Regulation-induced pollution substitution

Matthew Gibson*

July 9, 2016

Abstract

Regulations may cause firms to re-optimize over pollution inputs, leading to unintended consequences. 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 water emissions by 105 percent (72 log points). Regulation of an average plant increases air emissions at unregulated plants within the same firm by 13 percent. This leakage offsets 57 percent of emissions reductions by regulated firms. (JEL Q53, Q52, H23)

1 Introduction

Economic theory predicts firms will respond to environmental regulation by re-optimizing over pollution inputs. In the presence of unpriced or mispriced externalities, such responses can generate outcomes that are inefficient, unintended by policymakers, or both. The Clean Air Act (CAA) regulates particular air pollutants in particular counties, which creates incentives for firms to substitute among different forms of pollution. This paper tests two variants

*Department of Economics, Williams College, 24 Hopkins Hall Dr., Williamstown MA 01267; mg17@williams.edu. I would like to thank Maximilian Auffhammer, Peter Berck, Prashant Bharadwaj, Richard Carson, Andrew Chamberlain, Julie Cullen, Olivier Deschenes, Meredith Fowlie, Joshua Graff Zivin, Michael Greenstone, Carolyn Hayek, Kelsey Jack, Mark Jacobsen, Daniel Karney, Arik Levinson, Matthew Neidell, Jeffrey Perloff, Lynn Russell, Jeremy Schreifels, Ron Shadbegian, Glenn Sheriff, Jeffrey Shrader, Larry Sorrels, Junjie Zhang and participants in the UC San Diego environmental seminar for invaluable assistance with this project.

of this hypothesis: 1) Do firms respond to air pollution regulation by polluting other channels, like landfills and waterways? (cross-media substitution); and 2) Do multi-plant firms substitute toward pollution from less regulated plants? (spatial leakage). The existence and magnitude of such responses is important for both analysis of existing policies and design of future policies.

There is anecdotal evidence of such behavior. Duhigg (2009) describes a power plant that responded to air quality lawsuits by installing smokestack scrubbers, which spray water and chemicals into the stream of exhaust gases. The plant dumped the resulting liquid waste from the scrubbers into the Allegheny River. Previous empirical studies, however, have not found much evidence of cross-media substitution. 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 emissions of volatile organic compounds (VOC) by chemical manufacturers as a function of CAA regulation, proxying for output changes with employment changes. She finds no increased emissions into other media. Both Greenstone and Gamper-Rabindran model emissions differences as a function of county-level CAA regulation.

My approach builds on this work along several dimensions. First, I account for spatial heterogeneity in regulation. 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. My study follows the implications of Auffhammer et al. (2009), which finds that the effect of CAA non-attainment on the average monitor in a non-attainment county is zero, but that the effect on monitors above the CAA standard is -11 to -14 percent. Similarly, Bento et al. (2014) find that non-attainment affects home prices near non-attainment monitors, but not farther away. These findings suggest that regulators respond to non-attainment by focusing on problematic areas, rather than requiring uniform changes across a county. I demonstrate that only plants near non-attainment monitors are treated under the CAA. This pattern is consistent with a regulator whose objective function involves minimization of enforcement costs, either pecuniary or political (Amacher and Malik, 1996), rather than socially efficient abatement. My analysis of substitution accounts for this and so avoids averaging changes at treated plants with null responses from untreated plants in non-attainment counties.

Second, by estimating in log levels rather than log differences, I account for two important facts: 1) the CAA allows states and firms to respond slowly (over three years) to a non-attainment designation; and 2) many abatement decisions are discrete, producing a one-time change in the level of emissions. On average, air emissions fall in the first few regulated years and then stabilize. Specifications using log differences as the dependent variable average the negative growth rate effects in the first few years with the zero growth rate effects in the more numerous later years. As a result they are biased downward in magnitude. By pooling log levels across all regulated years, my estimating equations produce consistent estimates of the difference between pre- and post-treatment emissions. Third, motivated by a simple theoretical model, I show that one can use emissions ratios to recover the signs of net substitution elasticities among pollution inputs. Ratio estimation avoids confounding substitution and output effects.

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. My identification relies on the exogeneity of non-attainment status and monitor locations with respect to time-varying plant characteristics. The exogeneity of non-attainment status derives largely from the small share of point sources in particulate emissions (25 percent; Auffhammer et al., 2011). The exogeneity of monitor placement derives from EPA placement rules, which are based on population characteristics (e.g. average age) rather than industrial characteristics, and the prohibitively high cost of relocating a plant in response to monitor placement (Raffuse et al., 2007).

Both cross-media substitution and spatial leakage occur and responses can be large. Regulated plants increase their water emissions by 105 percent (72 log points). Regulation of an average plant increases air emissions at unregulated plants owned by the same firm by 13 percent, offsetting 57 percent of a firm's emissions reductions. The latter result recommends caution in studying CAA impacts with difference-in-differences designs.

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. Optimal 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 (Muller and Mendelsohn, 2009). While such an optimal policy might not be feasible or consonant with policymaker goals, patterns of substitution among pollutants are nonetheless a vital input into policy design.

This analysis contributes to the literature on regulation in the presence of mispriced inputs (e.g. Campbell, 1991). To the best of my knowledge, it is the first work to document regulation-induced cross-media pollution substitution. My findings are consistent with the theoretical work of Fullerton and Karney (2014) on pollution substitution. More generally, they complement the important work by Walker (2011, 2013) on labor input changes from CAA regulation. This study also contributes to the literature on pollution leakage. To date this literature has focused on international leakage (Levinson and Taylor, 2008; Davis and Kahn, 2010; Hanna, 2010) and simulated carbon leakage (Fowlie, 2009; Bushnell and Mansur, 2011). Both Henderson (1996) and Becker and Henderson (2000) find the CAA makes firms more likely to enter attainment counties, which might be considered a form of leakage. Fowlie (2010) demonstrates reallocation of NO_x emissions across plants in response to the NO_x Budget Program, but in that case such reallocation was among the aims of the policy. To the best of my knowledge, this is the first study to find evidence of unintended emissions leakage across existing domestic plants, and the first to show spatial heterogeneity in CAA-driven air pollution reductions at the plant level.

The rest of the paper is organized as follows. Section 2 provides background on regulations and abatement strategies. Section 3 discusses a simple theoretical model that informs my estimation. Section 4 describes the data, while Section 5 defines treatment and discusses identifying assumptions. Section 6 presents estimating equations and results and Section 7 explores their robustness. Section 8 concludes.

2 Background

2.1 The Clean Air Act and State Implementation Plans

Under the Clean Air Act, the EPA sets air quality standards for six criteria pollutants: carbon monoxide (CO), nitrogen dioxide (NO_2), particulate matter (PM), lead (Pb), sulfur dioxide (SO_2), and volatile organic compounds (VOC). For detailed information on particulate standards, which are the focus of this paper, see Appendix Table A3. A county violates the standard for a particular pollutant if at least one monitor exceeds the CAA standard in a given year.[1] In what follows, I refer to a monitor that exceeds the annual standard as a non-attainment monitor. A monitor violation triggers the following sequence of events (author's interview notes; Environmental Protection Agency, Undated):

1. Together EPA and the state go through a process to designate a county as non-attainment. This may take up to two years.

2. Non-attainment designation begins a process by which states submit a State Implementation Plan (SIP) to EPA. This may take 18 to 36 months.

3. SIPs are not federally enforceable until EPA approves them, but state authorities may enforce them prior to such approval. As a result actual regulation sometimes begins concurrent with a non-attainment designation, but often begins after a delay of a year or more.

[1] While EPA sometimes regulates smaller areas within counties, this far less common than county-level regulation (author's interview notes from conversations with EPA officials).

As a result of such lags, states have often drafted or even submitted SIPs before one of their constituent counties officially receives a non-attainment designation (see for example Missoula County Environmental Health Division, 1999).

SIPs detail steps that will bring the county into attainment. They may address both point and non-point sources of air pollution. Under a SIP, a state issues air emissions permits to plants. These typically include "lowest achievable emissions rates" (LAER) equipment requirements and plant-specific emissions limits (Becker and Henderson, 2000, 2001; Walker, 2013). SIPs may prescribe a specific control technology for a plant, but they often allow a plant to choose an abatement strategy (discussed in Section 2.3). In either case, a plant's permit under the SIP is the outcome of a negotiation between the state and the plant. That negotiation covers many potential abatement strategies and does consider private cost to the firm. In advocating a particular abatement technology, a state may think about its ability to monitor and enforce the permit as much as abatement effectiveness. The state typically does not want, for example, to permit a strategy that would allow a plant to repeatedly claim special circumstances and not abate. EPA or the state may provide guidance to the plant on probable costs of different technologies.[2]

State and EPA enforcement mechanisms include fines, inspections, and withholding of federal highway funds (Becker and Henderson, 2000; Chay and Greenstone, 2005). Once a state brings a non-attainment county back into compliance with CAA standards, it applies to EPA to have the county re-designated as an attainment county. That re-designation request must include a revised SIP, with a maintenance plan covering at least the next ten years (United States Code, 1990). In effect, such maintenance plans mean that most provisions of SIPs are made permanent.

Under the CAA, EPA periodically revises standards to reflect new research on the health effects of air pollution. For example, the agency finalized new PM_{10} standards[3] in 1987,

[2]This paragraph relies on the author's notes from telephone conversations and email exchanges with EPA officials, who generously commented on this paper.

[3]PM_{10} is particulate matter 10 microns or less in diameter.

but did not designate a county in non-attainment of the new standard until 1990. Revisions of CAA standards typically cause large numbers of counties to fall into non-attainment simultaneously. In my county-level data (described in Section 4), PM non-attainment lasts for an average of approximately 4 years. Conditional on PM non-attainment in at least one year, I observe on average .78 entries into non-attainment (some counties are already in non-attainment in the first year of my data) and .69 exits. Of the 299 counties observed in non-attainment, 126 remain in non-attainment through 2014, the last year of my data. For three counties I observe two entries into non-attainment.

2.2 Regulation of water and land emissions

CAA-induced substitution will reduce welfare, relative to an efficient policy, only if substitute emissions are unpriced or under-priced. Such is plausibly 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 689 chemicals. The PPL lists 126 chemicals (Environmental Protection Agency, 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 and have no "significant nexus" with navigable waters (Environmental Protection Agency, 2008).[4] EPA has limited this potential problem in the electric power sector, promulgating a 2015 rule on water discharges (Environmental Protection Agency, 2015b). The rule specifically targets secondary waste streams from air pollution abatement technologies.

The Resource Conservation and Recovery Act (RCRA) governs many forms of toxic disposal on land. Coal combustion residuals were exempt from many provisions of the RCRA during the period I study. In 2015, however, EPA issued a rule under the RCRA imposing technical

[4] Waters excluded from the federal CWA may still be regulated at the state level.

requirements on coal ash landfills and surface impoundment ponds (Environmental Protection Agency, 2015a). Some mining and petrochemical wastes remain exempt (Environmental Protection Agency, 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" (Environmental Protection Agency, Undated). While the incomplete regulations governing water and land emissions suggest cross-media substitution may reduce welfare relative to the first-best case, a full welfare analysis is beyond the scope of this paper.

2.3 Abatement strategies and variable costs

If abatement entailed no variable costs, plants would have no incentive to substitute in response. While abatement technologies usually have fixed costs, they also have large operating costs. Pollution control devices typically require substantial energy and may yield secondary wastes that require costly disposal. Processes that employ catalysts require periodic replacement of the catalyst. These variable costs range from 33 to 100 percent of capital cost for most abatement technologies (Environmental Protection Agency, Undated; Vatavuk et al., 2000; Farnsworth, 2011). For other abatement options like fuel switching and coal washing, the new fuel must be weakly more expensive than the old, or the plant would have been using it before. A similar logic applies to technological process changes.[5] Such costs mean that CAA non-attainment changes the relative price of air emissions for regulated plants.

