} p_o \left(c_A A^\rho + c_W W^\rho\right)^{1/\rho} - p_A A - p_w W$$

Note that using a CES function assumes constant returns to scale. While this might be implausible over the whole range of a firm's technically feasible output levels, it is reasonable if output responses to the CAA are modest. The choice of CES does not impose any strong assumptions on the nature of substitution.

Taking FOCs, one obtains an optimality condition:[1]

$$\left(\frac{c_A}{c_W}\right)\left(\frac{A^{*\rho-1}}{W^{*\rho-1}}\right) = \frac{p_A}{p_W}$$

Taking logs gives ratio of unconditional factor demands:

$$\ln\left(\frac{W^*}{A^*}\right) = \frac{1}{\rho-1}\ln\left(\frac{c_A}{c_W}\right) + \frac{1}{\rho-1}\ln\left(\frac{p_W}{p_A}\right) \tag{1}$$

If $\rho < 1$ the inputs are substitutes and the coefficient on the price ratio is negative. Treating CAA non-attainment as a decrease in $\frac{p_W}{p_A}$, theory then predicts an increase in the ratio of water to air pollution $\frac{W^*}{A^*}$.[2]

Modeling W^* as a function of prices alone will result in biased estimates because of the omitted variable A^*. Rearranging equation 1 to put A^* on the right hand side makes this apparent.

[1] The use of two inputs here is without loss of generality. In models with more than two inputs, one obtains an analogous optimality condition for each pair.

[2] For a model that treats CAA non-attainment as a limit on the quantity of air emissions, please see Appendix Section 10. The qualitative predications from that model are the same as those presented here.

$$\ln(W^*) = \frac{1}{\rho - 1}\ln\left(\frac{c_A}{c_W}\right) + \frac{1}{\rho - 1}\ln\left(\frac{p_W}{p_A}\right) + \ln(A^*) \qquad (2)$$

In the context of the CAA, suppose a plant is located in a county that falls into non-attainment. The plant has two emissions reduction options: 1) substitute toward another form of pollution W^* (e.g. by switching fuels or using existing pollution-control capital more intensively); 2) produce less output. If the plant does both, the level of W^* may fall even though the ratio $\frac{W^*}{A^*}$ has increased. The output reduction disguises the regulation-induced substitution. Avoiding this confusion requires controlling for A^*, Y^* (output), or both. I model the input ratio, thereby controlling for changes in A^*.

One might worry that this framework will capture a "mechanical" substitution effect. After all, if the CAA causes plants in non-attainment counties to reduce their air emissions and leave water emissions unchanged, the ratio $\frac{W^*}{A^*}$ will increase. But this is actually evidence of substitution, as apparent from figure 1. In the left-hand panel, the price of air emissions rises from p_{A0} to p_{A1}. Holding total cost TC and water emissions fixed, the firm's new input bundle is (W_1, A_1) at lower output Y_1. Water emissions are unchanged (by construction), but air emissions are lower. This change, however, incorporates both output and substitution effects. The right-hand panel removes the output effect by drawing a cost line (in green) at the new prices and the original output level Y_0. The input bundle is now (W_2, A_2), where $W_2 > W_0$. Holding output fixed, water emissions have actually increased.

The preceding discussion assumes a static production technology, with input substitution driven by exogenous price changes. This assumption may be incorrect if firms respond to regulation with both technology changes (e.g. installation of new pollution-control capital) and input substitution. In such a case both $\frac{p_A}{p_W}$ and $\frac{c_A}{c_W}$ may change. Under a CES functional form assumption this does not change the interpretation of my estimates. Factoring equation 1 yields the following.

$$\ln\left(\frac{W^*}{A^*}\right) = \frac{1}{\rho - 1}\left[\ln\left(\frac{c_A}{c_W}\right) + \ln\left(\frac{p_W}{p_A}\right)\right] \qquad (3)$$

Given a proxy for the quantity $\left[\ln\left(\frac{c_A}{c_W}\right)+\ln\left(\frac{p_W}{p_A}\right)\right]$, it is still possible to recover the substitution parameter $\frac{1}{\rho-1}$. If the CES assumption fails, my estimates can no longer be interpreted as the substitution parameter $\frac{1}{\rho-1}$. Instead they will capture the full effect of CAA regulation, which may be a function of multiple underlying parameters.

4 Data

My plant-level emissions and location data come from the EPA Toxic Release Inventory (TRI) 1987-2010. TRI records emissions of more than 500 chemicals by weight (in pounds or grams). TRI data encompass a broad set of industries, from electric power to soybeans. The top ten industries by total TRI-reportable emissions are listed in Table 1.

These data have several shortcomings, discussed in Hamilton (2005). Only large emitters are required to participate.[3] Firms typically report estimates derived from engineering models, rather than direct measurements. Gamper-Rabindran (2006) finds that the location variables are sometimes inaccurate. Under TRI there are penalties for false reporting, but not high emissions, which should ameliorate firm incentives to under-report emissions. The EPA has fined firms up to $27,000 per day for reporting problems in the past (Gamper-Rabindran 2009).

A subset of TRI chemicals are classified as particulates (PM).[4] The TRI data capture emissions in great detail, distinguishing for example between

[3] Reporting thresholds have varied over time and by chemical. Typically firms must report if they use or process more than 10,000 pounds of a TRI-listed chemical per year.

[4] Professor Michael Greenstone generously shared his mapping from TRI chemicals to CAA criteria pollutants. Details are available in Greenstone (2003). These data also include mappings to lead and VOCs, which I do not employ. I do not analyze lead emissions because of the small number of treated plants. The VOC mapping is problematic because VOCs are not directly regulated under the CAA. They are one of two primary precursors (the other is NOx) of ozone, which is a CAA criteria pollutant. While one would expect ozone non-attainment to affect VOC emissions, the link is much less clear than for particulates, as not all VOCs contribute substantially to ozone formation. EPA regulates PM10 (particles <10 microns in diameter) and PM2.5 (<2.5 microns in diameter) separately, but the Greenstone data do not allow me to separately identify these categories. TRI does not include emissions of CO, NO2, or SO2.

different types of underground wells. To simplify presentation and analysis I aggregate up to the categories described in Table 2.

Data on county attainment status come from the EPA Green Book. Monitor-level data on pollutant concentrations come from the EPA Air Quality System (AQS) 1993-2010. My analysis is based on the period of overlap between my TRI and AQS data, 1993-2010. For descriptive statistics see appendix table A1.

5 Estimation

5.1 Estimating equations

To estimate treatment effects on air emissions, I use the following specification, with i indexing plant and t year.

$$\ln(A_{it}) = \overline{\alpha_i} + \overline{\delta_t} + \beta treated_{it} + \varepsilon_{it} \tag{4}$$

The dependent variable is the log of a plant's air emissions. The equation includes plant fixed effects and year dummies, with the latter capturing changes in TRI reporting requirements and secular forces influencing emissions. As discussed in section 5.2 below, the variable $treated_{it}$ equals 1 for plants that were within 2km of a non-attainment monitor in year t-1. If CAA regulations are effective in reducing air emissions, I expect estimates of β to be negative.

