The B.E. Journal of Economic Analysis & Policy

Contributions

Volume 10, *Issue* 1 2010 *Article* 19

The Impact of Inspections on Plant-Level Air Emissions

Rema Nadeem Hanna* Paulina Oliva†

*Kennedy School of Government, Harvard University, rema_hanna@ksg.harvard.edu
†University of California at Santa Barbara, oliva.paulina@gmail.com

Recommended Citation
Rema Nadeem Hanna and Paulina Oliva (2010) "The Impact of Inspections on Plant-Level Air Emissions," *The B.E. Journal of Economic Analysis & Policy*: Vol. 10: Iss. 1 (Contributions), Article 19.
Available at: http://www.bepress.com/bejeap/vol10/iss1/art19

The Impact of Inspections on Plant-Level Air Emissions*

Rema Nadeem Hanna and Paulina Oliva

Abstract

Each year, the United States conducts approximately 20,000 inspections of manufacturing plants under the Clean Air Act. This paper compiles a panel dataset on plant-level inspections, fines, and emissions to understand whether these inspections actually reduce air emissions. We find plants reduce air emissions by fifteen percent, on average, following an inspection under the Clean Air Act. Plants that belong to industries that typically have low abatement costs respond more strongly to an inspection than those who belong to industries with high abatement costs.

KEYWORDS: environmental economics, EPA inspections

*We are in debt to John Borton, Mamie Miller and Phyllis Slaherty at the Environmental Protection Agency for providing key insights into the inspection process. We thank Samuel Stolper for research assistance. We also would like to thank the editor, three anonymous referees, Michael Greenstone, and John Bellows for their incredibly helpful comments. Any errors or omissions remain our own.

I. INTRODUCTION

Each year, the Environmental Protection Agency (EPA) and state governments spend roughly $40 million conducting inspections of manufacturing plants to determine their compliance with U.S. environmental regulations.[1] The effectiveness of these expenditures depends on whether firms reduce emissions in response to an actual inspection or to the threat of an inspection. In this paper, we construct a panel dataset of plant-level inspections, fines and emissions to analyze the effects of U.S. environmental regulatory activities. Specifically, we conduct an event analysis to determine how plants' emissions respond to inspections.

Industrial air emissions have been steadily falling over the last two decades. At the same time, plant-level inspections, which are mandated under the Clean Air Act, have risen slowly from 16,000 in 1985 to 23,000 in 2000. Despite the large role of inspections in U.S. and global environmental policy, the question of whether or not inspections have contributed to falling emissions remains unanswered. Inspections might improve the environmental performance of manufacturing plants through two mechanisms. First, plants may remain in compliance with regulations if both the *threat* of an inspection and the cost of non-compliance is high. Second, an *actual* inspection may reduce plant emissions if the recommendations that come from an inspection cause the plant to improve performance in order to avoid fines and other potential penalties.

On the other hand, for numerous reasons, we might expect inspections to have little long-run impact on environmental performance. The most obvious possible reasons are that the costs of non-compliance with environmental regulation could be set too low or that the inspections fail to catch the worst violations.[2] These scenarios are plausible. For example, a 1998 audit of regional and state air quality enforcement programs warned about the weaknesses in the identification and reporting of air standards violators. According to the report, these weaknesses are due to the fact that states either do not want to report violators or that the inspections are inadequate to detect them (Mid-Atlantic Audit Division Philadelphia, 1998). Moreover, the effect on overall emissions levels could be low if firms increase emissions through other media as a result of lowering air emissions. In particular, manufacturing plants may convert air

---

[1] This approximate cost is estimated from data obtained by the authors from the Environmental Protection Agency. The average cost of an inspection ranges from $2000 to $5000. To compute a conservative estimate of the total cost of the inspections program, we multiplied the minimum cost of an inspection ($2000) by the average number of inspections per year (20,000) between 1985 and 2001.

[2] Of course, even if the inspections have no financial bite, simply monitoring the plants may or may not result in lower emissions. In the context of education, however, Duflo, Hanna, and Ryan (2007) show that much of the effect of monitoring and incentives programs is driven by the incentive component of these types of programs.

pollutants into waste water and solids, which require treatment and disposal either on- or off-site. If the treatment and disposal process is inadequate, emissions may simply be shifted to other media as a result of a reduction in air pollution. As a result of this theoretical ambiguity, the impact of U.S. inspections on overall levels of emissions through air and non-air media is ultimately an empirical question.