I catalog the most common particulate air emissions abatement strategies in Table 1. Many abatement strategies produce secondary waste streams. For example, incineration decreases toxic air emissions but increases carbon emissions. Wet scrubbers "...can lead to water and

[5] Note that the set of available technologies may partially reflect the US policy environment, as the Clean Air and Clean Water acts have been in place, in approximately their current form, since the early 1970s.

solid waste pollution problems" (EC/R Incorporated, 1998). Theory predicts that firms will not consider the external costs of such secondary waste and so their abatement strategies may not be socially efficient. Even when a SIP prescribes a particular strategy, this may not correspond to the social optimum. A state regulator's objective function may differ from a hypothetical social planner's objective function, for example. Incomplete information could also lead to inefficient SIPs. In discussing potential environmental harm from cross-media substitution, EPA claims, "Such well-established adverse effects and their costs are normal and *assumed* to be reasonable and should not, in most cases, justify nonuse of the control technology" (Domike and Zacaroli, 2011, author's emphasis). If states do not have forecasts of welfare losses from secondary waste disposal, then SIPs may generate levels of pollution into landfills and waterways that exceed social optima.

3 Theory

The following simple model informs interpretation of my empirical results for cross-media substitution. Suppose a price-taking firm produces a single good using two pollution inputs A and W and a composite third input L comprised of labor, capital, land, etc. For discussion, let A be air emissions and W be water emissions. The CAA may be viewed as shift in relative input 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 2.3; and 2) non-pecuniary cost, for example the cost of incurring the displeasure of a regulator.[6]

Assume that the firm's cost function is multiplicatively separable into a function of quantity and a function of input prices.

$$C\left(Q, p_A, p_W, p_L\right) = f\left(Q\right) g\left(p_A, p_W, p_L\right)$$

The assumption of multiplicative separability is weaker than an assumption of constant re-

[6]For a model that treats CAA non-attainment as a limit on the quantity of air emissions, please see Appendix Section A.1.1. The qualitative predications from that model are the same as those presented here.

turns to scale (CRS; above this would imply $f(Q) = Q$). Like all cost functions, $C(\cdot)$ is homogeneous of degree one in input prices. Using the Envelope Theorem, one can differentiate the cost function to obtain conditional input demands.

$$\begin{aligned} W^* &= \frac{\partial C}{\partial p_W} = f(Q)\frac{\partial g}{\partial p_W} \\ A^* &= \frac{\partial C}{\partial p_A} = f(Q)\frac{\partial g}{\partial p_A} \end{aligned}$$

These input demands are homogeneous of degree zero in input prices, so they can be written as functions of price ratios. Dividing yields an optimal input ratio that is independent of output Q.

$$\frac{W^*}{A^*} = \frac{\frac{\partial g}{\partial p_W}}{\frac{\partial g}{\partial p_A}} \equiv h\left(\frac{p_A}{p_W}, \frac{p_A}{p_L}, \frac{p_L}{p_W}\right) \tag{1}$$

Under these conditions, therefore, one can learn about net substitution by estimating the relationship between the optimal input ratio and prices or proxies for prices. One does not have to control for quantity, as is common practice in estimation of translog cost functions (see for example Westbrook and Buckley (1990)).

The assumption of multiplicative separability in the cost function builds on previous pollution substitution work, which has typically employed a CRS assumption (see for example Fullerton and Karney (2014)). CRS implies multiplicative separability. Forms like CES and Cobb-Douglas, which are commonly used for aggregate production functions, satisfy this property irrespective of returns to scale.

For expositional convenience, I will temporarily assume production is CES. Then $\frac{W^*}{A^*} = h\left(\frac{p_A}{p_W}, \frac{p_A}{p_L}, \frac{p_L}{p_W}\right) = \left(\frac{c_W}{c_A}\frac{p_A}{p_W}\right)^\sigma$ (the derivation is in Section A.1.2), where c_A and c_W are technological constants. Taking logs yields the following expression.

$$\ln\left(\frac{W^*}{A^*}\right) = \sigma\ln\left(\frac{c_W}{c_A}\right) + \sigma\ln\left(\frac{p_A}{p_W}\right) \tag{2}$$

In this case a single global parameter σ represents the Morishima elasticity of substitution with respect to price p_A (Blackorby and Russell, 1989):

$$\sigma = M_{AW}\left(Y, p_A, p_W\right) = \varepsilon_{WA} - \varepsilon_{AA}$$

where ε_{WA} and ε_{AA} are cross- and own-price elasticities of factor demand. While this is the natural generalization of the Hicks elasticity, its asymmetry makes it different in one important respect: the elasticity M_{AW} is informative for changes in p_A but not for changes in p_W. The sign of M_{AW} is ambiguous because the sign of ε_{WA} is unknown when there are three or more inputs. If M_{AW} is positive, the inputs are net substitutes. If it is negative, they are net complements. Given a proxy for $\frac{p_A}{p_W}$, one can recover the sign of the Morishima elasticity. One need not assume CES, however. Under the assumption of multiplicative separability in the cost function, one can recover the sign this elasticity without observing output, but it may be local, rather than global as in the CES case.

Note that controlling for additional inputs (beyond A and W) would force the tradeoff back into the $A - W$ plane. As Blackorby and Russell (1989) argue, this measure of curvature is interesting but substantially less informative than the Morishima elasticity. My ratio-based regression models assume not that the change in p_A has no effect on other inputs, but rather that only p_A changes and other prices remain constant. If the plants under study are price takers in factor markets and CAA non-attainment does not produce general-equilibrium effects on other factor prices, then this assumption likely holds.

The ability to estimate Morishima elasticities without observing output is useful 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. Gross and net (Morishima) elasticities have different signs. Modeling ratios allows me to infer when pollution inputs are net substitutes in production, even if they are gross complements.

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 actually reveals the inputs are net substitutes. Figure 1 illustrates this for the two-input case. 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 $W_2 > W_0$. Had the firm held output fixed, water emissions would have increased.

The preceding discussion assumes a static production technology, with input substitution driven by exogenous price changes. This assumption could be incorrect if firms respond to regulation with both installation of new pollution-control capital and input substitution. If one thinks of the new pollution-control equipment as an increase in capital, the ratio approach still recovers the Morishima elasticity, which allows for responses in inputs other than the pair being considered. If instead one thinks of the new pollution-control equipment as a change in parameters or functional form, a ratio approach is potentially problematic. Only under a relatively strong CES functional form assumption can the ratio approach handle this case. Technological change can be modeled with changes in the constants c_W and c_A. Factoring equation 2 yields the following.

$$\ln\left(\frac{W^*}{A^*}\right) = \sigma\left[\ln\left(\frac{c_W}{c_A}\right) + \ln\left(\frac{p_A}{p_W}\right)\right] \tag{3}$$

Given a proxy for the quantity $\left[\ln\left(\frac{c_W}{c_A}\right) + \ln\left(\frac{p_A}{p_W}\right)\right]$, it is still possible to recover a scalar function of the Morishima elasticity. CES is the *only* functional form that factors such that the optimal input ratio responds identically to changes in the ratio of technological constants and changes in the ratio of prices.

None of the above is meant to suggest that the Morishima elasticity is the only object of interest in this context. Modeling W^*, rather than $\frac{W^*}{A^*}$, as a function of prices is informative about *gross* substitution, which may be the object of greater policy interest. I present both ratio and non-ratio results in Section 6.2.

4 Data

My plant-level emissions and location data come from the EPA Toxic Release Inventory (TRI) 1987-2014. As of this writing, the TRI records annual emissions of 689 chemicals by mass (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 2. The database also includes the Dun & Bradstreet DUNS number for the parent company of each plant.

These data have several shortcomings, discussed in Hamilton (2005). Only large facilities are required to participate.[7] Firms typically report estimates derived from engineering models, rather than direct measurements. There is no straightforward measure of output.[8] 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). In the early years of TRI data collection, reporting requirements changed dramatically. For example, reported pollution increased sixfold between 1990 and 1991 due to reporting changes required by the Pollution Prevention Act (Environmental Protection Agency, 2012). To avoid confounding such reporting changes with genuine emissions changes, I exclude the period 1987-1991 from my

[7] Reporting thresholds have varied over time and by chemical. Typically a plant must report if it meets all of the following 3 criteria: 1) manufactures 25,000 lb/year, processes 25,000 lb/year, or uses 10,000 lb/year of a TRI-listed chemical; 2) employs 10 or more FTE workers; 3) in a covered SIC code.

[8] The TRI does include a "production or activity ratio." In some cases this is equal to the ratio of output in year t to the ratio of output in year t-1. In others it is equal to the ratio of activity rates, e.g. the number of cleanings in year t divided by the number of cleanings in year t-1. Firms choose which of these ratios they report.

analysis.

A subset of TRI chemicals are classified as particulates (PM).[9] The TRI data capture emissions in great detail, distinguishing for example between different types of underground wells. To simplify presentation and analysis I aggregate up to the categories described in Table 3 by adding the mass of each chemical emitted (in pounds).

Data on county attainment status come from the EPA Green Book 1992-2014. Monitor-level data on pollutant concentrations come from the EPA Air Quality System (AQS) 1990-2014. For descriptive statistics see Appendix Table A1.

5 Treatment

5.1 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:

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

[9]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 VOC, 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 VOC are not directly regulated under the CAA. They are one of two primary precursors (the other is NO_x) 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 VOC contribute substantially to ozone formation. EPA regulates PM_{10} (particles $<$10 microns in diameter) and $PM_{2.5}$ ($<$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, NO_2, or SO_2.

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 counties that were in non-attainment in the previous year against the distance to the nearest non-attainment monitor in the previous year. Residuals are large and negative (roughly -50 log points) near the non-attainment monitor, indicating air emissions abatement. As distance to the monitor increases, the residuals rapidly rise to zero near one kilometer and remain there. This figure provides evidence that regulators indeed treat plants near non-attainment monitors intensively, while treating more distant plants lightly or not at all. The pattern is consistent with the hedonic results from Bento et al. (2014). This spatial heterogeneity does not imply that studies finding effects of county non-attainment on home prices (Chay and Greenstone, 2005) or health (Chay and Greenstone, 2003a) are biased. Rather, they report unbiased county-average effects that may conceal substantial within-county heterogeneity.

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 non-attainment monitor. A non-attainment monitor is one that violated the CAA standard in year $t-1$ or previously. Based on Figure 2, I use a threshold distance $\overline{D}$ of 1.07 kilometers, the distance at which I can no longer reject a null hypothesis of a zero treatment effect on air emissions (at the 5 percent level). In Section 7.5 I discuss the sensitivity of my empirical results to this threshold distance.[10]

I use lagged rather than contemporaneous non-attainment status because: 1) state regulations may not take effect in the first non-attainment year (see Section 2); and 2) some firm responses plausibly require substantial time to implement (e.g., existing contracts might limit fuel switching). This treatment variable forms the basis for all subsequent results. Defining treatment in this way invokes an additional identifying assumption, exogeneity of monitor placement with respect to plant-level scope for abatement and substitution, which I discuss

[10]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.

in Section 5.2. This spatial pattern is consistent with a regulator whose objective function involves minimization of enforcement costs, either pecuniary or political (Amacher and Malik, 1996). The qualitative evidence presented by Becker and Henderson (2000) on regulator-firm negotiations is also consistent with such an explanation. Because maintenance plans make most SIP regulations permanent (see Section 2.1), I assign a plant to the treatment group in any year after it is first treated, even if the county in which it is located is re-designated as in attainment of CAA standards.

I use a dummy treatment variable for two primary reasons. First, a dummy simplifies the relationship between estimates from equation 7 and the underlying net elasticities. Second, a dummy allows me to construct an easily interpretable, plausibly exogenous measure of firm-level regulatory exposure by counting treated plants (see equation 8). One could construct a more continuous firm exposure measure. For example, the count of treated plants could be weighted by inverse distance to the nearest non-attainment monitor, by the square of that inverse distance, or by pre-treatment air emissions. Such variables require additional researcher choices, however, and may invoke additional identifying assumptions.