In addition, to investigate the time pattern of effects, I estimate an event-study specification.

$$\ln(A_{it}) = \overline{\alpha_i} + \overline{\delta_t} + \sum_j \tau_j + \varepsilon_{ipt} \tag{5}$$

The variables τ_j are indicators for a time index defined relative to treatment. I include dummies for $\tau = 3$, $\tau = -2$, $\tau = -1$, $\tau = 0$, $\tau = 1$, $\tau = 2$, and $\tau = 3$. Tau equals 0 in the first treated year. This means that if a county violates the CAA particulate standard in the year when $\tau = -1$, it enters treatment the following year.

To test for cross-media substitution, I estimate the following.

$$\ln\left(\frac{W_{it}}{A_{it}}\right) = \overline{\alpha_i} + \overline{\delta_t} + \beta treated_{it} + \varepsilon_{it} \tag{6}$$

As before, I include plant fixed effects and year dummies. The quantity $\ln\left(\frac{W_{it}}{A_{it}}\right)$ is the plant's log emissions ratio, with the numerator emissions into another medium (e.g. water or land) and the denominator air emissions. The estimating equation closely parallels the ratio of unconditional factor demands from equation 1 above. The treatment dummy proxies for the unobservable shift in the price ratio $\frac{p_W}{p_A}$ or, alternatively, the combination of price changes and technology changes in $\frac{c_1}{c_2}$. If the CAA induces cross-media substitution, estimates of β will be positive.

To test for within-firm leakage, I estimate the following specification using all plants in attainment counties.

$$\ln(A_{it}) = \overline{\alpha_i} + \overline{\delta_t} + \gamma multiplant_{it} + \beta multiplant_{it} * other\_treated_{it} + \varepsilon_{it} \tag{7}$$

Again I include plant fixed effects and year dummies. The variable $multiplant_{it}$ is a dummy for being part of a multi-plant firm. The variable $other\_treated_{it}$ is a dummy for one or more treated plants within the same firm and 6-digit NAICS code. If the CAA induces spatial leakage, estimates of β will be positive.

5.2 Defining treatment

Past research on cross-media substitution has typically defined treatment as presence in a non-attainment county, but this conceals important spatial heterogeneity. Auffhammer et al (2009) find the effect of county non-attainment status on an average monitor is zero, but the effect on a non-attainment monitor is negative 11 to 14 percent. This suggests that regulators treat plants near non-attainment monitors intensively, while treating plants farther away lightly or not at all. I present evidence in support of this hypothesis. First

I estimate a simple regression of a plant's air emissions on plant fixed effects and year dummies:

$$\ln\left(A_{it}\right) = \overline{\alpha_i} + \overline{\delta_t} + \varepsilon_{it} \tag{8}$$

In this equation A denotes air emissions, while i indexes plant and t year. Figure 2 is a local linear regression fit to plant residuals from non-attainment counties against the distance to the nearest non-attainment monitor. It provides evidence that regulators indeed treat plants near non-attainment monitors intensively, while treating more distant plants lightly or not at all.

Based on this pattern, I define a variable $treated_{it} = Nonattain_{it-1} * 1\left\{Distance_{it-1} \leqslant \overline{D}\right\}$. That is, I consider a plant *treated* in year t if in the prior year its county was in non-attainment and the plant was located "close" to a nonattainment monitor. Based on Figure 2 I use a threshold distance $\overline{D}$ of 2km. I use lagged rather than contemporaneous nonattainment status because status for year t-1 is not known with certainty until the end of the year, so I expect little effect before year t. This treatment variable forms the basis for all subsequent results.[5] This pattern is consistent with a regulator whose objective function involves minimization of enforcement costs, either pecuniary or political (Amacher & Malik 1996). The qualitative evidence presented by Becker and Henderson (2000) on regulator-firm negotiations is also consistent with such an explanation.

5.3 Exogeneity of CAA non-attainment

I cannot recover the causal effects of CAA regulation unless it is exogenous to my plant-level outcomes of interest. Past literature has typically argued in support of this assumption.[6] Chay and Greenstone (2003, 2005) document that PM10 non-attainment counties do not differ systematically from attainment

[5] While this pattern holds on average, it need not hold for all industries and pollutants. Stack height provides one source of heterogeneity. If a plant has tall stacks, it exerts more influence on distant monitors than on those nearby (author's interview notes). In such a case, even if regulators focus on particular plants, they may not be the plants adjacent to non-attainment monitors.

[6] Examples include Henderson (1996), Becker and Henderson (2000), Greenstone (2002), Auffhammer et al (2011), and Walker (2011).

counties on observable dimensions (including economic shocks), either in levels or in changes.

Non-attainment is plausibly exogenous if a given firm produces a small portion of the ambient air pollution in a county. For the *average* plant in a non-attainment county, this is a tenable assumption. Motor vehicles typically account for the majority of PM pollution, especially in urban areas. The California Air Resources Board estimates that 74 percent of PM10 emissions come from non-point sources like road dust and from residential fuel combustion (Auffhammer et al 2011).

The spatial heterogeneity documented in Section 5.2, however, calls into question the exogeneity of CAA regulation for *treated* plants (plants actually affected by regulation). CAA regulations primarily affect plants within two kilometers of a non-attainment monitor. It might be that past emissions by a given plant were pivotal in pushing its county above the CAA standard. If that were the case, CAA regulation would be endogenous to past emissions by treated plants. For example, if a plant experienced particularly strong demand for its output in a given year, it might have emitted more air pollution than usual and pushed the nearby monitor above the CAA standard.

This potential problem provides additional motivation for my use of emissions ratios, rather than emissions levels, in my analysis of cross-media substitution. If the endogeneity of nonattainment with respect to past emissions stems from output shocks, then treatment should remain exogenous to emissions ratios. It is still possible, however, that a plant might push its county into nonattainment because of shocks to emissions ratios. This form of endogeneity is perhaps less plausible, but impossible to exclude in principle. For example, a plant's scrubber might fail in a given year, increasing its ratio of air to water emissions and pushing its county into nonattainment. The sign of the bias in such a case would depend on the autocorrelation in the shocks to emissions ratios. Endogenous past output could also bias my estimates of CAA treatment effects on the level of air emissions. For example, if output shocks were negatively autocorrelated, my estimates might overstate CAA treatment effects. If instead output shocks were positively autocorrelated, it might understate them.