In this paper, we study the relationship between plant-level inspections under the Clear Air Act and plant emissions in the years following an inspection. Few studies have empirically analyzed the impact of inspections on the environmental performance of polluters, primarily due to the lack of available data of plant-level environmental inspections, but also due to the difficulty of establishing causality in the relationship between inspections and emissions levels. We improve upon both the data limitations and endogeneity problems that have plagued the previous literature by analyzing data from a 17 year panel dataset of plant level inspections and emissions for about 17,200 manufacturing plants. Specifically, we exploit the panel nature of these data to estimate the causal impact of an actual inspection. The primary difficulty in establishing a causal relationship between inspections and emissions is that a plant with higher emissions levels may be more likely to be targeted for an inspection. Thus, a cross-sectional approach comparing plants that were inspected to plants that were not inspected might find that inspected plants have high emissions, and therefore underestimate the impact of an inspection on emissions. Instead, we use an event study research design. This design allows us to exploit the variation in the timing of inspections to understand whether, after controlling for fixed plant-specific and time-varying characteristics, a plant reduces its emissions levels in the period immediately after an inspection.

The results of the event study may reflect both the direct effect of a past inspection as well as the effect of the threat of a future inspection. For example, a plant that is inspected in one year may be reacting to both that inspection and to the high probability of inspection in its region that year. Keohane, Mansur and Voynov (2009) have shown that the mere threat of an EPA lawsuit had a significant effect on the emissions of electric utilities, and Laplante and Rilstone (1996) provide evidence that just the threat of an inspection reduces the emissions of plants in the pulp and paper industry. In order to separate the effect of an actual inspection on a plant's behavior from the effect of the threat of inspection, we estimate and control for the predicted probability of inspection.

In the cross-section, we find that plants with more frequent air emissions inspections tend to have higher air emissions rates. This result is unsurprising, since plants with higher emissions levels are likely to be targeted for inspections. However, when we control for plant-level heterogeneity, we find that plants significantly reduce their air emissions in the four years following an inspection.

The magnitude of the effect is quite large: an actual inspection reduces the air emissions of a plant by 15 percent, with little to no resulting increase in emissions from other media. Interestingly, we find no effect of the threat of inspection on plant emissions. Overall, these findings suggest that the estimated return to inspections in terms of emissions averted is roughly 1.8 to 4.5 pounds per dollar. These results are robust to specifications that include the predicted probability of inspection.

In Section II, we describe the EPA's regulatory activities regarding air inspection, and also provide a short literature review. We describe our data and empirical strategy in Section III. In Section IV, we present the results, and we provide an interpretation of the results in Section V. We conclude in Section VI.

II. BACKGROUND

United States Regulatory Activities

Under the Clean Air Act of 1963, the EPA and state governments have the authority to inspect manufacturing plants in order to monitor their emissions levels. State inspections represent a majority—97 percent—of the total inspections, according to the data used in this paper. Regional EPA offices conduct additional inspections as a tool for targeting certain sectors or as a response to a major reporting problem. All inspections are unannounced, regardless of who conducts the inspection or the frequency of a plant's inspections.[3]

The frequency and priority of inspections for a given plant are based on an EPA inspection-targeting model—a computational model that produces a ranking of sources to be inspected within the estimated resource available. The model uses as inputs source-specific targeting data such as plant emissions, compliance information, and air quality factors. When specific input information is not available, the inspection frequency guide (IFG) determines the inspection rate for a plant. According to the IFG, Class A1 sources (facilities that have the potential to emit 100 tons of an air contaminant per year) must be subject to a Full Compliance Evaluation (FCE) annually, whereas Class A2 sources (potential emitters of 100 tons per year, not considering air-cleaning devices) must undergo a FCE biannually. Alternative inspection frequency schemes may be proposed by the state as long as their cost is not lower and all Class A2 sources are inspected at

[3] Some plants, however, belong to a category of plants that are inspected each and every year (2.6 percent of our sample). Despite the fact that these are unannounced, we assume that the plant managers recognize this pattern. Thus, the annual inspection becomes essentially a fixed characteristic of that firm, which renders it impossible to estimate the effect of any one inspection. We drop plants that fall into this category from the subsequent analysis.

least once every five years (Senate Bill No. 866 and EPA Compliance Monitoring Strategy, 1989). An FCE consists of "a review of all required reports and the underlying records; an assessment of air pollution control devices and operating conditions; observing visible emissions; a review of facility records and operating logs; an assessment of process parameters, such as feed rates, raw material compositions, and process rates; and a stack test if there is no other way to determine compliance with the emission limits" (EPA Compliance Monitoring Strategy, 1989). The EPA also performs Partial Compliance Evaluations (PCE), which include some form of emissions monitoring, and may include one or more, but not all, of the above listed tasks. In this paper, we consider both full and partial evaluations to be inspections.