5.2 Treatment exogeneity

I cannot recover the causal effects of treatment unless it is exogenous to my plant-level outcomes of interest. Concretely, I assume exogeneity of: 1) county-level attainment status; and 2) distance to the nearest non-attainment monitor. As for the first assumption, past literature has typically argued that county non-attainment is exogenous.[11] Chay and Greenstone (2003a,b, 2005) document that PM_{10} non-attainment counties do not differ systematically from attainment counties on observable dimensions (including economic shocks), either in levels or in changes. Appendix Table A1 shows that the emissions profiles of plants in PM attainment and non-attainment counties are not statistically different in my data.

Non-attainment is plausibly exogenous if a given firm produces a small portion of the ambient

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

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 PM_{10} 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, however, calls into question the exogeneity of CAA regulation for *treated* plants (plants actually affected by regulation). CAA regulations primarily affect plants within one kilometer 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. Endogenous past output could bias my estimates of CAA treatment effects on log emissions. For example, if output shocks were negatively autocorrelated, my estimates might overstate the magnitude of CAA treatment effects. If instead output shocks were positively autocorrelated, it might understate them.

To investigate the possibility of endogenous entry into treatment, I estimate an event-study specification for air emissions.

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

The variables τ_j are indicators for a time index defined relative to treatment. I include dummies for $\tau = -5$, $\tau = -4$, $\tau = -3$, $\tau = -2$, $\tau = -1$, $\tau = 0$, $\tau = 1$, $\tau = 2$, $\tau = 3$, and $\tau \geq 4$, so the reference category is the average of years for which $\tau < -5$. A county receives a non-attainment designation in year $\tau = -1$ and plants within one kilometer of a non-attainment monitor enter treatment in the following year ($\tau = 0$). Figure 3 presents coefficient estimates (the corresponding numbers are in Table 4). If the figure showed higher air emissions at $\tau = -1$, that would be evidence of endogenous entry into treatment. The figure shows no such pattern. Air emissions are roughly flat in the pre-treatment period and decline after treatment begins.

The second identifying assumption is exogeneity of distance to the nearest non-attainment monitor. Violations of this assumption could spring from two sources: firm location decisions and state monitor placement decisions. Given the relatively low cost of new monitors, firms are unlikely to profit by strategically locating away from existing monitors. The state monitor location decision warrants more discussion. States design monitoring networks, which must follow EPA rules and which EPA must approve (CFR, 2015). EPA may also suggest changes to planned networks. Importantly in this setting, the agency's placement rules largely depend on population characteristics, not firm characteristics. For example, EPA requires monitors in areas of high population density (Bento et al., 2014) and near large sensitive populations (e.g. asthmatic children Raffuse et al., 2007). Two types of monitoring sites raise potential endogeneity concerns: "Sites located to determine the impact of significant sources or source categories on air quality" and "Sites located to determine the highest concentrations expected to occur in the area covered by the network" (CFR, 2015). Monitors placed under these two rules could be correlated with unobservable time-varying characteristics of plants, as discussed below. States are prohibited from putting monitors in locations that do not meet scientific criteria. In most cases it is illegal for a state to move a monitor, and EPA allows relocation only if the new site is better under its scientific criteria. Should a state fail to follow these rules, EPA may file suit against it (Chay and Greenstone, 2005).

Note that my identifying assumption is *exogeneity of distance to the nearest non-attainment monitor,* not distance to the nearest monitor. The former is a weaker assumption, particularly given the event-study evidence that the plants in my data are not pivotal in putting their counties into non-attainment. Nonetheless, to investigate potential endogeneity, I first regress log distance to the nearest non-attainment monitor on a set of 317 dummies for six-digit NAICS codes, omitting the constant term. Figure A2 displays the probability density function of the coefficient estimates. While the distribution is roughly normal around a mean of 2.1, some coefficients are statistically distinguishable from that mean in both the positive and negative directions. Industries in the tails show no clear pattern. They include, for example, beet sugar manufacturing, prisons, and national defense. The R^2 from the regression is .78, indicating that industry explains a substantial fraction of the variation in distance to the nearest non-attainment monitor. This suggests that plant fixed effects are necessary to my identification strategy, but even with plant fixed effects the possibility of non-zero covariance between time-varying plant unobservables and monitor distance remains. To evaluate this threat to identification, I regress the log distance to the nearest non-attainment monitor on a vector of year dummies and the changes in log emissions into various media for untreated plant-years (pre-treatment or farther than two kilometers from the nearest non-attainment monitor). A negative coefficient is consistent with states strategically placing monitors near faster-growing emissions sources. Table A5 shows that all eight estimates are zero to two decimal places and are not statistically significant. Emissions growth rates in untreated plant-years generally do not systematically predict distance from eventual non-attainment monitors. Appendix Table A4 presents a version of this specification using emissions levels instead of growth rates. Estimates are practically large and statistically significant. Like Figure A2, they imply that plant fixed effects are necessary to my identification strategy.

6 Empirical strategy and results

6.1 Air emissions

To estimate reduced-form treatment effects on emissions into various media, I use the following specification, with i indexing plant and t year.

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

The dependent variable is the log of a plant's emissions into a particular medium, e.g. air or water. The equation includes plant fixed effects and year dummies, with the latter capturing secular forces influencing emissions. If CAA regulations are effective in reducing air emissions, the estimate of β will be negative for air emissions. If firms employ cross-media substitution in response, the estimates of β will be positive for other media like water and landfills. Because this specification does not control for output or other inputs, it captures the full effect of the policy, including both output and substitution effects.

Table 4 presents my estimate of the CAA treatment effect on airborne particulate emissions, where treatment is defined as in Section 5. Treated plants decrease their air emissions by 38 percent (49 log points, statistically significant at the one percent level). 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 state linear time trends. This reduces the magnitude of the estimate modestly, to 33 percent (41 log points), but it remains statistically significant at the five percent level. Column 3 adds NAICS-year fixed effects, and the estimate is practically unchanged at -48 log points. If there is substantial general-equilibrium leakage to untreated plants, my estimated effects on air emissions will be biased upward in magnitude. I investigate this possibility in Section 7.1 and provide evidence this is not a substantial concern. Bento et al. (2014) show these air quality improvements disproportionately benefit low-income people, at least in the short run.

Column 4 presents the results from an event-study specification (equation 5), where again the dependent variable is log air emissions. The time pattern suggests that most of the emissions reductions occur when τ is 1 (the second treated year). This helps motivate my use of fixed-effects models in levels. Estimates based on changes in treatment status would be biased toward zero because one does not see meaningful emissions declines at $\tau = -1$ (the first non-attainment year) or $\tau = 0$ (the first treated year). At approximately -50 log points, the event-study estimates are close in magnitude to my primary result (-49 log points). Together these results demonstrate that treated plants do indeed reduce airborne particulate emissions. Sample size falls in the event study specification because some plants are already in non-attainment counties in the first year of my data. Because treatment is permanent and my specifications include plant fixed effects, such plants do not help identify the treatment effect in columns 1 through 3. Thus the identifying variation is the same in all four columns, despite the different sample size in column 4.

6.2 Cross-media substitution, all industries

Panel A in Table 5 shows estimated treatment effects from equation 6, by medium across all industries. Treated plants increase water emissions by 105 percent (72 log points). This is evidence that water and air emissions are gross substitutes in production. Effects on other media, including releases to recycling and treatment, are not statistically significant. ("Onsite other" emissions include waste piles, leaks, and spills.) Panel B adds state linear time trends and the estimate for water increases slightly, to 77 log points. This increase in water emissions imposes social costs, which are difficult to quantify, given the relative scarcity of well-identified studies on the health and productivity effects of water pollution. Nonetheless one can say something about the likely efficiency properties of such substitution. If water emissions are mispriced, such substitution may be inefficient. The incompleteness of water and regulations (see Section 2.2), coupled with firm incentives to minimize private abatement cost, suggests this may be the case. It is possible that states set socially optimal relative prices in their SIPS (see Section 2.1). In order to do so, however, states must share a social planner's objective function and have complete information about the welfare effects of substitution toward water.

To investigate *net* (Morishima) elasticities of substitution across media, I estimate the following.

$$\ln\left(\frac{W_{it}}{A_{it}}\right) = \overline{\alpha_i} + \overline{\delta_t} + \beta treated_{it} + \varepsilon_{it} \tag{7}$$

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 conditional factor demands from equation 1 above. The treatment dummy proxies for the unobservable increase in the price ratio $\frac{p_A}{p_W}$. The coefficient $\beta = \nu\sigma$ is a scalar function of the Morishima elasticity of substitution σ, where ν is the percentage increase in relative prices produced by treatment. If air and water emissions are net substitutes, theory predicts the CAA will induce cross-media substitution and estimates of β will be positive. If instead air and water emissions are complements, estimates of β will be negative.

Equation 6 and 7 both involve log dependent variables and thus exclude plants reporting zero emissions into a given medium. Plants that do not emit into the air are beyond the scope of this work. There are no plants that do not emit into a given medium before treatment, then begin to do so after treatment.

Panel A in Table 6 presents effects on emissions ratios, based on equation 7 (again 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 imply positive net elasticities of substitution. There is evidence of statistically significant substitution toward onsite water emissions $\left(\hat{\beta} = 1.02\right)$ and offsite water emissions $\left(\hat{\beta} = .58\right)$. The negative estimates for onsite other and offsite other emissions demonstrate that the ratio approach does not assume positive net elasticities of substitution. Panel B in Table 6 adds state linear time trends and estimates are essentially unchanged from panel A.

As mentioned above, these ratio estimates are scalar multiples of underlying net substitution elasticities. Assuming treatment increases the price ratio $\frac{p_A}{p_W}$, the estimates and the underlying elasticities have the same sign. By itself this fact is informative. Note for example that Table 5 shows decreased emissions to recyclers (the two inputs are gross complements). One might erroneously infer that air emissions and emissions to recyclers are net complements. In the ratio specification (Table 6), however, the estimate is positive (albeit not statistically significant), suggesting these two forms of emissions may be net substitutes. To obtain the elasticity from one of these ratio estimates, one must divide by the percentage change in relative prices produced by treatment, which is unobserved. Given $\hat{\beta} = 1.02$ for onsite water emissions, if the increase in relative prices is less than 100 log points, then $\sigma > 1$. That suggests that in the aggregate US production function, there is a good deal of net substitutability between onsite air pollution and onsite water pollution.

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 7 presents estimates for the three industries with the largest treated sample sizes: primary metals, wood products, and utilities. (Appendix Table A12 shows effects on log emissions ratios by 2-digit NAICS, while Appendix Table A13 shows effects on log emissions by 3-digit NAICS.) Estimates again come from equation 6. In the discussion that follows, note that I cannot reject the null hypothesis of equal coefficients in many cases; the evidence of heterogeneity is merely suggestive. Primary metals show only a 4 percent decrease in air emissions, while wood products and utilities show large decreases, 65 and 58 percent respectively (-106 and -86 log points). In water emissions, primary metals show a 219 percent increase (116 log points).[12] The corresponding estimate for wood products is small, negative, and imprecise. The corresponding estimate for utilities is large and positive (68 log points), but again imprecise. Wood products and utilities increase onsite land emissions by 214 and 474 percent respectively (115 and 175 log points), while primary metals decrease such emissions by 78 percent (152 log points). Utilities substantially decrease their offsite water emissions (-85 percent or -191 log points).

[12] For primary metals, and for many other industries examined in this paper, in proportional terms increases in emissions into other media are much larger than decreases in air emissions. This is because baseline air emissions are generally much larger than baseline emissions into other media. For example, Appendix Table A1 shows air emissions are roughly six times greater than water emissions in both attainment and non-attainment counties.

6.4 Leakage

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

$$\ln(A_{it}) = \overline{\alpha_i} + \overline{\delta_t} + \beta other\_treated_{it} + \varepsilon_{it} \tag{8}$$

Again I include plant fixed effects and year dummies. The variable $other\_treated_{it}$ is a dummy for one or more treated plants within the same firm, year, and 2-digit NAICS code. If the CAA induces spatial leakage, estimates of β will be positive.