Figure 3 investigates the possibility of endogenous entry into treatment using an event study framework (estimates from equation 5). I define a new time index τ relative to treatment. A county violates the CAA in year $\tau = -1$ and plants within 2km of a non-attainment monitor enter treatment in the following year ($\tau = 0$). If the figure showed either higher air emissions or a lower ratio of other emissions to air emissions at $\tau = -1$, that would be evidence of endogenous entry into treatment. Instead the figure shows the opposite pattern. Air emissions fall in the final pre-treatment year and the ratio of other emissions to air emissions increases. I attribute this pattern to firm expectations of non-attainment. There are two particulate standards, one based on the annual average at a monitor and another based on the 98th percentile 24-hour mean. If a county violates the 24-hour standard in year $\tau = -1$, firms might plausibly learn about it before the official non-attainment designation at the end of the year. A firm might also be able to anticipate its county violating the annual standard. Any such anticipatory behavior by firms will bias the magnitudes of my estimates downward. Note however that for many plants I observe long pre- and post-treatment periods. By estimating in levels with 18 years of data, I partially mitigate this bias.

Finally, the importance of location in my analysis raises the potential for endogeneity springing from monitor placement. The EPA does not place monitors randomly. Indeed EPA rules require monitors, for example, near large sensitive populations (e.g. asthmatic children; Raffuse et al 2007). For monitors present throughout the study period, plant fixed effects should remove any potential bias from endogenous placement. EPA does sometimes relocate monitors and introduce new ones over time, however. This creates potential endogeneity between monitor distance and within-plant variation in emissions if the EPA's new-monitor placement decisions depend on: 1) plant-level scope for air emissions abatement; 2) plant-level scope for cross-media substitution; or 3) plant-level scope for shifting output to other plants within the same firm. Of these three potential problems, (1) is the most plausible. Suppose EPA is more likely to place monitors near high-emitting plants. By virtue of their high emissions levels, these plants may have more scope for air emissions abatement. I address this potential problem by analyzing proportional changes in air emissions.

6 Results

6.1 Air emissions

Table 3 presents my estimate of the CAA treatment effect on airborne particulate emissions. Treated plants decrease their air emissions by 25 percent. This is larger than the 11 to 14 percent effect on non-attainment monitors reported by Auffhammer et al (2011) because: 1) plant emissions become diluted as they mix with surrounding air; and 2) the treated plants in my sample are not the only factor influencing ambient air pollution. Column 2 adds county linear time trends. This reduces the magnitude of the estimate modestly, from 25 to 20 percent, but it remains statistically significant at the five percent level.

Column 3 presents the results from an event-study specification (equation 5). The time pattern suggests that most of the emissions reductions occur when τ is 0 or 1: the year prior to a non-attainment designation and the first year following a non-attainment designation. (For discussion of the pre-treatment decline, see section 5.2.) This motivates my use of fixed-effects models in levels elsewhere. Estimates based on changes in treatment status would be biased toward zero because of the emissions decline at $\tau = -1$. With a relatively long pre-treatment period, however, a model in levels averages this pre-treatment decline with other untreated years, partially mitigating the bias. At approximately -36 percent, the event-study estimates are close in magnitude to my primary result (-25 percent). Together these results suggest that treated plants do indeed reduce airborne particulate emissions.

6.2 Cross-media substitution, all industries

Panel A in table 4 shows estimated treatment effects from equation 6, by medium across all industries. The dependent variable is a log emissions ratio, with emissions into a given medium (indicated in the column heading) in the numerator, and air emissions in the the denominator. Positive estimates are evidence of cross-media substitution. Point estimates for treatment are uniformly positive and large in magnitude, especially for on-site forms of disposal. There is evidence of statistically significant substitution toward onsite water pollution, recycling, and "onsite other" emissions. "Onsite other" emissions include waste piles, leaks, and spills.

Panel B in table 4 adds county linear time trends to my model of cross-media substitution. The estimates for onsite water (45 percent) and recycling (44 percent) are essentially unchanged from my primary results in panel A. The estimate for onsite other, however, falls approximately by half and loses statistical significance.

The large increase in recycling highlights the fact that not all substitution responses reduce social welfare. The increased water emissions, however, impose social costs. The magnitude of those costs is difficult to quantify, given the relative scarcity of well-identified studies on the health and productivity effects of water pollution. Spills and waste piles similarly impose social costs that are difficult to estimate. On a net basis, cross-media substitution need not reduce welfare. Suppose a plant responds to non-attainment by reducing output and substituting toward water emissions. Gross water emissions may end up below their initial level. In such a case the CAA may still improve welfare, but substitution attenuates the gains.

6.3 Cross-media substitution, by industry

It is difficult to analyze substitution patterns at the industry level due to the small number of treated plants: recall that not all plants in non-attainment counties are treated. Moreover not all plants report emissions into all media. Nonetheless, to illustrate the heterogeneity in substitution responses, Table 6 presents estimates for the three industries with the largest treated sample sizes: primary metals, wood products, and utilities. (Appendix table A6 presents more disaggregated estimates at the 3-digit NAICS level.) Estimates again come from equation 6. In the discussion that follows, note that I cannot reject the null hypothesis of equal coefficients in most cases; the evidence of heterogeneity is merely suggestive. Wood products and utilities show large decreases in air emissions, 48 and 36 percent respectively, while primary metals shows only a 12 percent decrease. Similarly, wood products and utilities increase their ratios of water to air emissions by 71 and 118 percent, while primary metals increase this ratio by only 30 percent. These two industries also increase their use of waste piles by more than 200 percent. Only utilities increase their ratio of offsite land disposal to air emissions, but the effect is very large at 174 percent. Both wood products and utilities increase their ratio of recycling to air emissions by much more than does primary metals, again suggesting the former two industries are more intensively regulated.

6.4 Leakage

Intuition predicts that firms might respond to treatment of a plant in one county by shifting output to a plant in another county. Table 8 provides evidence they do so. Estimates correspond to equation 7. For the average plant in an attainment county, treatment of another plant within the same firm and 6-digit NAICS code increases air emissions by 17 percent. Column (2) adds county linear time trends and the estimate is slightly smaller at 15 percent. Treating the number of other treated plants as a continuous variable (column 3), estimated leakage is 12 percent per treated plant. With the addition of county linear time trends in column 4, the estimate is again slightly smaller at 10 percent, and no longer statistically significant. This leakage has associated health, mortality, and productivity costs. As a robustness check, I estimate the same model grouping plants by firm and 2-digit NAICS code and report results in table A4. Estimates are modestly smaller than in my preferred specification, though still positive and significant. This is reasonable, as the coarser classification groups plants that may not be close substitutes for each other.

The average treated plant in my data is part of a firm with approximately

3 plants that are candidates for leakage-driven increases: they share the same six-digit NAICS code and are located in attainment counties.[7] Average air emissions at eventually treated plants prior to treatment are 4768 pounds, while average baseline emissions at candidates are 2753 pounds. The estimates from table 8 imply the following net change in emissions from treating an average plant. The treated plant reduces emissions by $.25 * 1454 = 1192$ pounds. The 3 candidate plants together increase emissions by $3 * .17 * 2753 = 1404$ pounds. On net, then, CAA treatment of an average plant increases particulate emissions by $1404 - 1192 = 212$ pounds. This result should be interpreted with several important caveats in mind. First, the TRI data cover only large plants, which may be more likely to belong to multi-plant firms and thus may have more scope for within-firm leakage. Second, these estimates describe only TRI-reportable particulate emissions. Third, leakage patterns might differ for other CAA-regulated pollutants (e.g. SO2). Fourth, industrial sources account for approximately 25 percent of particulate emissions in an average county (Auffhammer et al 2011), so the implied changes in ambient pollution are much smaller than the emissions changes I estimate at the plant level.