The duration of an inspection varies widely among facilities. For small plants, the inspection is usually completed in a single day. For some plants, particularly refineries, it is not however uncommon for an inspector to be stationed on-site for the whole year. Not surprisingly, given this variation in the duration of inspections, the cost for a typical inspection also varies, from about $2,000 to $5,000.[4]

Anecdotal evidence suggests that some plants try to comply with EPA recommendations following an inspection, especially when the inspection results in a fine. For example, Shenango Incorporated, a coke and chemicals company, failed nine out of fifteen inspections for leaking oven doors and faced more than $100,000 in fines in 2001. As a result, the plant agreed to install equipment designed to remove sulfur from emissions and convert the sulfur for use in other processes.[5] In other cases, however, actions may take awhile to be implemented because complete information on the plant emissions may require comprehensive investigations that last several years. For example, Columbus Steel Castings in Ohio was shown in 2005 to be a likely source of air pollution according to the Ohio EPA, but the order to eliminate air pollution problems still had not been issued as of 2007 because the report was not yet complete.[6]

Literature Review and Contributions

A few key studies have expanded our understanding of the impact of environmental monitoring on plant emissions. Many of these studies have focused on the paper industry and have predominantly found that inspections and increased enforcement result in higher compliance (Magat and Viscusi, 1990; Laplante and Rilestone, 1996; Nadeau, 1997; Gray and Shadbegian, 2005; Gray and Shadbegian, 2004; and Shimshack and Ward, 2008). Outside the paper

[4] Author estimates from data on the cost components of an inspection provided by the EPA.
[5] The Associated Press State and Local Wire, Pittsburgh, May 19, 2001.
[6] The Columbus Dispatch, Ohio, February 11, 2007.

industry, Eckert (2004) finds that inspections deter violations for petroleum storage sites, but that this effect is small in magnitude.

Our paper builds on the existing literature by studying the impact of inspections on plant behavior. We make two key contributions. First, while the samples of plants in previous studies have tended to be both small in size and constrained to a given industry, we have carefully constructed a complete database of the inspections history and emissions behavior for nearly 17,200 plants over 17 years in multiple industries. The size and scope of our dataset gives our results greater external validity, and the completeness of the data allows us to investigate both how industry characteristics affect the probability of inspection and how the effect of inspections varies across industry characteristics.

Second, the panel dataset can be used to exploit both the cross-sectional and time-series variation in inspections and emissions in order to estimate the causal effect of an inspection event. The event study methodology is an improvement over previous studies, which have relied on instrumental variables. These instruments have included lagged emissions, plant size, or changes in the productive capability of a plant, which only meet the exclusion restriction under very restrictive assumptions. For example, while we expect changes in the productive capability of a firm to affect the probability of inspection, we also believe that it can affect plant level emissions directly as well. In contrast, the event study research design allows us to control non-parametrically for plant-specific and time-varying unobserved determinants of inspections. Hence, the causal effect is identified under the plausible assumption that no other plant-specific events coincide with the timing of the inspection and are also correlated with plant's emissions.

III. DATA AND EMPIRICAL STRATEGY

This paper brings together a variety of data sources to understand the relationship between the Clean Air Act inspections and plant-level emissions. This section describes the sources and structure of the data, as well as our empirical strategy.

Data and Descriptive Statistics

We have constructed a panel dataset of plant-level emissions, inspections and fines from two key sources. Panel data on air, land and water emissions come from the Toxic Releases Industries (TRI).[7] Information on Clean Air Act (CAA) inspections and fines comes from the Integrated Data for Enforcement Analysis

[7] Most air pollutants in the TRI can be classified as part of the CAA pollutants (Greenstone, 2003). Therefore, the TRI provides a reasonable outcome measure in which to study the effect of the inspections program.

(IDEA). Both the TRI and the IDEA data include a Facility Registry System (FRS) identification number of individual plants, which allows us to merge the emissions and inspections data.

The TRI data include emissions for all facilities that are required to report under the Emergency Planning and Community Right-to-Know Act of 1986 (EPCRA) and Pollution Prevention Act of 1990. Specifically, plants are required to report their emissions for 21 to 30 chemical categories (275 to 523 substances) if they manufacture or process more than 25,000 pounds of a regulated substance per year or if they use more than 10,000 pounds of one of the regulated substances per year in the production process.

For the purposes of this paper, we reorganized the TRI data in several ways. First, we restricted our attention to the manufacturing sector (SIC digits 20 – 40), which comprises 93 percent of the records in the TRI. We impose this restriction because the inspections process may be very different for non-manufacturing plants. Second, the list of substances on which facilities need to report has changed over time for a variety of reasons, including concerns about new chemicals. Thus, to ensure consistency in our measure of air emissions, we analyzed only the 243 substances that were reported in all years between 1987 and 2001. Third, the TRI only includes a plant-year observation if the plant used regulated substances above the threshold in that particular year. An analysis using only these data would fail to capture, for example, a firm that experienced an inspection and reduced its emissions to zero, and would therefore generate an upwardly-biased estimate of the effect of regulation. To remove this potential bias, we completed the panel: if the plant was captured in the TRI in previous and subsequent years, we assigned the plant "zero" emissions for the missing middle years.