Intuition predicts that firms might respond to treatment of a plant in one county by shifting emissions to a plant in another county. Table 8 provides evidence they do so. Estimates correspond to equation 8. For the average plant in an attainment county, treatment of one or more plants within the same firm and 2-digit NAICS code increases air emissions by 15.8 percent. Column (2) adds state linear time trends and the estimate is slightly smaller at 15.3 percent. Treating the number of other treated plants as a continuous variable (column 3), estimated leakage is 12.9 percent per treated plant. With the addition of state linear time trends in column 4, the estimate is again slightly smaller at 12.4 percent, but remains statistically significant at the five percent level. This leakage has associated health, mortality, and productivity costs. I take the TRI parent company identifiers as given. If they are defined at a level below the ultimate corporate parent, my estimates will likely understate the true amount of leakage. Likewise, if there is general-equilibrium leakage to plants owned by other firms in attainment counties, my estimates will be biased downward (see Section 7.1). This model will not capture within-firm leakage to plants located in non-attainment counties, but beyond the threshold distance.

The identifying assumptions for this model are modestly stronger than for my model of cross-media substitution and warrant brief discussion. Limiting the sample to attainment-county plants changes the interpretation of the estimates, but is not in itself problematic, especially if attainment status is exogenous. Interpreting the estimates in Table 8 as causal, however, also requires that the leakage plants do not differ from other attainment-county plants in time-varying, unobservable ways. Appendix Table A2 shows that emissions profiles for leakage and non-leakage plants are not significantly different, which is reassuring but does not exclude the possibility of endogeneity.

The average treated firm in my data includes 1.4 treated plants and 21 leakage candidates: they share the same two-digit NAICS code and are located in attainment counties.[13] Average air emissions at eventually treated plants prior to treatment are 8970 pounds, while average baseline emissions at leakage candidates are 717 pounds. The estimate from column three of Table 8 implies the following net change in emissions from treating an average firm. The firm's treated plants reduce emissions by $1.4 * .38 * 8970 = 4772$ pounds.[14] The 21 candidate plants together increase emissions by $21 * .13 * 1.4 * 717 = 2740$ pounds. On net, then, the average firm treated under the CAA decreases particulate air emissions by $4772 - 2740 = 2032$ pounds. Roughly 57 percent of reductions at treated plants are offset by leakage. 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, relative to a first-best policy, because attainment-county emissions are unpriced (unregulated). This leakage need not imply a net welfare loss relative to baseline, however. Leakage-driven emissions increases occur in

[13] This is the average number of leakage candidates over all treated firms, including single-plant firms that have zero leakage candidates by definition.

[14] The 38% reduction is the percentage change corresponding to the estimated treatment effect on log air emissions: $e^{-.485} - 1 = .38$.

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.[15] 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. My estimated effects on air emissions (Table 4) will be biased upward in magnitude and my estimated effects on other emissions (Table 5) will be biased downward in magnitude. Such bias may be small if within-firm leakage is a minor determinant of ambient pollutant concentrations in attainment counties. Provided the optimal input ratio is independent of scale, spatial leakage will not bias my ratio-based estimates (Table 6).

7 Additional results, robustness & placebos

7.1 Air emissions

It is possible CAA regulation induces general-equilibrium leakage, with output reallocated from treated plants to attainment-county plants not owned by the same firm. If this is the case, my estimated effects on log emissions at treated plants will be biased upward in magnitude. My estimated within-firm leakage effects will be biased downward in magnitude. It is impossible to test directly for general-equilibrium leakage, since all plants are potentially affected by CAA regulation through general-equilibrium mechanisms. One can however test indirectly for general-equilibrium leakage by modeling the air emissions at untreated plants[16] as a function of the number of treated plants "nearby." To that end, I estimate the following

[15] Author's calculation from 2010 Census data.

[16] i.e. Attainment-county plants and plants farther than 1.07km from a non-attainment monitor in a non-attainment county

equation using plants in attainment counties.

$$\ln\left(A_{it}\right) = \overline{\alpha_i} + \overline{\delta_t} + \beta total\_treated_{jt} + \varepsilon_{it} \tag{9}$$

In this specification the variable $total\_treated_{jt}$ is a count of treated plants at either the state-year level or the state-year-NAICS2 level. If general-equilibrium leakage is occurring, estimates of β will be positive. Results appear in Appendix Table A6. In both specifications, the estimated coefficient on the number of treated plants in the same state-year is negative, insignificant, and zero to two decimal places. These estimates suggest general-equilibrium leakage is not a first-order source of bias in my estimates.

It is possible that intra-firm 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 6), controlling for intra-firm leakage as in equation 8. Reported in Appendix Table A7, the estimates are unchanged. This is likely because identification of the coefficients on the year dummies comes primarily from plants that are not leakage recipients.

Both Henderson (1996) and Becker and Henderson (2000) show that CAA non-attainment influences plant entry and exit decisions, and this is a potential source of bias. A Heckman correction would be inappropriate, as I do not have any variables that would enter the selection equation but not the outcome equation. Instead I restrict the sample to plants present throughout the study period and estimate treatment effects on air emissions (see Appendix Table A8). While the smaller sample reduces precision, at $-.45$ (statistically significant at 5 percent) and $-.34$ (not statistically significant), the estimates are close to the results in Table 4. This suggests selection does not meaningfully bias my main results.

Lastly I estimate specifications similar to those employed previously in this literature (Greenstone, 2003; Gamper-Rabindran, 2009) and present results in Appendix Table A9. Column 1 uses emissions differences as the dependent variable and defines treatment as I do in my primary analysis. Column 2 uses emissions levels as the dependent variable and lagged county non-attainment as the treatment. Column 3 combines these approaches, modeling emissions differences as a function of lagged county non-attainment. Estimates range from $-.013$ to $-.078$, roughly similar to the Greenstone (2003) and Gamper-Rabindran (2009) results. This demonstrates the importance of modeling air emissions in levels, rather than growth rates, and accounting for spatial heterogeneity in treatment intensity.

7.2 Cross-media substitution

Appendix Table A11 estimates a variant of my ratio specification, with log air emissions employed as a right-hand-side control rather than a denominator in the dependent variable. Results are similar in sign and significance to those from the ratio specification, but smaller in magnitude; the estimate for water emissions is .77, rather than 1.02 in the ratio specification. This specification has the benefit of allowing log air emissions to enter the cross-media model more flexibly, but at the cost of including an endogenous variable on the right-hand side of the equation.

To test whether intra-firm leakage influences my cross-media results, I estimate my leakage model using an emissions ratio as the dependent variable and report results in Appendix Table A14. Estimates are generally near zero and statistically insignificant, with one important exception: the estimate for onsite water is positive 22 percent (statistically significant at the 10 percent level). This leakage will produce a downward bias in the estimated effect on onsite water emissions from my cross-media model, because it increases water emissions in the control group. Such bias should be small, however, as leakage candidates constitute a small fraction of control-group plants.

In addition, Appendix Table A15 shows estimates for toxicity-weighted emissions into non-air media. Estimates for onsite water releases are larger in magnitude, but much less precise. There is also statistically significant evidence that firms are shifting some of their most toxic releases into waste piles (onsite other) and offsite water. I do not employ toxicity weights in my preferred specifications for the following reasons: 1) toxicity weights for a given chemical can vary by three orders of magnitude, depending on the method used (Hertwich et al., 1998); 2) toxicity weights rely on assumptions that some chemicals are not carcinogenic, but epidemiological evidence suggests such assumptions may not hold (Hendryx et al., 2012); 3) toxicity weights are not available for all TRI-listed chemicals. Finally, Appendix Table A16 presents estimated effects of county non-attainment. These are small and statistically indistinguishable from zero, because they average non-zero responses at treated plants with a larger number of zero responses at untreated plants in non-attainment counties. Estimates are similar to the Greenstone (2003) and Gamper-Rabindran (2009) results.

7.3 Leakage

As a robustness check on my leakage results I estimate the same model, grouping plants by firm and 6-digit NAICS code, and report results in Table A18. Estimates are similar to those from my preferred specification, but no longer statistically significant, as this grouping results in far fewer treated plants within a group. In Table A19 I include controls for firm size and the results are unchanged.

7.4 Placebos

Treatment should have no direct effect on plants that do not emit any air pollution, and the results from Appendix Table A6 suggest that general-equilibrium treatment effects are negligible. Table 9 tests this hypothesized null effect by estimating a variant of equation 6, where 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. Importantly the estimated effect on onsite water emissions from plants without air emissions is near zero. This suggests the estimated increase in onsite water emissions in Table 5 does not arise from gross misspecification.

Table 10 reports results from a placebo test of my leakage model. I construct variables based on placebo “treated” plants: plants within the same firm and 2-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 and general-equilibrium effects are not apparent, one 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 10 are in the range from negative one to positive two percent and are not statistically significant. This suggests that the leakage results in Table 8 do not spring from an omitted variable problem.

7.5 Distance threshold

Because the treatment variable I employ relies on an estimated threshold distance, it is important to explore the robustness of my findings with respect to changes in that distance. Appendix Table A10 shows effects on air emissions at five threshold distances, with two smaller and two larger than the 1.07 kilometers used in my primary analyses. As expected given Figure 2, smaller thresholds increase estimate magnitudes. This is consistent with plants closer to non-attainment monitors being regulated more intensively. Larger thresholds decrease estimate magnitudes, as the models begin to include potentially untreated plants in the treatment group. Appendix Table A17 performs a similar exercise for my models of onsite water emissions. While the estimates are less precise, the same broad pattern holds, with lower-magnitude treatment effects as threshold distance increases. Finally Appendix Table A20 estimates intrafirm leakage, again varying threshold distance. Again estimate magnitudes decline with increased threshold distance. In none of these tables is the sign or rough magnitude of my estimate appreciably altered by the choice of threshold. In Tables A10 and A20 statistical significance is also unaffected. In Table A17, however, some of the alternative thresholds do result in the loss of statistical significance.

Appendix Figure A1 demonstrates that the threshold is not sensitive to the range over which one plots the local polynomial. The left panel extends the range of the horizontal axis to 8.5 kilometers (50th percentile) and the right to 53 kilometers (95th percentile). In both cases the threshold distance is extremely close to the one used in my primary analysis. There is no evidence of a treatment effect on plants beyond the threshold distance.

The rule I employ to estimate threshold distance also warrants some discussion. The 1.07 kilometer threshold is the distance at which I can no longer reject a null hypothesis of zero effect on air emissions (at the five percent level). There are many other possible decision rules, e.g. the distance at which the local polynomial first achieves a zero value or the distance at which the local polynomial takes on a zero slope, but my choice is conservative in the following sense. My rule will tend to estimate a lower threshold distance than many alternative rules. Figure 2 shows that treatment intensity declines with distance to the nearest non-attainment monitor. If my rule errs, it does so by assigning treated plants to the control group. (Recall that plants in non-attainment counties, but beyond the 1.07 kilometer threshold, are part of the control group in my empirical models.) In a difference-in-differences design, such mis-assignments will bias the magnitudes of my estimates downward. That is, this rule makes it less likely my hypothesis tests will produce false positives.

8 Conclusion

While economists have long recognized the potential for substitution responses to location-specific, single-medium pollution regulation, empirical studies have found little evidence of such effects. Using specifications motivated by classical firm optimization theory, this study provides evidence of regulation-induced pollution substitution in response to the Clean Air Act. Estimates from EPA Toxic Release Inventory data show that CAA-regulated plants increase their onsite water emissions by 105 percent. Particulate regulation of an average plant increases air emissions at unregulated plants owned by the same firm by 13 percent. At the firm level, such leakage offsets 57 percent of emissions reductions at regulated plants. This paper examines only two possible types of pollution substitution. In addition, new source performance standards mean that new plants may use non-air pollution inputs more intensively, and locate more frequently in attainment counties (the latter is documented in Henderson (1996) and Becker and Henderson (2000)). Thus industry- or economy-wide responses may be larger in magnitude than the plant- and firm-level responses identified in this study.

The welfare effects of such substitution present an interesting subject for future research. Air pollution regulations can have large benefits (Chay and Greenstone, 2003b; Currie and Neidell, 2005). In particular, EPA estimated the 1970-1990 benefits of the Clean Air Act (CAA) at $22 trillion (Environmental Protection Agency, 2011). While social costs from firm re-optimization are plausibly smaller, they may be large in absolute terms.