Leakage reduces the welfare gains from CAA regulation, but need not imply a net welfare loss. Leakage-driven emissions increases occur in attainment counties, which by definition have lower ambient air pollution. In addition, the average attainment county population is approximately $1/3$ of the average non-attainment county population (author's calculation). Particularly if the social damage function for air pollution is convex, the net welfare effect from CAA treatment of the plants in my data may be positive. Leakage does present a potential problem in using difference-in-differences designs to evaluate the CAA, as it is a spillover from the treatment group (typically non-attainment counties) to the control group (attainment counties). The spillovers identified in table 8 imply that such analyses overstate CAA benefits in non-attainment counties and fail to account for some of the costs in attainment counties. This provides additional motivation for my use of an emissions ratio, rather than

[7] This includes plants that are not part of a multi-plant firm.

a level, in my analysis of cross-media substitution. To test whether spillovers influence my cross-media results, I estimate my leakage model using an emissions ratio as the dependent variable and report results in appendix table A5. Estimates are generally near zero and statistically insignificant, with one exception: the estimate for "offsite other" is negative 76 percent. This spillover will tend to bias my cross-media model toward finding evidence of substitution toward the "offsite other" channel. It is also possible that leakage causes my treatment model to overestimate the air emissions reductions undertaken by treated plants. To evaluate this possibility, I estimate a variant of my air emissions model (equation 4), controlling for spillovers as in equation 7. Reported in table A3, the estimates are unchanged.[8]

7 Additional robustness & placebos

7.1 Cross-media substitution

Table 9 panel A moves from a ratio specification based on equation 1 to a more flexible specification based on equation 2, with log air emissions on the right-hand side. This is not my preferred specification because air emissions are endogenous and likely bias the estimates of the treatment effect. Nonetheless I include this specification to test the importance of the CES functional form assumptions that affect interpretation of my primary results. The broad pattern of results is unchanged, with large estimates for on-site media and recycling, but smaller estimates for off-site media. Substitution toward onsite other and recycling remains statistically significant. While estimated substitution toward water remains large and positive at 25 percent, it is no longer statistically significant.

Panel B shows estimates from a specification without any control for air emissions. Most estimates remain large and positive. There are statistically

[8]The potential bias in my main specification would come from the influence of the spillover plants on the estimates for year dummies. The failure to control for spillovers has no practical import because only a small number of plants are affected by spillovers. Identification of the year dummies comes primarily from plants that do not receive spillovers.

significant increases in “onsite other” releases (e.g. waste piles) and recycling. The close agreement between panel A and panel B provides indirect evidence that false reporting is not a major problem in my setting. A treated firm might conceivably have weak incentives to under-report air emissions, even though TRI data are not used for CAA enforcement. It has no incentives to misreport non-air emissions. The agreement between panels A and B shows that the inclusion of an air control has minimal influence on estimate magnitudes in these alternative specifications. This suggests that reporting problems are not driving my results.

7.2 Placebos

In partial equilibrium, treatment should have no effect on plants that do not emit any air pollution. Table 11 tests this hypothesis by estimating a variant of equation 6 with two changes: 1) the dependent variable is log emissions into a given medium (the lack of air emissions precludes a ratio); and 2) treatment is interacted with a dummy indicating zero air emissions. If my model is well specified, it should find no effect of CAA regulation on these plants. The estimates are indeed insignificant, with the exception of the one for onsite land emissions. While the latter is statistically significant, it is negative. If there were some omitted variable decreasing land emissions at plants near non-attainment monitors, it would work against finding cross-media substitution (it would bias my primary estimates downward).

Table 12 reports results from a placebo test of my leakage model. I construct variables based on placebo “treated” plants: plants within the same firm and 6-digit NAICS code that are located in non-attainment counties, but farther than eight kilometers from the nearest non-attainment monitor. As these plants are not treated, we should not see increased air emissions by attainment-county plants in the same firm and NAICS code. If my leakage model is capturing, for example, changes in the geographic distribution of output that happen to be correlated with treatment, this placebo test should return large positive estimates. Instead the estimates in Table 12 are in the one to two percent range and are not statistically significant. This suggests that the leakage results in table 8 do not spring from an omitted variable problem.

8 Conclusion

The Clean Air Act is the most prominent environmental regulation in the United States. While economists have long recognized the potential for substitution responses to single-medium pollution regulation, empirical studies have found little evidence of such effects. The paucity of available data and the difficulty of controlling for scale have made firm responses hard to detect.

Using specifications motivated by classical firm optimization theory, this study provides evidence of regulation-induced pollution substitution in response to the CAA. Estimates from 18 years of EPA Toxic Release Inventory data show that CAA-regulated plants increase their ratio of water to air emissions by 42 percent. In multi-plant firms, particulate regulation of an average plant increases air emissions at unregulated plants by 17 percent. This results in a net emissions increase. Responses of this magnitude plausibly have social costs and should be considered in policy design. The welfare effects of such substitution present an interesting subject for future research. Additionally, I document spatial heterogeneity in regulatory intensity, which suggests that regulators seek to minimize costs (political or pecuniary) in implementing the CAA.

These findings might helpfully inform the design of future pollution control policy. A theoretically optimal policy, with emissions into every medium and location priced according to marginal damage, would be difficult to achieve. But policy could move toward the first-best simply by adjusting prices according to county population and ensuring that water and land emissions are priced approximately at marginal damage.

Such improvements in policy design would have economically significant consequences. There is growing evidence that air pollution has costly long-run effects on exposed individuals. Isen et al (2014), for example, finds that in-utero and early childhood air pollution exposure depresses earnings for workers ages 29-31. Given these costs, the returns to improved pollution regulation may be large.