The CAA inspections data come from the Integrated Data for Enforcement Analysis (IDEA). These data are copied by the EPA from many different compliance and enforcement databases, so that we can query a combination of databases using a variety of criteria to obtain integrated results. We obtained the data from the Environmental Protection Agency under a Freedom of Information Act Request in 2002. We collected data on all inspections conducted from 1985 to 2001, including all FCE and PCE conducted by either the state or the regional EPA. It is important to note that many states drop the inspections records for plants that no longer exist, which creates a complication for our estimation strategy. For example, suppose we believe that the dirtiest plants are more likely to be inspected and are also more likely to shut down as a result of environmental regulations. Our analysis would then fail to capture the effect of an inspection on a plant that has recently shut down, but had reduced its emissions before shutting down. To address the resulting bias, we focus on a balanced sample of plants, i.e. plants that appear at both the beginning and at the end of our period of study.

In Table 1, we provide basic descriptive statistics for the final, balanced sample. About 78 percent of plants were inspected at least once during the sample period, and about a third of all plant-year observations have a recorded inspection (Panel A). On average, a single plant emits 56,269 pounds of pollutants into the air each year, the equivalent of 344 passenger cars.[8] In Panel B, we provide information on the fines imposed under the Clean Air Act. While the probability of paying a fine is relatively low (17 percent), the size of the fine, conditional on paying a fine, is about $89,151. This is smaller than the $148,198 that plants spend, on average, in air pollution abatement costs (Becker, 2005).

Table 1: Summary Statistics

Panel A: Inspections and Emissions	
Fraction of plants with at least one inspection	0.78
Fraction of plant-year combinations with an inspection	0.34
Average air emissions per plant (in pounds)	56,269
Panel B: Fines	
Fraction of plants with at least one fine	0.17
Average fine per inspection rate ($)	4,567
Average fine conditional on having one ($)	89,151
Panel C: Characteristics	
Fraction of plants defined as clean	0.47
Fraction of plants with low abatement costs	0.30
Fraction of plants that belong to an industry that has a low probability of inspection in 1985 to 1988	0.54

Notes:
1. This table provides sample statistics for the 17,208 plants in the balanced sample for the years 1987 to 2001.
2. Inspection and fine data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

[8] An average passenger car produces in average 108 pounds of hydrocarbons and 55.8 pounds of nitrogen oxides per year.

Finally, in Panel C, we provide information on the characteristics of the plants. Forty-seven percent of the plants in our sample belong to an industry that is defined as clean ("low polluting"). Note that we define the dirtiest industries as the 10 industries with the highest emissions per plant for 1988 at the 2 digit-SIC level. Thirty percent of the plants in our sample belong to industries with low pollution abatement costs.[9] Finally, 54 percent of the plants in our sample belong to industries that were infrequently inspected between 1985 and 1988.[10]

What Determines Inspections?

In the analysis of the impact of an inspection on plant emissions, one of the key empirical challenges is to address the possibility of selection bias. Plants with higher emissions or plants in high-polluting industries may be targeted for inspections, which would confound the estimation of the impact of inspections on emissions. In order to address this concern, it is important to understand which characteristics of a plant, industry, or region are correlated with the inspection event. We start by providing a descriptive picture that includes some sample statistics on the inspection rate across industries and areas. We then use regression analysis to pinpoint which factors are correlated with a plant-level inspection.

Table 2 presents the yearly inspection rate per plant by 2-digit SIC industry for all manufacturing plants that report to the TRI. The paper industry (SIC 26) and tobacco industry (SIC 21) have the highest inspection rates at 65 percent per plant per year, closely followed by the petroleum industry (SIC 29) at 61 percent and the stone, clay and glass industry (SIC 32) at 60 percent. These industries tend to contain the most heavily polluting plants, and therefore we would expect that the government targets its limited resources to monitoring these industries.[11] On the other end of the spectrum, industries that are relatively clean – including fabricated metals (SIC 34), machinery (SIC 35), electronics (SIC 36), and instruments (SIC 38) – are among the least frequently inspected, with 24 percent of plants inspected in a given year.

[9] Pollution abatement cost data come from the 1999 Pollution Abatement Cost and Expenditures Survey. These costs relate to all emissions including air, water and solid waste. The high abatement cost industries are defined as Machinery (SIC 35), Rubber (SIC 30), Instruments (SIC 38), Transportation (SIC 37), Food (SIC 20), Stone, clay, glass (SIC 32), Chemicals (SIC 28), Paper (SIC 26), Primary Metals (SIC 33), and Petroleum (SIC 29). All remaining industries are defined as low abatement cost.

[10] Industries that are infrequently inspected between 1985 and 1988 include: Food, Apparel, Rubber, Fabricated Metal, Machinery, Electronics, Transportation, Instruments and Miscellaneous.