A policy with efficiency among its goals should account for these firm responses. A maximally efficient policy, with emissions into every medium and location priced according to marginal damage, would be difficult to achieve and might not be desirable for normative reasons. The primary goal of the Clean Air Act is not efficiency, but rather safeguarding human health (Environmental Protection Agency, 2011). Given any set of policy goals, however, it is easier to formulate effective policy when policymakers have well-identified estimates of firm responses. For example, my cross-media results suggest that restricting water emissions or increasing water quality monitoring in CAA non-attainment counties might be important for protecting public health. They also suggest EPA's recent rules on coal ash disposal and water pollution from power plants are likely to constrain firms. My estimates of within-firm leakage imply that applying PM-style regulations to carbon emissions would be largely ineffective.

Additionally, I document spatial heterogeneity in regulatory intensity. Most plants in non-attainment counties show no evidence of being regulated, but plants near non-attainment monitors show large air emissions decreases. This pattern is consistent with theoretical models in which regulators seek to minimize costs (political or pecuniary) in implementing the CAA, but there are other possible explanations. Questions concerning state implementation of federal environmental regulations warrant additional research, building on work like Levinson (2003), Helland (1998), and Sigman (2003, 2005). Legislators might also want to consider the regulator behavior implied by my spatial heterogeneity result when designing future policy.

Such improvements in policy design would likely have economically significant consequences.

While environmental economics research initially focused on the mortality effects of air pollution, especially for infants and the elderly, there is growing evidence that air pollution has costly effects on healthy adults. Isen et al. (2014), for example, find that in-utero and early childhood air pollution exposure depresses earnings for workers ages 29-31. Zivin and Neidell (2012) find air pollution decreases worker productivity. Given these large costs, the returns to improved pollution regulation may be large.

References

Amacher, G. S., Malik, A. S., 1996. Bargaining in environmental regulation and the ideal regulator. Journal of Environmental Economics and Management 30 (2), 233–253.

Auffhammer, M., Bento, A. M., Lowe, S. E., Jul. 2009. Measuring the effects of the Clean Air Act Amendments on ambient concentrations: The critical importance of a spatially disaggregated analysis. Journal of Environmental Economics and Management 58 (1), 15–26.

Auffhammer, M., Bento, A. M., Lowe, S. E., 2011. The city-level effects of the 1990 clean air act amendments. Land Economics 87 (1), 1–18.

Becker, R., Henderson, V., 2000. Effects of Air Quality Regulations on Polluting Industries. Journal of Political Economy 108 (2), 379–421.

Becker, R. A., Henderson, J. V., 2001. Costs of air quality regulation. In: Behavioral and distributional effects of environmental policy. University of Chicago Press, pp. 159–186.

Bento, A., Freedman, M., Lang, C., 2014. Who benefits from environmental regulation? evidence from the clean air act amendments. Review of Economics and Statistics (0).

Blackorby, C., Russell, R. R., 1989. Will the Real Elasticity of Substitution Please Stand Up? (A Comparison of the Allen/Uzawa and Morishima Elasticities). The American Economic Review 79 (4), 882–888.

Bushnell, J. B., Mansur, E. T., May 2011. Vertical Targeting and Leakage in Carbon Policy. American Economic Review 101 (3), 263–267.

Campbell, H. F., 1991. 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), 262–274.

CFR, 2015. 40 cfr part 58, appendix d to part 58 - network design criteria for ambient air quality monitoring.

Chay, K. Y., Greenstone, M., 2003a. Air quality, infant mortality, and the clean air act of 1970. Tech. rep., National Bureau of Economic Research.

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

Chay, K. Y., Greenstone, M., 2005. Does Air Quality Matter? Evidence from the Housing Market. Journal of Political Economy 113 (2).

Currie, J., Neidell, M., 2005. Air Pollution and Infant Health: What Can We Learn from California's Recent Experience? The Quarterly Journal of Economics 120 (3), 1003–1030.

Davis, L. W., Kahn, M. E., 2010. International trade in used vehicles: The environmental consequences of NAFTA. American Economic Journal: Economic Policy 2 (4), 58–82.

Department of Energy, 2014. Cleaning up coal. Accessed January 16, 2014.
URL http://www.fossil.energy.gov/education/energylessons/coal/coal_cct2.html

Domike, J. R., Zacaroli, A. C. (Eds.), 2011. The Clean Air Act Handbook. American Bar Association.

Duhigg, C., 2009. Cleansing the Air at the Expense of Waterways.

EC/R Incorporated, October 1998. Stationary source control techniques document for fine particulate matter. Tech. rep., EC/R Incorporated, EPA Contract No. 68-D-98-026, Work Assignment No. 0-08.

Environmental Protection Agency, 1999. Land disposal restrictions for hazardous wastes. Tech. Rep. EPA530-F-99-043, Environmental Protection Agency.

Environmental Protection Agency, 2008. Clean Water Act Jurisdiction Following the U.S. Supreme Court's Decision in Rapanos v. United States & Carabell v. United States. Tech. rep., Environmental Protection Agency.

Environmental Protection Agency, 2011. The benefits and costs of the clean air act from 1990 to 2020 final report, rev . a. Tech. rep., Environmental Protection Agency.

Environmental Protection Agency, 2012. Factors to consider when using toxics release inventory data. Tech. rep., Environmental Protection Agency.

Environmental Protection Agency, 2013. CWA Methods. Accessed June 8, 2013.
URL `http://water.epa.gov/scitech/methods/cwa/pollutants.cfm`

Environmental Protection Agency, 2015a. 40 CFR Part 257.

Environmental Protection Agency, 2015b. 40 CFR Part 423.

Environmental Protection Agency, Undated. Air pollution control technology fact sheet. Tech. Rep. EPA -452/F-03-032, Environmental Protection Agency.

Farnsworth, D., 2011. Preparing for EPA Regulations: Working to Ensure Reliable and Affordable Environmental Compliance. Tech. rep., Regulatory Assistance Project.

Fowlie, M., 2009. Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage. American Economic Journal: Economic Policy 1 (2), 72–112.

Fowlie, M., 2010. Emissions Trading, Electricity Restructing, and Investment in Pollution Abatement. The American Economic Review 100 (3), 837–869.

Fullerton, D., Karney, D. H., 2014. Multiple Pollutants, Unregulated Sectors, and Suboptimal Environmental Policies.

Gamper-Rabindran, S., Jul. 2006. 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), 391–410.

Gamper-Rabindran, S., 2009. The Clean Air Act and volatile organic compounds: Did plants reduce their health-indexed air emissions or shift their emissions into other media?

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

Greenstone, M., 2003. Estimating Regulation-Induced Substitution: The Effect of the Clean Air Act on Water and Ground Pollution. American Economic Review: AEA Papers & Proceedings 93 (2), 442–449.

Hamilton, J., 2005. Regulation through revelation: the origin, politics, and impacts of the Toxics Release Inventory Program. Cambridge University Press.

Hanna, R., 2010. US environmental regulation and FDI: evidence from a panel of US-based multinational firms. American Economic Journal: Applied Economics 2 (3), 158–189.

Helland, E., 1998. The Revealed Preferences of State EPAs: Stringency, Enforcement, and Substitution. Journal of Environmental Economics and Management 35 (3), 242–261.

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

Hendryx, M., Conley, J., Fedorko, E., Luo, J., Armistead, M., Jan. 2012. Permitted water pollution discharges and population cancer and non-cancer mortality: toxicity weights and upstream discharge effects in US rural-urban areas. International Journal of Health Geographics 11 (1), 9.

Hertwich, E. G., Pease, W. S., McKone, T. E., 1998. Evaluating Toxic Impact Assessment Methods: What Works Best? Environmental Science and Technology.

Isen, A., Rossin-Slater, M., Walker, W. R., 2014. Every Breath You Take-Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970. Tech. rep., National Bureau of Economic Research.

Levinson, A., 2003. Environmental Regulatory Competition: A Status Report and Some New Evidence. National Tax Journal 56 (1), 91–106.

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

Missoula County Environmental Health Division, 1999. History of missoula's air quality program. Accessed October 10, 2014.
URL `http://www.co.missoula.mt.us/airquality/MissoulasAir/aqhistory.htm`

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

Raffuse, S., Sullivan, D., McCarthy, M., Penfold, B., Hafner, H., 2007. Analytical techniques for technical assessments of ambient air monitoring networks. Tech. rep., Environmental Protection Agency.

Sigman, H., 1996. Cross-Media Pollution: Responses on Chlorinated Solvent Releases to Restrictions. Land Economics 72 (3), 298–312.

Sigman, H., 2003. Letting States Do the Dirty Work: State Responsibility for Federal Environmental Regulation. National Tax Journal 56 (1), 107–122.

Sigman, H., 2005. Transboundary spillovers and decentralization of environmental policies. Journal of Environmental Economics and Management 50 (1), 82–101.

United States Code, 1990. 42 USC, Title 42, Chapter 85, Subchapter I, Part D, subpart 1, Sec. 7505a - Maintenance plans.

Vatavuk, W., van der Vaart, D., Spivey, J., 2000. Voc controls. Tech. Rep. EPA/452/B-02-001, Environmental Protection Agency.

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

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

Westbrook, M. D., Buckley, P. A., 1990. Flexible Functional Forms and Regularity: Assessing the Competitive Relationship between Truck and Rail Transportation. The Review of Economics and Statistics 72 (4), 623–630.

Zivin, J. G., Neidell, M., 2012. The impact of pollution on worker productivity. The American Economic Review 102 (7), 3652–3673.

Figures and tables

Figures

Figure 1: Pollution changes, holding output fixed, 2-input case

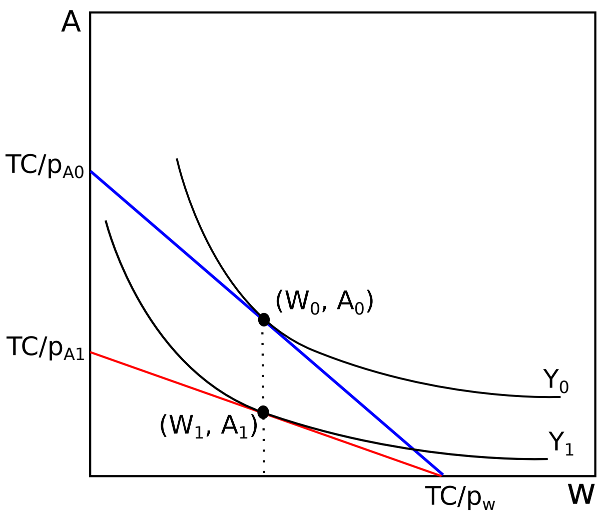

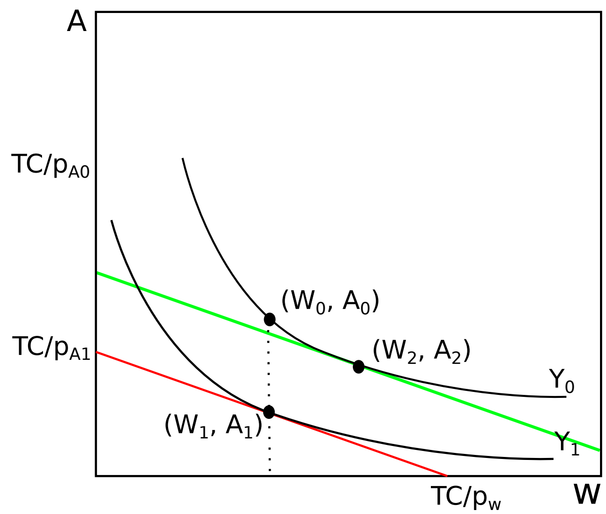

Figure 2: Residual onsite air emissions by distance from nearest non-attainment monitor