9 Figures and tables

Figure 1: Pollution changes, holding output fixed

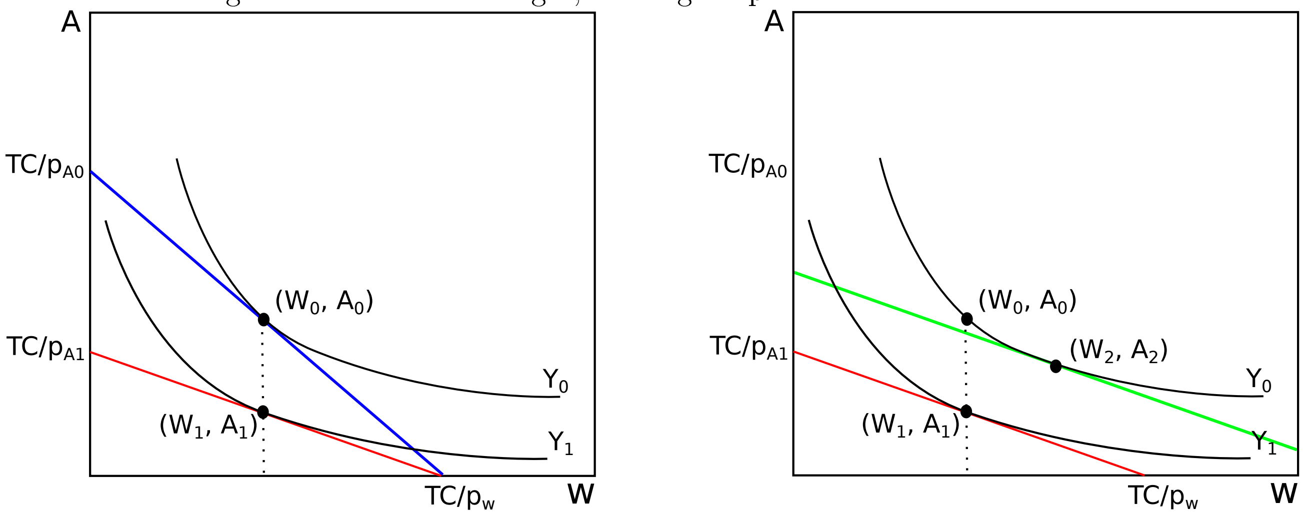

Table 1: Top ten industries, by TRI-reportable emissions

Rank	NAICS code	Industry
1	221112	Fossil electric power
2	325188	Inorganic chemicals
3	212231	Pb & Zn mining
4	212234	Cu & Ni mining
5	212221	Au mining
6	331111	Iron & steel
7	325199	Organic chemicals
8	322121	Paper
9	562211	Hazardous waste
10	324110	Petroleum Refining

Table 2: Aggregated TRI emissions categories

Aggregated category	Included TRI components
Onsite air	Fugitive air, stack air
Onsite water	Onsite water
Onsite land	Landfills, impoundment ponds, underground wells
Onsite other	Waste piles, leaks, spills
Offsite water	Public/private water treatment
Offsite land	Landfills, impoundment ponds, underground wells
Offsite other	Residual emissions*, waste brokers, incinerators and storage facilities
Recycled or treated	Recycled, recovered, treated

Figure 2: Residual air emissions by distance from nearest non-attainment monitor

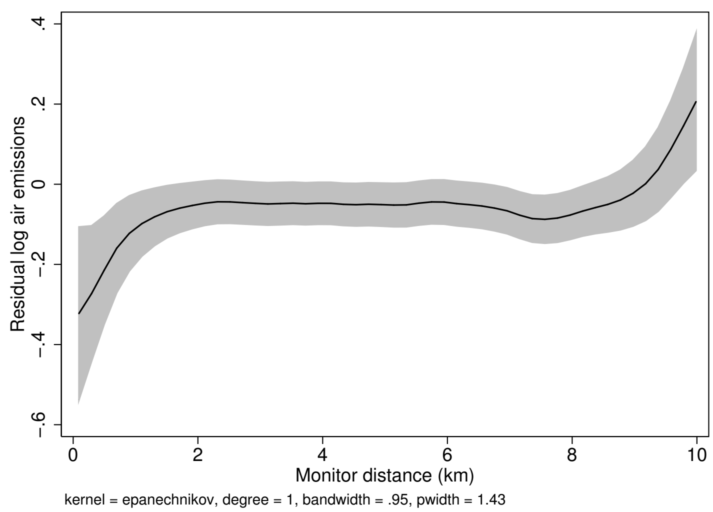

Underlying residuals from equation 8, a panel model with year dummies and plant fixed effects. Fitted line from a local linear regression. Shaded area is the 95% confidence interval.

Figure 3: Event study estimates

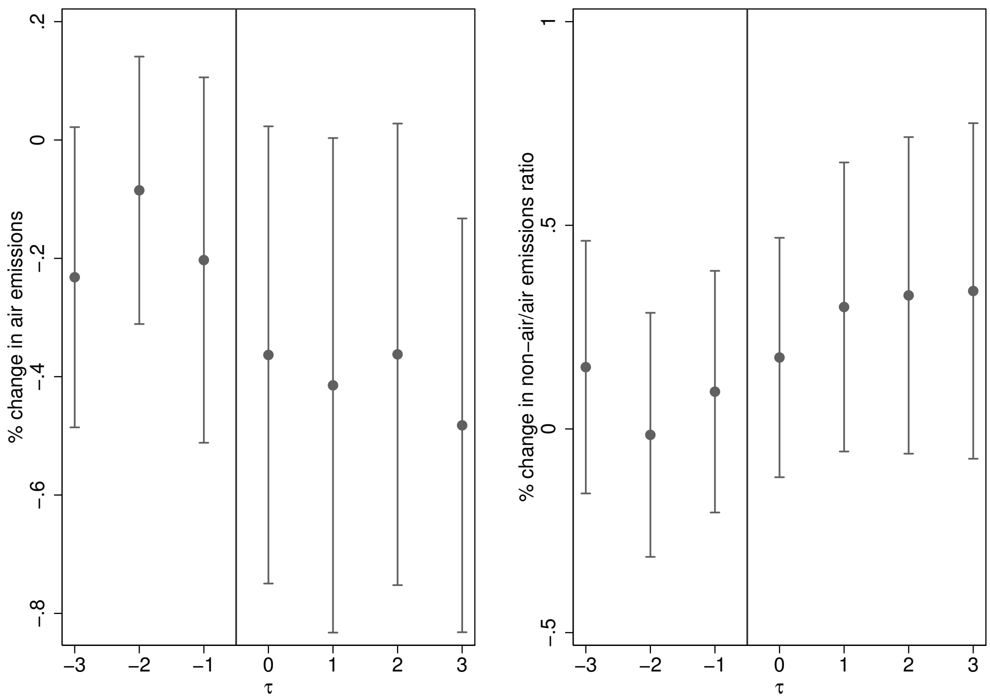

Estimates from equation 5, with point estimates reported in column 3 of table 3. A county violates the CAA in year tau=-1 and plants within 2km of a non-attainment monitor enter treatment in the following year (tau=0). Dependent variable is log air emissions. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table 3: Effect on air emissions

	(1) Onsite air	(2) Onsite air	(3) Onsite air
Treated	-0.251**	-0.201**	
	(0.111)	(0.0990)	
Tau=-3			-0.232*
			(0.129)
Tau=-2			-0.0851
			(0.115)
Tau=-1			-0.203
			(0.158)
Tau=0 (1st treated year)			-0.363*
			(0.197)
Tau=1			-0.415*
			(0.213)
Tau=2			-0.362*
			(0.199)
Tau=3			-0.482***
			(0.178)
County linear trends	No	Yes	No
Year dummies	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes
Observations	123918	123918	150808

Standard errors in parentheses

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

Estimates in columns 1-2 correspond to equation 4, while estimates in column 3 correspond to equation 5. Dependent variable is log air emissions. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table 4: Effect on emissions ratios

Panel A: Main specification							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Treated	0.420**	0.614	1.348***	0.180	0.143	0.177	0.468***
	(0.190)	(0.775)	(0.382)	(0.177)	(0.125)	(0.288)	(0.146)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	30035	14417	6128	35678	51052	27736	56023
Panel B: County linear trends							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Treated	0.450**	0.726	0.608	0.177	0.0562	0.136	0.435***
	(0.176)	(0.776)	(0.844)	(0.137)	(0.146)	(0.289)	(0.136)
County linear trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	30035	14417	6128	35678	51052	27736	56023

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

Estimates correspond to equation 6. Dependent variable is log emissions ratio, with the numerator indicated atop the column and the denominator air emissions in all columns. Specification includes year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media. “Onsite other” emissions include waste piles, leaks, and spills.