[11] We compiled a ranking of the most heavily polluted industries from the 1988 TRI.

Table 2: Inspection Rate Per Industry

Paper (26)	0.65
Tobacco (21)	0.65
Petroleum (29)	0.61
Stone, Clay and Glass (32)	0.60
Furniture (25)	0.58
Printing (27)	0.52
Textiles (22)	0.49
Primary Metal (33)	0.42
Lumber (24)	0.42
Apparel (23)	0.42
Chemicals (28)	0.38
Transportation (37)	0.37
Food (20)	0.35
Leather (31)	0.33
Miscellaneous (39)	0.29
Rubber (30)	0.28
Fabricated Metal (34)	0.24
Machinery (35)	0.24
Electronics (36)	0.24
Instruments (38)	0.24

Notes:
1. This tables provides the average yearly inspection rate per plant in each 2-digit SIC industry, using data drawn from the IDEA database.

Regional differences in the inspection rate are even starker than differences across industries. Figure 1 graphs the inspection rate per region (bar graph) and the yearly air emissions per plant (line graph) by ten EPA regions. Inspection rates are generally higher in regions with higher pollution per plant, with some important exceptions. Region 8 (MT ND SD WY UT CO), for example, has the highest pollution per plant, but also has an inspection rate that is close to the average rate. On the other hand, Region 3 (PA WV VA DE MD DC) has the highest inspection rate, while its pollution per plant is below average.

In order to determine which factors are statistically related to the probability of inspection for a given plant each year, we focus on four main categories that may affect the probability that a plant will be inspected: past inspection rates, changes in anti-pollution policies, past emissions rates, and

Figure 1: Inspection Rate and Air Emissions By Region

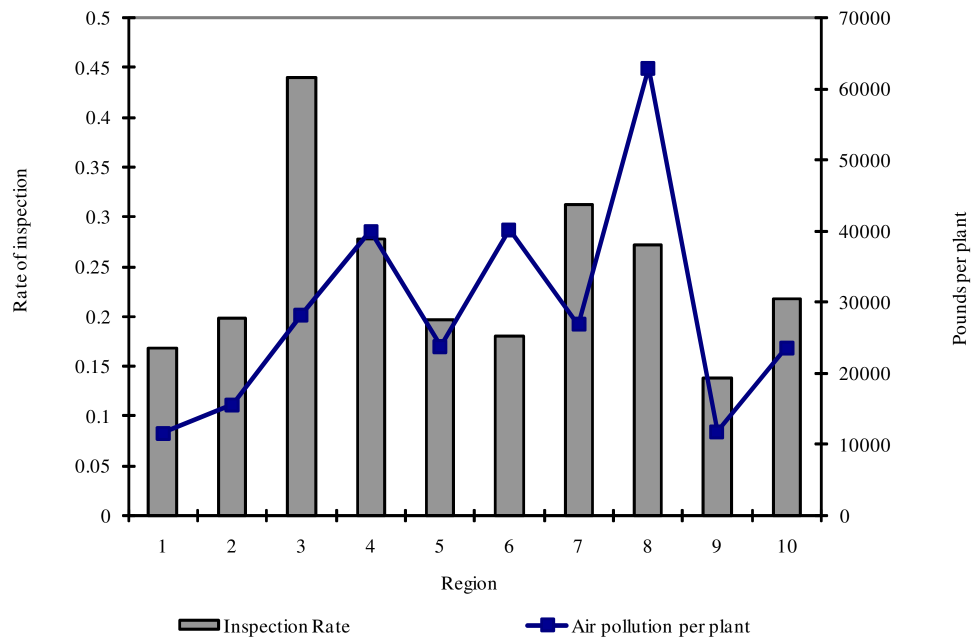

Region
1: NH ME VT MA RI CT
2: NY NJ
3: PA WV VA DE MD DC
4: KY NC TN SC MS AL GA FL
5: MN WI MI IL IN OH
6: NM OK AR LA TX
7: NE IA KS MO
8: MT ND SD WY UT CO
9: CA NV AZ
10: WA OR ID

Notes:

1. The bar graph shows the inspection rate per region in the period 1987-2001 from the IDEAS database. The line plots average air emissions per plant per year in pounds and is calculated from the Toxic Release Inventory.

socio-demographic characteristics of the region in which the plant is located.[12] Table 3 presents the results of a linear probability model in which the dependent variable is 100 if the plant was inspected in the current year, and zero otherwise.[13]

[12] These variables were chosen based on the theoretical and empirical findings of previous studies; for example, please see Gray and Shadbegian (2004) and Nadeau (1997) for a discussion of the determinants of inspections.