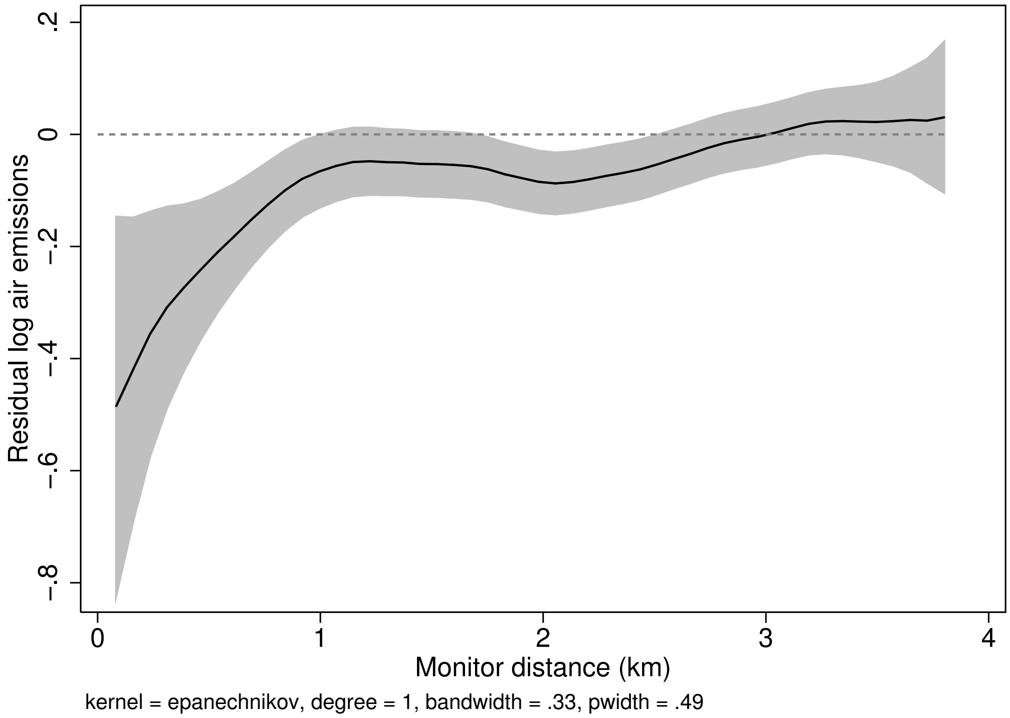

Underlying residuals from equation 4, a panel model of log air emissions (lbs) with year dummies and plant fixed effects. The fitted line represents a local linear regression estimated over residuals for plants in non-attainment counties. Shaded area is the 95% confidence interval. 3.8km is 25th percentile of distance distribution.

Figure 3: Event study estimates, onsite air emissions

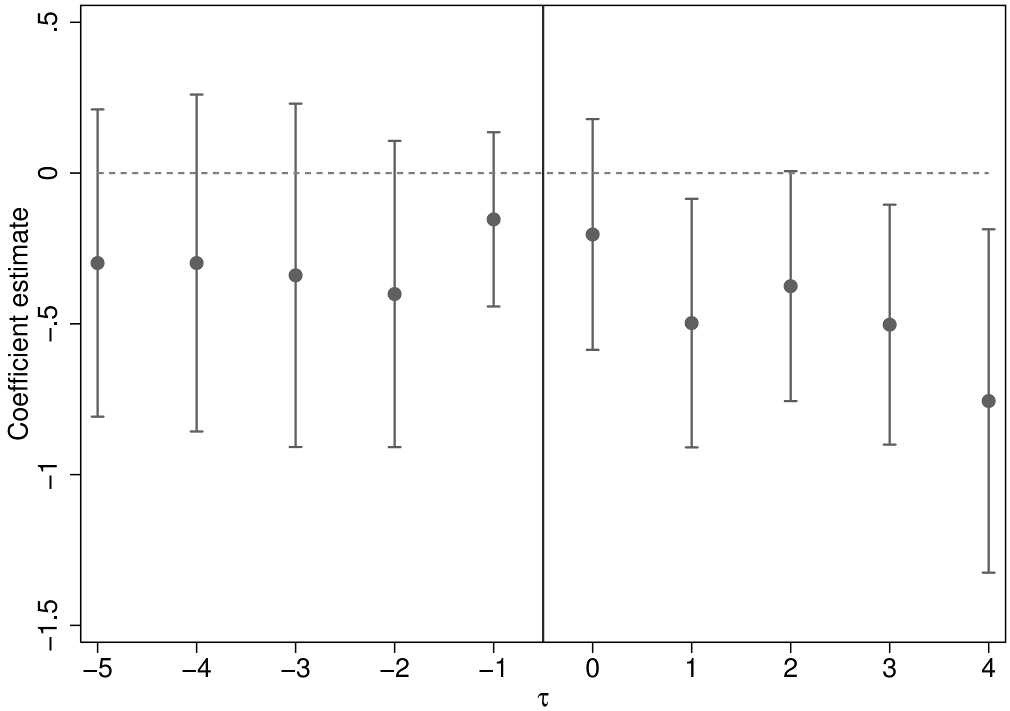

Estimates from equation 5, also reported in column 3 of Table 4. Dependent variable is log air emissions (lbs). Reference category is years for which $\tau < -5$. A county enters non-attainment in year tau=-1 and plants within ~1km of a non-attainment monitor enter treatment in the following year ($\tau = 0$). 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.

Tables

Table 1: Particulate abatement strategies

Name	Category	Description	Variable Costs	Secondary wastes
Output reduction	-	-	-	-
Reduce exhaust temp./pressure	-	Lower reaction temperature generates fewer particulates	Efficiency loss	-
Fuel switching	-	Switch to washed coal, oil, or natural gas	Added fuel cost	Coal slurry (offsite)
Process modification	-	e.g. Changing furnace type or cooling system	-	-
Flue gas conditioning	Pretreatment	Chemistry/temp./moisture modified to aid collection	Absorbant, electricity	Sulfates
Precollection	Pretreatment	Collectors use gravity/inertia to gather particles	Electricity	Solid waste
Electrostatic precipitation	End-of-pipe	Field charges particles, collected by electrode	Electricity, water	Liquid/solid waste
Fabric filters	End-of-pipe	Tightly woven fabric and dust layer trap particles	Electricity, filters	Solid waste
Wet scrubbers	End-of-pipe	Liquid (often sprayed) traps particles	Electricity, water	Liquid/solid waste
Incineration	End-of-pipe	Emissions burned at 300-2000°F, sometimes catalyzed	Fuel, catalyst	CO_2, N_2, H_2O
Ventilation	Fugitive control	e.g. Vacuum hoods, building enclosure	Electricity	Solid waste
Road paving	Fugitive control	-	Maintenance	-
Water spraying	Fugitive control	Wet down sources of fugitive emissions, e.g. coal piles	Water	Coal slurry

Sources: Department of Energy (2014); EC/R Incorporated (1998); Environmental Protection Agency (Undated); Farnsworth (2011); Vatavuk et al. (2000). Variable costs range from 33 to 100 percent of capital cost for most "end-of-pipe" abatement technologies. Incineration is typically used only for waste streams containing both PM and VOCs.

Table 2: 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 3: 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

Table 4: Effect on log air emissions

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
Treated	-0.485***	-0.407**	-0.476***	
	(0.177)	(0.174)	(0.175)	
Tau=-5				-0.266
				(0.260)
Tau=-4				-0.298
				(0.285)
Tau=-3				-0.339
				(0.290)
Tau=-2				-0.401
				(0.259)
Tau=-1				-0.154
				(0.147)
Tau=0 (1st treated year)				-0.204
				(0.195)
Tau=1				-0.498**
				(0.210)
Tau=2				-0.375*
				(0.194)
Tau=3				-0.503**
				(0.203)
Tau>=4				-0.756***
				(0.291)
State linear trends	No	Yes	No	No
NAICS*Year FE	No	No	Yes	No
Year dummies	Yes	Yes	No	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	152951	152951	152951	129283

Standard errors in parentheses

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

Estimates in columns 1-3 correspond to equation 6, while estimates in column 4 correspond to equation 5. Dependent variable is log air emissions (lbs). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. Sample size falls in the event study specification because some plants are already in non-attainment counties in the first year of my data. Because treatment is permanent and my specifications include plant fixed effects, such plants do not help identify the treatment effect in columns 1 through 3. Thus the identifying variation is the same in all four columns, despite the different sample size in column 4.

Table 5: Effect on log emissions, other media

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	0.719**	0.192	-0.00728	-0.0677	-0.0949	-0.770	-0.153
	(0.337)	(0.610)	(0.682)	(0.248)	(0.272)	(0.528)	(0.229)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	39592	18989	9755	51294	71048	43220	91806

Panel B: State 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	0.774***	-0.00749	-0.149	0.0269	-0.0431	-0.773	-0.0630
	(0.297)	(0.564)	(0.592)	(0.245)	(0.273)	(0.527)	(0.243)
State 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	39592	18989	9755	51294	71048	43220	91806

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

Estimates correspond to equation 6. Dependent variable is log emissions (lbs), with the medium 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. "Onsite other" emissions include waste piles, leaks, and spills.

Table 6: Effect on log emissions ratios, other media

Panel A: Main specification	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
Treated	1.022***	0.492	-0.373	0.578*	0.262	-0.944	0.119
	(0.374)	(0.647)	(0.608)	(0.317)	(0.284)	(0.593)	(0.420)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	36264	17940	8647	41876	60947	35720	75108

Panel B: State linear trends	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
Treated	1.020***	0.295	-0.405	0.616*	0.214	-1.073*	0.109
	(0.355)	(0.550)	(0.559)	(0.315)	(0.269)	(0.592)	(0.450)
State 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	36264	17940	8647	41876	60947	35720	75108

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

Estimates correspond to equation 7. Dependent variable is log emissions ratio (lbs), 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 7: Effect on log emissions, 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	-0.0437	1.160***	-1.516**	1.364***	0.459	-0.347	-0.603	-0.0658
	(0.203)	(0.406)	(0.602)	(0.140)	(0.303)	(0.381)	(0.625)	(0.255)
Observations	79884	18289	3431	2701	38279	40686	25871	69938
Wood products	-1.056***	-0.178	1.145*	-0.350	-0.709**	0.0455	-1.037	-0.438
	(0.322)	(0.556)	(0.628)	(0.400)	(0.326)	(0.409)	(0.932)	(0.468)
Observations	53166	13770	8132	2664	10116	22419	12180	16359
Utilities	-0.858	0.678	1.747***	-2.553**	-1.911***	0.751	-0.284	-0.138
	(0.799)	(0.567)	(0.498)	(1.173)	(0.377)	(1.052)	(0.541)	(1.580)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8266	4940	5094	1295	538	3742	2255	1877

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

Includes the three 2-digit NAICS industries with the largest treated sample sizes. All columns based on equation 6. Dependent variable is log emissions (lbs). 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 & 2-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other treated plants	0.158**	0.153**		
	(0.0672)	(0.0667)		
Count other treated			0.129**	0.124**
			(0.0511)	(0.0511)
State linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	128465	128465	128465	128465

Standard errors in parentheses

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

Estimates correspond to equation 8, where "other treated plant" is a treated plant within the same firm and 2-digit NAICS code. Dependent variable is log air emissions (lbs). 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. Parent firm identifiers come from TRI data.

Table 9: Placebo effect on log emissions

	(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	0.0758	0.142	0.546	-0.415	-0.944	-1.175	-0.381
	(0.210)	(0.428)	(0.667)	(0.426)	(0.646)	(0.828)	(0.304)
Treated*air emissions	0.790**	0.202	-0.361	0.0909	0.0667	-0.615	-0.0990
	(0.349)	(0.639)	(0.632)	(0.214)	(0.233)	(0.496)	(0.233)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	41520	20100	10773	61449	82066	51509	117844

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

Estimates correspond to equation 6, but 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 10: Placebo leakage effect, within firm & 2-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other placebo plants	0.0127	0.0180		
	(0.0330)	(0.0329)		
Count placebo plants			-0.0141	-0.0119
			(0.00939)	(0.00924)
State linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	128465	128465	128465	128465

Standard errors in parentheses

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

Estimates correspond to equation 8, but using variables based on placebo treated plants: plants within the same firm and 2-digit NAICS code, located in non-attainment counties, but farther than 8km from the nearest non-attainment monitor. Dependent variable is log air emissions (lbs). 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. Parent firm identifiers come from TRI data.