Table 6: Effect on emissions ratios, by 2-digit NAICS code

	(1) Onsite air	(2) Onsite water	(3) Onsite land	(4) Onsite other	(5) Offsite water	(6) Offsite land	(7) Offsite other	(8) Recycled or treated
Primary metals	-0.116	0.301	0.131	2.050***	0.269	0.00572	0.167	0.259
	(0.156)	(0.231)	(0.795)	(0.393)	(0.234)	(0.177)	(0.356)	(0.175)
Observations	64715	13955	3121	2030	25711	28861	16755	41122
Wood products	-0.483***	0.705*	2.344	2.158***	0.111	0.175	0.176	1.048**
	(0.153)	(0.389)	(2.064)	(0.509)	(0.222)	(0.180)	(0.412)	(0.410)
Observations	44858	11234	6668	1945	8137	17181	8106	11889
Utilities	-0.355	1.180	-0.998	-0.159		1.742*	2.036	11.10***
	(0.426)	(0.893)	(1.140)	(0.593)		(0.987)	(1.457)	(0.286)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5410	3213	3319	838	368	2445	1452	1188

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

Includes the three 2-digit NAICS industries with the largest treated sample sizes. Column 1 (onsite air) corresponds to equation 4, remaining columns to equation 6. Dependent variable is log air emissions in column 1, otherwise log emissions ratio, with the numerator indicated atop the column and the denominator air emissions in all columns. All specifications include year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media. "Onsite other" emissions include waste piles, leaks, and spills.

Table 8: Leakage effect, within firm & 6-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other treated plants	0.172**	0.146*		
	(0.0773)	(0.0775)		
Count other treated			0.124**	0.102
			(0.0627)	(0.0632)
County linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Multiplant dummy	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	111902	111902	111902	111902

Standard errors in parentheses

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

Estimates correspond to equation 7, where "other treated plant" is a treated plant within the same firm and 6-digit NAICS code. Dependent variable is log air emissions. Specification includes year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Sample restricted to plants in attainment counties.

Table 9: Effect on emissions ratios, alternative specifications

Panel A: Air emissions on RHS	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Treated	0.247	0.617	1.002***	-0.0615	0.0158	0.128	0.228**
	(0.188)	(0.691)	(0.340)	(0.143)	(0.140)	(0.219)	(0.0889)
Log air emissions	0.188***	0.264***	0.259***	0.199***	0.217***	0.160***	0.141***
	(0.0127)	(0.0307)	(0.0521)	(0.0115)	(0.0145)	(0.0186)	(0.0119)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	30035	14417	6128	35678	51052	27736	56023

Panel C: No air control	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Treated	0.207	0.619	0.880**	-0.122	-0.0193	0.118	0.189**
	(0.194)	(0.666)	(0.352)	(0.158)	(0.156)	(0.211)	(0.0876)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	30035	14417	6128	35678	51052	27736	56023

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

Panel A estimates correspond to equation 6, but with log air emissions moved to the right-hand side of the equation. Note that air emissions are endogenous, so this is not my preferred specification. Panel B removes the control for air emissions from the right hand side of the model. Dependent variable is log emissions, with the medium indicated atop the column and the denominator air emissions in all columns. All specifications include year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media. "Onsite other" emissions include waste piles, leaks, and spills.

Table 11: Placebo effect on emissions levels

	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
Treated*no air emissions	-0.647	-0.761***	0.171	0.243	0.0984	-0.774	-0.127
	(0.454)	(0.0474)	(0.139)	(0.376)	(0.294)	(0.581)	(0.162)
Treated*air emissions	0.273	0.620	0.877***	-0.0544	0.0724	0.144	0.172*
	(0.192)	(0.666)	(0.323)	(0.185)	(0.154)	(0.200)	(0.0899)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	34544	16393	7307	51753	69004	39451	85433

Standard errors in parentheses

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

Estimates correspond to equation 6, but with 2 changes: 1) the dependent variable is log emissions (not a ratio); and 2) estimates for "Treated*no air emissions" report the effect of placebo treatment (being near a non-attainment monitor) on plants with no air emissions, which should not be affected by the CAA. Estimates for "Treated*air emissions" are for actually treated plants; they are not placebos. The medium is indicated atop the column. All specifications include year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media.

Table 12: Placebo leakage effect, within firm & 6-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other placebo plants	0.0178	0.00797		
	(0.0491)	(0.0494)		
Count placebo plants			0.0193	0.0104
			(0.0290)	(0.0299)
County linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Multiplant dummy	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	128147	128147	128147	128147

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Estimates correspond to equation 7, but using variables based on placebo treated plants: plants within the same firm and 6-digit NAICS code, located in non-attainment counties, but farther than 8km from the nearest non-attainment monitor. Dependent variable is log air emissions. Specification includes year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Sample restricted to plants in attainment counties.

References

[1] Amacher, G. S., & Malik, A. S. "Bargaining in Environmental Regulation and the Ideal Regulator." *Journal of Environmental Economics and Management* (1996), 233–253.

[2] Auffhammer, Maximilian, Antonio M. Bento, and Scott E. Lowe. "Measuring the effects of the Clean Air Act Amendments on ambient PM10 concentrations: The critical importance of a spatially disaggregated analysis." *Journal of Environmental Economics and Management* 58.1 (2009): 15-26.

[3] Auffhammer, Maximilian, Antonio M. Bento, and Scott E. Lowe. "The City-Level Effects of the 1990 Clean Air Act Amendments." *Land Economics* 87.1 (2011): 1-18.

[4] Becker, Randy, and Vernon Henderson. "Effects of air quality regulations on polluting industries." *Journal of Political Economy* 108.2 (2000): 379-421.

[5] Becker, Randy A., and J. Vernon Henderson. "Costs of air quality regulation." *Behavioral and distributional effects of environmental policy.* University of Chicago Press, 2001. 159-186.

[6] Bushnell, J. B., & Mansur, E. T. "Vertical Targeting and Leakage in Carbon Policy." *American Economic Review Papers & Proceedings,* 101:3 (2011): 263–267.

[7] Campbell, Harry Fleming. "Estimating the elasticity of substitution between restricted and unrestricted inputs in a regulated fishery: a probit approach." *Journal of Environmental Economics and Management* 20.3 (1991): 262-274.