[13] With plant fixed effects, the probit model suffers from the incidental parameter problem. Thus, we use a linear probability model. We have estimated a specification without the plant fixed

The linear probability specification allows the estimated coefficients to be interpreted as percentages. The independent variables are listed in the first column of the table. All specifications include industry-year, industry-region and region-year fixed effects, which addresses the problem of strong correlations between industry classification and inspection frequency, and between region and inspection frequency. In Column 2 of Table 3, plant fixed effects are included in the linear probability model. Standard errors are clustered at the plant level in both specifications.

Panel A of Table 3 reports the coefficient estimates of the lagged mean EPA-inspection rate, the lagged mean state-conducted inspection rate, and the lagged mean fine at the country-industry-year level. These variables were constructed from the IDEA database. In the base specification, both the lagged EPA and state inspection rate appear to be a good predictor of current inspections (Column 1). However, when controlling for plant fixed effects, the lagged effect of EPA-inspections becomes indistinguishable from zero, and the magnitude of the effect of lagged state-inspection rate falls by more than half (Column 2).

Panel B of Table 3 provides evidence on how the nonattainment polices of the Clean Air Act Amendment affect the inspection rate. Specifically, Panel B reports the coefficient estimates of the effect of belonging to a highly polluting industry for one of the three major regulated categories of pollutants (VOC, PM, Pb) and residing in a county that is nonattainment for the corresponding pollutant (see Greenstone (2002) for a description of the nonattainment policies). Dummy variables for being located in a nonattainment county are also included in all of the regressions, but these coefficients are omitted from the table for conciseness (industry fixed effects absorb the effect of belonging to a dirty industry). Surprisingly, belonging to a top-five dirty industry for a specific pollutant and being located in a county that is in nonattainment for that pollutant has little to no effect on the probability of inspections.[14] Several factors could potentially explain this lack of correlation. First, states may not necessarily choose inspections as the policy tool to reduce emissions in nonattainment counties. At the extreme, it is possible that activities that substitute for inspections may drain funding away from inspections activities. Second, plants that shut down are not included in our inspections data; if inspections cause the dirtiest plants to shut down and the dirtiest plants are in the nonattainment counties, we will underestimate the effect of nonattainment on the inspections rate.[15]

---

effects using both the probit and the linear probability model, and the results are very similar. The results of this analysis are available from the authors on request.

[14] We also looked at specifications for the nonattainment effect that did not include the additional control variables listed in Panel A, C, and D of Table 3. The results do not significantly change, and are available upon request.

[15] Some evidence suggests that environmental policies do not speed the closure of the existing dirty plants. For example, List, Millimet and McHone (2004) find that the New Source Review

Panel C reports the coefficient estimates on two county level variables: the total number of plants in the plant's county and the lagged air emissions in a given plant's county. We include these county level explanatory variables to provide information on whether the government focuses inspections in industrial areas and whether they target plants that are located in counties that contain the dirtiest plants. We find that, as the number of plants in the county increases by 100, a plant's probability of inspection increases by about 0.07 of a percentage point (Column 2). This is suggestive evidence implying that inspections may be clustered in industrial areas. In terms of both magnitude and significance, lagged emissions in the county have little effect on the inspections rate.

Finally, the coefficients reported in Panel D provide evidence on how the socioeconomic and political characteristics of a county or a state affect the probability of inspection.[16] Specifically, Panel D reports coefficients for per capita income, county population, and political power (when including plant fixed effects in column 2, per capita income and county population are dropped from the regression, as they are collinear with the plant fixed effects). These results suggest that plants located in richer counties are less likely to be inspected (Column 1) and that having a Democratic or Independent Governor increases the rate of inspection relative to having a Republican Governor (Columns 1 and 2).

Overall, the analysis confirms that inspections are *not* independent of the specific characteristics of the plant. All of the specifications in our empirical strategy, which is detailed in the next section, include controls for industry-specific, county-specific, and plant-specific characteristics in addition to the predicted probability of inspection.

Methodology

The fact that plants with higher overall levels of emissions may be more likely to be targeted for an inspection poses a significant challenge to estimating the impact of an actual inspection on plant emissions. Because of the targeting of inspections, a cross-sectional approach that compares plants that have been inspected with those that have not been inspected might underestimate the impact of an inspection on emissions. In order to address the bias in the cross-sectional approach, we exploit variation in the timing of an inspection within a given plant by using an event study design.

---

requirements—a key component of the Clean Air Act and its subsequent amendments—did not hasten the closure of existing dirty plants. However, we are unable to test this using the data in this paper.

[16] These administrative variables were compiled from The Book of States and State Elective Officials and the Legislatures, published by the Council of State Governments.