Appendix A

A.1 Alternative theoretical models

A.1.1 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:

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

Taking FOCs, one obtains an optimality condition:

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

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:

$$\ln\left(\frac{W^*}{A^*}\right) = \frac{1}{1-\rho}\ln\left(\frac{c_W}{c_A}\right) + \frac{1}{1-\rho}\ln\left(\frac{p_A + 0}{p_W}\right) \tag{10}$$

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:

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

If ρ is finite and $\rho \leq 1$, then the coefficient on the last term is positive. The positive shadow price λ causes an increase in the last term. Theory then predicts an increase in the ratio of water to air pollution $\frac{W^*}{A}$. This prediction is the same as the one from the model treating CAA non-attainment as a relative price change. The 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 11 yields:

$$\ln\left(W^*\right) = \frac{1}{1-\rho}\ln\left(\frac{c_W}{c_A}\right) + \frac{1}{1-\rho}\ln\left(\frac{p_A+\lambda}{p_W}\right) + \ln\left(\overline{A}\right) \tag{12}$$

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

A.1.2 Three production inputs

Suppose a nested CES production function, including a third input L. As in Fullerton and Karney (2014), this input may be regarded as labor or as a composite of non-pollution inputs like labor, land and capital. The firm problem then becomes:

$$\max_{A,W,L} p_O c_2 \left\{ c_P \left[c_1 \left(c_A A^\rho + c_W W^\rho \right)^{1/\rho} \right]^\theta + c_L L^\theta \right\}^{1/\theta} - p_A A - p_w W$$

The constants c_1, c_2, c_A, c_W, c_P and c_L reflect a firm's technology. Taking first order conditions on A and W, then dividing, yields:

$$\frac{p_O}{p_O}\frac{c_2}{c_2}\frac{\{\cdot\}^{1/\theta-1}}{\{\cdot\}^{1/\theta-1}}\frac{c_P}{c_P}\frac{[\cdot]^{\theta-1}}{[\cdot]^{\theta-1}}\frac{c_1}{c_1}\frac{(\cdot)^{1/\rho-1}}{(\cdot)^{1/\rho-1}}\frac{c_W}{c_A}\frac{W^{*\rho-1}}{A^{*\rho-1}} = \frac{p_W}{p_A}$$

This produces the optimality condition presented in Section 3.

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

Intuitively, this is because the firm substitutes over the air-labor and water-labor input pairs in the same way, so changes in the third factor do not affect the ratio of A and W. Under my multiplicative separability assumption, the omission of output (and other inputs) from my ratio regression specifications will not prevent inference of properties of the parameter $\sigma = \frac{1}{1-\rho}$. Nested CES is not the only functional form with this property, but it illustrates the character of the required assumptions in a three-input case.

A.2 Additional figures

Figure A1: Residual air emissions by distance from nearest non-attainment monitor, extended distance

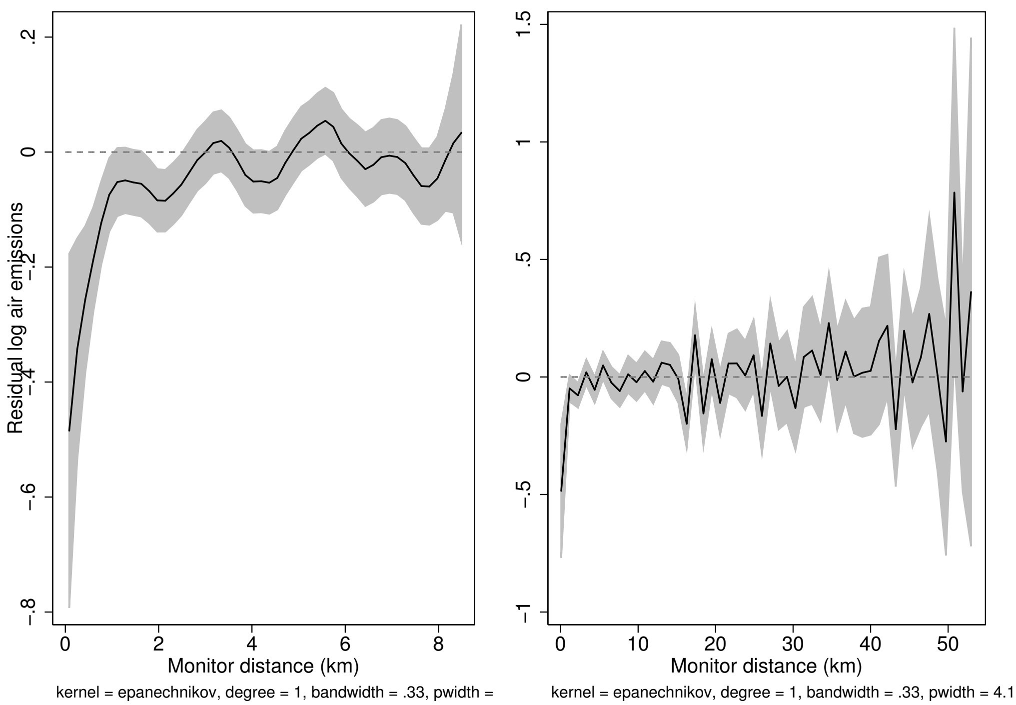

Underlying residuals from equation 4, a panel model of log air emissions (lbs) with year dummies and plant fixed effects. The fitted line represents a local linear regression estimated over residuals for plants in non-attainment counties. Shaded area is the 95% confidence interval. 8.5km is the 50th percentile of the distance distribution and 53km is the 95th.

Figure A2: PDF of NAICS6 coefficients

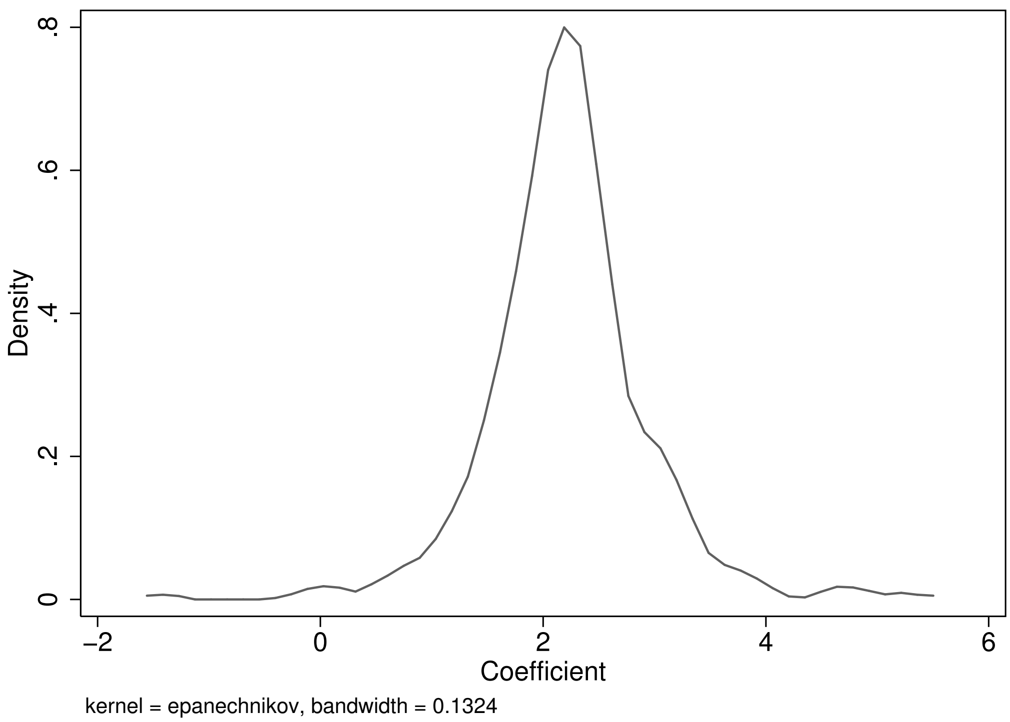

Probability distribution function of estimates from a regression of distance to nearest non-attainment monitor on year dummies and 317 dummies for six-digit NAICS codes. Regression does not include a constant. $R^2 = .6$. Industries in the right tail show no clear pattern. They include, for example, beet sugar manufacturing, prisons, and national defense.

A.3 Additional tables

A.3.1 Descriptive tables

Table A1: TRI PM descriptive statistics, by attainment status

	Attainment counties		Nonattainment counties	
	Mean	Stdev	Mean	Stdev
Onsite air	4,671.90	423,203.47	1,700.19	56,683.82
Onsite water	772.29	10,723.20	337.84	7,191.41
Onsite land	48,140.87	1,182,686.23	38,496.45	553,928.50
Offsite other	43,545.33	2,516,588.01	63,116.05	2,107,012.01
Offsite water	394.49	19,961.01	851.40	41,676.99
Offsite land	11,827.68	131,003.25	16,005.52	151,376.62
Offsite other	3,773.98	60,737.40	4,039.81	60,594.14
Recycled or treated	77,869.44	1,027,343.45	87,694.84	1,471,169.80
Dist. to nonattain monitor (km)	19.84	26.75	12.21	15.28
Treated	0.00	0.05	0.05	0.21
Observations	168191		36417	

Emissions measured in pounds. Unit of observation is a plant-year. Treated has a non-zero standard deviation in attainment county plant-years because plants remain treated even after their counties return to attainment of CAA standards. The distance to the the nearest non-attainment monitor exists for some attainment-county plant-years because of the delay between violation of CAA standards and the official non-attainment designation for a county.

Table A2: TRI PM descriptive statistics, by leakage dummy

	Other plants		Leakage plants	
	Mean	Stdev	Mean	Stdev
Onsite air	4,854	433,802	1,085	5,365
Onsite water	793	10,962	354	3,538
Onsite land	48,791	1,205,882	35,328	553,543
Offsite other	45,731	2,579,608	466	13,622
Offsite water	400	20,441	290	3,979
Offsite land	11,325	124,076	21,743	227,809
Offsite other	3,652	60,388	6,180	67,212
Recycled or treated	76,804	1,048,653	98,877	427,698
Observations	160071		8120	

Emissions measured in pounds. Unit of observation is a plant-year. "Other plants" are plants in attainment counties that have no treated plants within the same firm-year. "Leakage plants" are plants in attainment counties that have at least one treated plant within the same firm-year.

Table A3: Historical CAA particulate standards

Final rule	Type	Averaging time	Standard (µg/m3)	Form
1987	PM_{10}	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	$PM_{2.5}$	24hr	65	98th percentile, averaged over 3 years
		Annual	15	Annual arithmetic mean, averaged over 3 years
	PM_{10}	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	$PM_{2.5}$	24hr	35	98th percentile, averaged over 3 years
		Annual	15	Annual arithmetic mean, averaged over 3 years
	PM_{10}	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.

A.3.2 Monitor distance

Table A4: Monitor distance and emissions levels

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln(Dist.)	ln(Dist.)	ln(Dist.)	ln(Dist.)	ln(Dist.)	ln(Dist.)	ln(Dist.)	ln(Dist.)
Onsite air	-0.000358							
	(0.00667)							
Onsite water		0.0157						
		(0.0130)						
Onsite land			0.0524***					
			(0.0138)					
Onsite other				0.0240*				
				(0.0131)				
Offsite water					-0.0118*			
					(0.00696)			
Offsite land						-0.0135*		
						(0.00728)		
Offsite other							-0.0161**	
							(0.00654)	
Recycled or treated								-0.00663
								(0.00572)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43528	9934	3676	2611	17350	19914	14259	26960

Standard errors in parentheses

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

Estimates from a regression of distance to the nearest non-attainment monitor (in km) on log emissions. Sample is untreated plant-years. SEs clustered at the county level, which is the level of exogenous variation. Observation counts differ across columns because not all plants report emissions into all media. "Onsite other" emissions include waste piles, leaks, and spills.