[8] Chay, Kenneth Y., and Michael Greenstone. "Air quality, infant mortality, and the Clean Air Act of 1970." No. w10053. National Bureau of Economic Research (2003).

[9] Chay, Kenneth Y., and Michael Greenstone. "The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession." T*he Quarterly Journal of Economics* 118.3 (2003): 1121-1167.

[10] Chay, Kenneth Y., and Michael Greenstone. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy* 113.2 (2005).

[11] Davis, L. W., & Kahn, M. E. "International Trade in Used Vehicles: The Environmental Consequences of NAFTA." *American Economic Journal: Economic Policy* 2:4 (2010), 58–82.

[12] Department of Energy. "Cleaning Up Coal." http://www.fossil.energy.gov/education/energylessons/coal/coal_cct2.html. Accessed January 16, 2014.

[13] Environment News Service (ENS), "Court Orders EPA to Impose Power Plant Water Pollution Rule." April 23, 2013.

[14] Environmental Protection Agency. "Reducing Toxic Air Pollutants." http://www.epa.gov/airquality/peg_caa/toxics.html. Accessed December 3, 2013.

[15] Environmental Protection Agency. "Land Disposal Restrictions for Hazardous Wastes." EPA530-F-99-043 (1999).

[16] Environmental Protection Agency. "Air Pollution Control Technology Fact Sheet." EPA -452/F-03-032 (undated).

[17] Environmental Protection Agency. "Hazardous Air Pollutant Emissions for Miscellaneous Organic Chemical Manufacturing: Supplementary Information Document for Proposed Standards." Emission Standards Division (2001).

[18] Environmental Protection Agency. "Clean Water Act Jurisdiction Following the U.S. Supreme Court's Decision in Rapanos v. United States & Carabell v. United States." (2008).

[19] Environmental Protection Agency. "Disposal of Coal Combustion Residuals From Electric Utilities; Proposed Rule." *Federal Register* 75. 118 (2010).

[20] Environmental Protection Agency. "The Benefits and Costs of the Clean Air Act from 1990 to 2020 Final Report – Rev . A." 2011.

[21] Environmental Protection Agency. “Toxics Release Inventory Basic Plus Data File Format Documentation v10.” February 7, 2012.

[22] Environmental Protection Agency. “CWA Methods.” http://water.epa.gov/scitech/methods/cwa/pollutants.cfm. Accessed June 8, 2013.

[23] Environmental Protection Agency. “Electrostatic Precipitator Operation.” EPA 2.0-2/98.

[24] Farnsworth, David. “Preparing for EPA Regulations: Working to Ensure Reliable and Affordable Environmental Compliance.” Regulatory Assistance Project, July 2011.

[25] Fowlie, Meredith L. "Incomplete environmental regulation, imperfect competition, and emissions leakage." *American Economic Journal: Economic Policy* (2009): 72-112.

[26] Gamper-Rabindran, Shanti. "Did the EPA's voluntary industrial toxics program reduce emissions? A GIS analysis of distributional impacts and by-media analysis of substitution." *Journal of Environmental Economics and Management* 52.1 (2006): 391-410.

[27] Gamper-Rabindran, Shanti. "The Clean Air Act and volatile organic compounds: Did plants reduce their health-indexed air emissions or shift their emissions into other media?." Mimeo (2009).

[28] Graff Zivin, Joshua, and Matthew Neidell. "The impact of pollution on worker productivity." *American Economic Review* 102 (2012): 3652-3673.

[29] Greenstone, Michael. "The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures." *Journal of Political Economy* 110.6 (2002): 1175-1219.

[30] Greenstone, Michael. "Estimating regulation-induced substitution: The effect of the Clean Air Act on water and ground pollution." *The American Economic Review* 93.2 (2003): 442-448.

[31] Hamilton, J.T. "Regulation Through Revelation: The Origin and Impacts of the Toxics Release Inventory Program." New York: Cambridge University Press (2005).

[32] Henderson, J. V. "Effects of air quality regulation." *American Economic Review* 86.4 (1996).

[33] Isen, A., Walker, W. R., & Rossin-Slater, M. "Every Breath You Take – Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970." NBER Working Paper No. 19858, January 2013.

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

[35] Raffuse, S., Sullivan, D., McCarthy, M., Penfold, B., & Hafner, H. "Analytical Techniques for Technical Assessments of Ambient Air Monitoring Networks." U.S. Environmental Protection Agency, 2007.

[36] Sanders, Nicholas J. "What doesn't kill you makes you weaker prenatal pollution exposure and educational outcomes." *Journal of Human Resources* 47.3 (2012): 826-850.

[37] Sigman, Hilary. "Cross-media pollution: Responses to restrictions on chlorinated solvent releases." *Land Economics* (1996): 298-312.

[38] Smith, V. Kerry, et al. "General equilibrium benefits for environmental improvements: projected ozone reductions under EPA's Prospective Analysis for the Los Angeles air basin." *Journal of Environmental Economics and Management* 47.3 (2004): 559-584.

[39] Vatavuk, William; van der Vaart, Donald; and Spivey, James. "VOC Controls." EPA/452/B-02-001 (2000).

[40] Walker, W. Reed. "Environmental Regulation and Labor Reallocation: Evidence from the Clean Air Act." *The American Economic Review* 101.3 (2011): 442-447.

[41] Walker, W. Reed. "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce." *The Quarterly Journal of Economics* (2013), 1787–1835.

Appendix for online publication

10 Modeling the CAA as a quantity restriction

Suppose two pollution inputs: A ~ air emissions, W ~ water emissions. Treat the CAA as an exogenous quantity restriction $\overline{A}$ on air emissions. The object of policy interest is unconditional factor demand W^*, incorporating firms' possible output response to regulation. Suppose a CES production function, so the firm problem becomes:

$$\max_{A,W} p_o \left(c_1 A^\rho + c_2 W^\rho\right)^{1/\rho} - p_A A - p_w W + \lambda\left[\overline{A} - A\right]$$

Taking FOCs, one obtains an optimality condition:

$$\left(\frac{c_1}{c_2}\right)\left(\frac{A^{*\rho-1}}{W^{*\rho-1}}\right) = \frac{p_A + \lambda}{p_W}$$

If the constraint does not bind prior to CAA non-attainment, the shadow price λ is zero. Taking logs gives ratio of unconditional factor demands:

$$\ln\left(\frac{W^*}{A^*}\right) = \frac{1}{\rho-1}\ln\left(\frac{c_1}{c_2}\right) + \frac{1}{\rho-1}\ln\left(\frac{p_W}{p_A+0}\right) \tag{9}$$

Treat CAA non-attainment as a decrease in $\overline{A}$ such that it binds. This changes the value of λ from zero to an unknown positive number. The optimality condition then becomes:

$$\ln\left(\frac{W^*}{\overline{A}}\right) = \frac{1}{\rho-1}\ln\left(\frac{c_1}{c_2}\right) + \frac{1}{\rho-1}\ln\left(\frac{p_W}{p_A+\lambda}\right) \tag{10}$$

If $\rho < 1$, then the coefficient on the last term is negative. The positive shadow price λ causes a decrease in the last term. Theory then predicts an increase in the ratio of water to air pollution $\frac{W^*}{\overline{A}}$. This prediction is the same as the one from the model treating CAA non-attainment as a relative price change. The crucial difference is that under this model, a regression that fails to control for output will not produce biased estimates if $\overline{A}$ is truly exogenous. Rearranging equation 10 yields:

$$\ln(W^*) = \frac{1}{\rho-1}\ln\left(\frac{c_1}{c_2}\right) + \frac{1}{\rho-1}\ln\left(\frac{p_W}{p_A+\lambda}\right) - \ln\left(\overline{A}\right) \tag{11}$$

If regulators consider plant characteristics when deciding on the constraint $\overline{A}$, however, the potential for bias in a non-ratio specification returns.