Table 3: Exploring determinants of Inspections

	(1)	(2)
Panel A: Lagged inspection rates at the county-industry-year level		
Lagged mean EPA inspection rate	10.9363	-0.2811
	(1.3025)***	(1.2325)
Lagged mean state inspection rate	52.2228	19.5093
	(0.4546)***	(0.5257)***
Lagged mean fine rate	13.878	6.7542
	(1.1766)***	(1.1869)***
Panel B: Nonattainment county status and top polluting industry indicators		
Five top VOC × VOC Nonattainment	-0.0278	-1.7325
	(0.8183)	(1.0392)*
Five top PM × PM Nonattainment	-1.4971	-2.7027
	(1.0504)	(1.6506)
Five top Pb × Pb Nonattainment	2.6325	6.2942
	(3.4177)	(3.5117)*
Panel C: Number of plants and emissions per county		
Number of plants in county	0.0089	0.0673
	(0.0039)**	(0.0119)***
Lagged VOC air emits in county	0.3258	0.3493
	(0.0504)***	(0.0579)***
Lagged PM air emits in county	0.0455	-0.1522
	(0.2863)	(0.3101)
Lagged Pb air emits in county	23.2691	-2.3151
	(10.6466)*	(8.7033)
Panel D: County sociodemographic and political characteristics		
Per capita income at county in 1990	-0.231	
	(0.0510)***	
1990 County Population	-0.0015	
	(0.0005)***	
Democratic governor	1.0205	1.2891
	(0.3536)***	(0.3434)***
Other non-Republican governor	4.5015	4.9649
	(1.0147)***	(0.9824)***
Year-Industry Fixed Effects	X	X
Year-Region Fixed Effects	X	X
Industry-Region Fixed Effects	X	X
Plant Fixed Effects		X
N	239407	239407

Notes:

1. Each column presents the coefficients of a separate linear model for the probability of being inspected. The dependent variable is dummy variable for whether an individual plant received an inspection on a given year multiplied by 100. Coefficients, thus, should be interpreted as percentages.

2. Controls include year-industry, industry-region and year-region fixed effects.

3. Population in 1990 and income per capita are scaled by a thousand. Emissions per county is measured in millions of tons.

4. Standard errors in parentheses. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.

5. The inspections data are drawn from the IDEA database. Emissions data come from the TRI. The administrative variables were compiled from the Book of States and State Elective Officials and Legislatures, published by the Council of State Governments.

An event analysis usually exploits the combination of a panel structure of the data with a treatment (i.e. "event") that occurs at different points in time for different individuals or firms (e.g. Jacobson et al. (1993)). In our case, the event is an unannounced EPA inspection, which occurs at different times for different plants. A key benefit of this design is that we are able to include time fixed effects, which absorb the effect on emissions of all variables—such as changes in economic activity, changes in public expenditures, and so forth—that may have a simultaneous effect on the inspections and the emissions of all plants within a county or an industry. Note that the time fixed effects cannot be included in situations where the treatment starts at the same time for all plants in the treatment group, because in this case they would be collinear with the treatment variable.

A second advantage of an event analysis is that it offers a simple but compelling check on the exclusion restriction. As pointed out by Ashenfelter and Card (1985), difference-in-difference methodologies may be subject to mean reversion bias when the treatment is not random and may be triggered by a lagged high (or low) value of the dependent variable. The possibility of mean reversion is high in our setting, as inspections may be triggered by high emissions in the previous year. The coefficients on the dummy variables for several lags and leads of the treatment provide evidence on whether inspections follow particularly high-emission years. In addition, these coefficients provide evidence on whether the post inspection emissions are in fact smaller when compared to mean emissions or whether they are just reverting to the mean after a particularly high emitting year. Again, we are able to include these dummy variables because inspections occur at different times for different plants.

For the purposes of this study, an observation in the data is a plant (indexed by *i*) in a given year (indexed by *t*). The empirical specification is as follows:

$$e_{ircn} = a + \sum_{k=-4}^{4} \beta_k D_{t-k,ircnt} + w_i + u_{nt} + s_{rt} + j_{rn} + \gamma \hat{z}_{ircnt} + \varepsilon_{ircnt} \quad [1]$$

where e_{ircnt} is the air emissions of a plant *i* in EPA region *r*, county *c*, industry *n* and year *t*. The variables of interest in the above equation are denoted $D_{t-k,\, ircnt}$ for $k \in \{-4,-3,\ldots,0,\ldots 4\}$, where $D_{t-k,ircnt}$ takes the value of one if plant *i* was inspected at time *t-k*. Hence, the coefficients β_{-4} to β_{-1} measure the emissions associated with having an inspection in the future, β_0 measures the emissions associated with having an inspection in the current year, and β_1 to β_4 measure the emissions associated with inspections in the past years. In contrast with a classic cross sectional specification, where a single coefficient on the inspection indicator denotes the differential in emissions for inspected plants with respect to non inspected plants, in Specification [1] each of the β_k coefficients denotes the differential in emissions of plants inspected in year *t-k* relative to firms not inspected in year *t-k*. Hence, the estimated plant's emission response to an unannounced inspection is given by the *post-inspection* emission differentials, coefficients β_0 to β_4, relative to the *pre-inspection* emission differentials, coefficients β_{-4} to β_{-1}. We included coefficient β_0 in the post-inspection differentials group because the inspection may occur before emissions are reported within the year of inspection.