Table A5: Monitor distance and emissions growth rates

	(1) ln(Dist.)	(2) ln(Dist.)	(3) ln(Dist.)	(4) ln(Dist.)	(5) ln(Dist.)	(6) ln(Dist.)	(7) ln(Dist.)	(8) ln(Dist.)
D.Onsite air	0.00443							
	(0.00332)							
D.Onsite water		-0.00578						
		(0.00640)						
D.Onsite land			0.00785					
			(0.0119)					
D.Onsite other				0.000574				
				(0.0100)				
D.Offsite water					0.00211			
					(0.00343)			
D.Offsite land						0.000647		
						(0.00343)		
D.Offsite other							-0.00152	
							(0.00305)	
D.Recycled or treated								0.000817
								(0.00316)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	39644	8870	3133	2151	15512	17194	11542	23602

Standard errors in parentheses

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

Estimates from a regression of distance to the nearest non-attainment monitor (in km) on changes in log emissions. Sample is untreated plant-years. SEs clustered at the county level, which is the level of exogenous variation. Observation counts differ across columns because not all plants report emissions into all media. “Onsite other” emissions include waste piles, leaks, and spills.

Table A6: General-equilibrium spillover test

	(1) Onsite air	(2) Onsite air
Num. treated plants (state)	-0.00206	
	(0.00325)	
Num. treated plants (state and NAICS2)		-0.00611
		(0.00623)
Year dummies	Yes	Yes
Plant FEs	Yes	Yes
Observations	151156	151156

Standard errors in parentheses

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

Estimate corresponds to equation 9. Dependent variable is log air emissions (lbs). 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. "Num. treated plants (state)" is the number of treated plants in a given state-year. "Num. treated plants (state and NAICS2)" is the number of treated plants in a given state, year, and two-digit NAICS code.

A.3.3 Air emissions

Table A7: Effect on air emissions, intrafirm leakage controls

	(1) Onsite air	(2) Onsite air
Treated	-0.480***	-0.402**
	(0.177)	(0.174)
Spillover controls	Yes	Yes
State linear trends	No	Yes
Year dummies	Yes	Yes
Plant FEs	Yes	Yes
Observations	152951	152951

Standard errors in parentheses

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

Estimates in columns 1-2 correspond to equation 6, while estimates in column 3 correspond to equation 5, but with the inclusion of a leakage control from equation 8: the number of treated plants within the same firm. Dependent variable is log air emissions (lbs). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table A8: Effect on air emissions, plants open 1993-2010

	(1) Onsite air	(2) Onsite air
Treated	-0.452**	-0.342
	(0.214)	(0.225)
State linear trends	No	Yes
Year dummies	Yes	Yes
Plant FEs	Yes	Yes
Observations	39073	39042

Standard errors in parentheses

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

Estimates correspond to equation 6. Dependent variable is log air emissions (lbs). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table A9: Effect on air emissions, alternative specifications

	(1) D.Onsite air	(2) Onsite air	(3) D.Onsite air
Treated	-0.0466*		
	(0.0274)		
Non-attainment (t-1)		-0.0783*	-0.0130
		(0.0418)	(0.00889)
Year dummies	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes
Observations	132365	152951	132745

Standard errors in parentheses

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

Estimates correspond to equation 6, but with the dependent variable replaced by the year-on-year difference in logs (the growth rate) in columns 1 and 3. In columns 2 and 3, lagged county non-attainment is the treatment of interest. These specifications are similar to those used by Greenstone (2003) and Gamper-Rabindran (2009). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year.

Table A10: Effect on air emissions, varying threshold distance

	(1) <.97km	(2) <1.02km	(3) <1.07km	(4) <1.12km	(5) <1.17km
Treated	-0.560***	-0.516***	-0.485***	-0.417**	-0.363**
	(0.196)	(0.186)	(0.177)	(0.171)	(0.162)
Year dummies	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes
Observations	152951	152951	152951	152951	152951

Standard errors in parentheses

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

Estimates correspond to equation 6. Dependent variable is log air emissions (lbs). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. The threshold used elsewhere throughout the paper is 1.07km, the distance at which one can no longer reject a null hypothesis of a zero effect on air emissions.

A.3.4 Cross-media substitution

Table A11: Effect on log non-air emissions, controlling for log air emissions

	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
Treated	0.774**	0.382	-0.648	-0.0463	0.0159	-0.885	-0.0956
	(0.353)	(0.676)	(0.693)	(0.237)	(0.239)	(0.558)	(0.260)
Log air emissions	0.220***	0.369***	0.310***	0.194***	0.257***	0.170***	0.164***
	(0.0139)	(0.0267)	(0.0450)	(0.0129)	(0.0131)	(0.0167)	(0.0103)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	36264	17940	8647	41876	60947	35720	75108

Standard errors in parentheses

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

All columns based on equation 7, with air emissions as a right-hand side control rather than a denominator for the dependent variable. Dependent variable is log emissions(lbs), with the medium 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. "Onsite other" emissions include waste piles, leaks, and spills.

Table A12: Effect on log emissions ratios, by 2-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
Primary metals	1.268***	-1.066	1.432***	0.383	-0.467	-1.038	0.111
	(0.336)	(0.816)	(0.168)	(0.333)	(0.313)	(0.712)	(0.414)
Observations	16528	3094	2394	30966	34496	21164	56274
Wood products	0.411	1.083*	-0.381	0.835	0.980**	-0.807	0.386
	(0.800)	(0.576)	(0.564)	(0.520)	(0.425)	(1.117)	(1.157)
Observations	12701	7738	2511	8752	19384	10058	14221
Utilities	1.317*	2.569***	-2.183***	0.437	1.818	-0.762	-0.767
	(0.796)	(0.548)	(0.729)	(0.299)	(1.122)	(0.526)	(1.155)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4914	5075	1294	523	3726	2238	1854

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

Includes the three 2-digit NAICS industries with the largest treated sample sizes. All columns based on equation 7. Dependent variable is log emissions ratio (lbs), 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 A13: Effect on log emissions, by 3-digit NAICS code

	Onsite air	Onsite water	Onsite land	Onsite other	Offsite water	Offsite land	Offsite other	Recycled or treated
Primary metals	0.0950	0.984**	-1.882***		0.749**	-0.571	-0.556	0.196
	(0.238)	(0.431)	(0.522)		(0.357)	(0.408)	(0.861)	(0.389)
Chemicals	-1.145***	-0.932			-0.857**	-0.197	-1.447	-0.498
	(0.431)	(0.594)			(0.348)	(0.530)	(1.121)	(0.599)
Fabricated metals	-0.456	2.157***	-1.375***		0.314	-0.328	-1.133*	-0.595**
	(0.341)	(0.608)	(0.400)		(0.550)	(1.028)	(0.653)	(0.270)
Nonmetallic mineral products	-0.800	-0.121	2.053***	-0.275		0.495**	0.0232	-2.028***
	(0.523)	(0.196)	(0.522)	(0.467)		(0.223)	(0.211)	(0.169)
Transportation equipment	-0.442		-0.644*		0.169	-0.0898	1.344	0.408
	(0.936)		(0.381)		(0.227)	(0.123)	(1.391)	(1.055)
Petroleum and coal	-1.326	1.125			0.558*	0.558	0.170	-0.294
	(0.815)	(0.943)			(0.308)	(0.564)	(0.890)	(1.079)
Utilities	-0.858	0.678	1.747***	-2.553**	-1.911***	0.751	-0.284	-0.138
	(0.799)	(0.567)	(0.498)	(1.173)	(0.377)	(1.052)	(0.541)	(1.580)
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. All columns correspond to equation 6. Dependent variable is log emissions (lbs). 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 A14: Intra-firm leakage effect on non-air emissions, within firm & 2-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	0.220*	0.0770	0.544	0.0990	-0.203*	-0.00218	0.0403
	(0.114)	(0.250)	(0.367)	(0.118)	(0.110)	(0.148)	(0.0922)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	34080	17066	8139	41708	60066	34979	76482

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

Estimates correspond to equation 8, where "other treated plant" is a treated plant within the same firm and 2-digit NAICS code, but dependent variable is log emissions ratio (lbs). 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. Parent firm identifiers come from TRI data.

Table A15: Effect on toxicity-weighted log emissions

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	1.606	-0.468	1.151***	1.294*	0.373	-0.407	-0.206
	(1.048)	(0.304)	(0.282)	(0.745)	(0.347)	(1.067)	(0.397)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18115	6467	4576	29688	38171	20880	64721
Panel B: State 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	1.808*	-0.598**	0.942**	1.304*	0.522	-0.536	-0.138
	(0.968)	(0.274)	(0.431)	(0.696)	(0.330)	(0.998)	(0.376)
State 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	18115	6467	4576	29688	38171	20880	64721

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

Estimates correspond to equation 7. Dependent variable is log toxicity-weighted emissions (unitless), with the medium indicated atop the column. 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. EPA ingestion toxicity weights applied to all emissions.

Table A16: Effect of county non-attainment on log emissions, other media

	(1) Onsite water	(2) Onsite land	(3) Onsite other	(4) Offsite water	(5) Offsite land	(6) Offsite other	(7) Recycled or treated
County non-attainment (t-1)	-0.0948	0.0342	0.0816	-0.0731	-0.0175	0.0122	0.0495
	(0.0788)	(0.112)	(0.188)	(0.0619)	(0.0621)	(0.0948)	(0.0554)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	39592	18989	9755	51294	71048	43220	91806

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

Estimates correspond to equation 6. Dependent variable is log emissions (lbs), with the medium 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. "Onsite other" emissions include waste piles, leaks, and spills.

Table A17: Effect on onsite water emissions, varying threshold distance

	(1) <.97km	(2) <1.02km	(3) <1.07km	(4) <1.12km	(5) <1.17km
Treated	0.499	0.643*	0.719**	0.541	0.380
	(0.366)	(0.342)	(0.337)	(0.344)	(0.316)
Year dummies	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes
Observations	39592	39592	39592	39592	39592

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

Estimates correspond to equation 6. Dependent variable is log onsite water emissions (lbs). SEs clustered at the county level, which is the level of exogenous variation. Unit of observation is a plant-year. The threshold used elsewhere throughout the paper is 1.07km, the distance at which one can no longer reject a null hypothesis of a zero effect on air emissions.

A.3.5 Leakage

Table A18: Leakage effect, within firm & 6-digit NAICS code

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other treated plants	0.130	0.0932		
	(0.0965)	(0.0958)		
Count other treated			0.132	0.0963
			(0.0852)	(0.0850)
State linear trends	No	Yes	No	Yes
Year dummies	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	128465	128465	128465	128465

Standard errors in parentheses

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

Estimates correspond to equation 8, where "other treated plant" is a treated plant within the same firm and 6-digit NAICS code. Dependent variable is log air emissions (lbs). 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. Parent firm identifiers come from TRI data.

Table A19: Leakage effect, continuous firm size controls

	(1) Onsite air	(2) Onsite air	(3) Onsite air	(4) Onsite air
1+ other treated plants	0.165**		0.159**	
	(0.0680)		(0.0674)	
Count other treated		0.135***		0.130**
		(0.0517)		(0.0512)
Plants in firm	Yes	Yes	No	No
Plants in firm and NAICS	No	No	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Observations	128465	128465	128465	128465

Standard errors in parentheses

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

Estimates correspond to equation 8, where "other treated plant" is a treated plant within the same firm and 2-digit NAICS code. "Plants in firm" is a count of all plants in a given firm-year. "Plants in firm and NAICS" is a count of plants within firm-year and 2-digit NAICS code. Dependent variable is log air emissions (lbs). 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. Parent firm identifiers come from TRI data.

Table A20: Leakage effect, varying threshold distance

	(1)	(2)	(3)	(4)	(5)
	<.97km	<1.02km	<1.07km	<1.12km	<1.17km
1+ other treated plants	0.163**	0.178***	0.158**	0.130**	0.131**
	(0.0794)	(0.0675)	(0.0672)	(0.0638)	(0.0631)
Year dummies	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes
Observations	128465	128465	128465	128465	128465

Standard errors in parentheses

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

Estimates correspond to equation 8, where "other treated plant" is a treated plant within the same firm and 2-digit NAICS code. Dependent variable is log air emissions (lbs). 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. Parent firm identifiers come from TRI data. The threshold used elsewhere throughout the paper is 1.07km, the distance at which one can no longer reject a null hypothesis of a zero effect on air emissions.