11 Additional tables

Table A1: TRI PM descriptive statistics

	Mean	Stdev	Min	Max
Onsite air	6327.85	570243.85	0.00	1.10e+08
Onsite water	633.85	12441.29	0.00	3361865.00
Onsite land	34801.68	936337.82	0.00	1.10e+08
Offsite other	26081.85	1658407.06	0.00	2.53e+08
Offsite water	534.15	27202.02	0.00	6063868.00
Offsite land	11437.75	128437.66	0.00	12870510.00
Offsite other	3629.44	60310.04	0.00	4371760.00
Recycled or treated	63509.98	668200.67	0.00	1.27e+08
Dist. to nonattain monitor (km)	0.90	4.40	0.00	99.74
PM nonattainment	0.08	0.27	0.00	1.00
Treated	0.01	0.09	0.00	1.00
Observations	197717			

Emissions measured in pounds. Unit of observation is a plant-year.

Table A2: Historical CAA particulate standards

Final rule	Type	Averaging time	Standard (µg/m3)	Form
1987	PM10	24hr	150	Not to be exceeded more than once per year on average over a 3-year period
		Annual	50	Annual arithmetic mean, averaged over 3 years
1997	PM2.5	24hr	65	98th percentile, averaged over 3 years
		Annual	15	Annual arithmetic mean, averaged over 3 years
	PM10	24hr	150	Not to be exceeded more than once per year on average over a 3-year period
		Annual	50	Annual arithmetic mean, averaged over 3 years
2006	PM2.5	24hr	35	98th percentile, averaged over 3 years
		Annual	15	Annual arithmetic mean, averaged over 3 years
	PM10	24hr	150	Not to be exceeded more than once per year on average over a 3-year period

Adapted from http://www.epa.gov/ttn/naaqs/standards/pm/s_pm_history.html. Accessed March 19, 2014.

Table A3: Effect on air emissions, spillover controls

	(1) Onsite air	(2) Onsite air	(3) Onsite air
Treated	-0.256**	-0.205**	
	(0.110)	(0.0989)	
Tau=-3			-0.150
			(0.157)
Tau=-2			-0.0574
			(0.125)
Tau=-1			-0.189
			(0.161)
Tau=0 (1st treated year)			-0.341*
			(0.196)
Tau=1			-0.409*
			(0.220)
Tau=2			-0.326
			(0.209)
Tau=3			-0.335**
			(0.164)
Spillover controls	Yes	Yes	Yes
County linear trends	No	Yes	No
Year dummies	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes
Observations	123918	123918	123918

Standard errors in parentheses

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

Estimates in columns 1-2 correspond to equation 4, while estimates in column 3 correspond to equation 5, but with the inclusion of spillover controls from equation 7. Dependent variable is log air emissions. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table A4: Leakage effect, within firm & 2-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other treated plants	0.125**	0.101**		
	(0.0511)	(0.0495)		
Count other treated			0.0705**	0.0581*
			(0.0356)	(0.0334)
County linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Multiplant dummy	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	111902	111902	111902	111902

Standard errors in parentheses

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

Estimates correspond to equation 7, where "other treated plant" is a treated plant within the same firm and 2-digit NAICS code. Dependent variable is log air emissions. Specification includes year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Sample restricted to plants in attainment counties.

Table A5: Leakage effect on emissions ratios, within firm & 6-digit NAICS code

	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
1+ other treated plants	0.0362	0.0354	-0.0578	-0.144	0.226	-0.760***	-0.0306
	(0.209)	(0.181)	(0.260)	(0.187)	(0.142)	(0.221)	(0.149)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27725	13669	5613	31566	46511	24426	50099

Standard errors in parentheses

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

Estimates correspond to equation 7, where "other treated plant" is a treated plant within the same firm and 6-digit NAICS code, but dependent variable is log emissions ratio. Numerator indicated atop column and denominator is air emissions in all columns. Specification includes year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media. "Onsite other" emissions include waste piles, leaks, and spills. Sample restricted to plants in attainment counties.

Table A6: Effect on emissions ratios, by 3-digit NAICS code

	Onsite air	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Primary metals	0.136	0.184	0.148	2.236**	-0.167	-0.0984	0.269	0.116
	(0.198)	(0.272)	(0.832)	(0.905)	(0.487)	(0.280)	(0.380)	(0.257)
Chemicals	-0.387*	-0.142	4.568	2.757***	0.0138	-0.00659	0.162	1.143*
	(0.200)	(0.321)	(4.031)	(0.626)	(0.245)	(0.190)	(0.466)	(0.641)
Fabricated metals	-0.310	0.865			0.277	0.427	-0.596	0.287
	(0.330)	(0.770)			(0.280)	(0.474)	(0.435)	(0.244)
Nonmetallic mineral products	-0.685***		1.925***		3.567***	0.791***	0.269	-0.0952
	(0.206)		(0.509)		(0.259)	(0.237)	(0.479)	(0.226)
Transportation equipment	-0.896**	1.970	-0.493		0.632*	-0.130	2.762*	1.050
	(0.373)	(1.991)	(0.797)		(0.342)	(0.554)	(1.413)	(0.909)
Petroleum and coal	-1.370***	2.537***			6.282**	1.381**	-1.284	0.498
	(0.477)	(0.505)			(3.075)	(0.618)	(1.518)	(0.784)
Utilities	-0.355	1.180	-0.998	-0.159		1.742*	2.036	11.10***
	(0.426)	(0.893)	(1.140)	(0.593)		(0.987)	(1.457)	(0.286)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

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

Includes the seven 3-digit NAICS industries with the largest treated sample sizes. Column 1 (onsite air) corresponds to equation 4, remaining columns to equation 6. Dependent variable is log air emissions in column 1, otherwise log emissions ratio, with the numerator indicated atop the column and the denominator air emissions in all columns. All specifications include year dummies and plant fixed effects. SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Observation counts differ across columns because not all plants report emissions into all media. "Onsite other" emissions include waste piles, leaks, and spills.