The panel structure of the data allows for the inclusion of plant fixed effects, which absorb the unobserved heterogeneity in the determinants of air emissions that are common to a particular plant. The inclusion of plant fixed effects (w_i) is important if we believe that a plant's exposure to regulation is potentially correlated with factors inherent to a plant (such as technology or size). We also include industry-by-year (u_{nt}), region-by-year (s_{rt}), and region-by-industry (j_{rn}) fixed effects. The inclusion of industry-by-year fixed effects removes shocks to air emissions that are common to all plants within an industry in a particular year. Including industry-by-year fixed effects is especially important if certain industries decreased emissions in that particular year for reasons unrelated to environmental regulation (*e.g.* new technologies that improve manufacturing in an industry and also reduce emissions). Similarly, the inclusion of region-by-year fixed effects removes the mean emissions across all plants in a region in a particular year—thereby controlling for unobserved factors that affect the emissions of all plants within a single region equally (*e.g.* a regional recession or a regional change in EPA's level of enforcement). Region-by-industry fixed effects remove the mean emissions by region and industry, which is particularly important if we believe that some industries are targeted differently across different regional EPA offices.

Lastly, we include a measure of the "inspection threat" ($\hat{z}_{ircnt}$) that a plant faces. The inspection threat is defined as the predicted probability of inspection from the specification in Table 3, Column 1.[17] By controlling for the inspection threat, we are able to address the concern that plants that have faced a higher probability of inspection over time are already more likely to have made adjustments in their emissions in order to avoid fines from a possible inspection, which would make them less likely to respond to an actual inspection.

One factor that complicates the analysis is that plants may have different frequencies of inspections. The existence of multiple inspections for a given plant may confound the effects of any single inspection. First, plants that are inspected may have been more likely to be inspected in previous years and therefore may have already changed their emissions levels as a result of previous inspections. In this case, an analysis that looks at only a single inspection may lead to an

[17] We use the specification without plant fixed effects for the predicted probability of inspection, as the permanent differences in inspection rates are absorbed by the plant fixed effects in Equation 1.

underestimation of the inspection effect. We address this concern by controlling for all other previous inspections. To see this, notice that in equation [1] multiple indicator variables ($D_{t\text{-}k}$ for multiple k's) can simultaneously take the value of one for a single plant, so $\sum_{k=-4}^{4} D_{t-k,ircnt}$ may be larger than one for a single plant-year observation. Second, having had an inspection in a given year may change the probability of a future inspection. In particular, a plant that is inspected in the current year may assume that it has a higher probability of inspection in future years. In order to understand how this could affect our analysis, we report the share of plants that receive an additional inspection one year after receiving an inspection, two years after receiving an inspection, and so on.

Table 4 reports the probability of inspection for the years before and after the current year for plants that have had an inspection in the current year (Column 1) and for plants that have not had an inspection in the current year (Column 2). The difference in these probabilities is listed in Column 3. The difference in probabilities for the current year is, by construction, 100 percent. In all years immediately before and after the current year, the difference in the probability of inspection for those that were inspected in the current year and those that were not ranges from about 38 to 52 percent. These probabilities are consistent over the time period prior to and subsequent to the inspection. They suggest that having had an inspection in the current year does not have a differential impact on the probability of inspection in subsequent years compared to previous years, and thus validate our event study approach. These probabilities also provide information on the "treatment effect": a single inspection raises the probability of inspection from about 50 percent to 100 percent.

The basic event study proposed in Equation [1] offers us a convenient check for the exclusion restriction that inspections do not happen after unusually high emitting years and that a fall in emissions that follows an inspection does not merely follow from mean reversion. For this exclusion restriction to be met, the estimates of coefficients β_{-4} to β_{-1} should not be significantly different from each other. However, interpreting the treatment effect is not straightforward. In Specification [1], the emission response in year K to an inspection in year 0 is given by the difference between the coefficient β_{K} and the last coefficient on the pre-inspection period β_{-1}. An alternative and perhaps more convenient representation of the effect of inspections on emissions can be obtained from the following specification:

$$\log(e_{ircnt}) = a + \beta_{-4to4} D_{-4to4,ircnt} + \beta_{0to4} D_{0to4,ircnt} + w_i + u_{nt},$$
$$+ s_{rt} + j_{rn} + \gamma \hat{z}_{ircnt} + \varepsilon_{ircnt} \quad [2]$$

where ε_{ircnt} is the stochastic error term and $D_{-4to4,ircnt}$ is an indicator variable that equals 1 if plant i was inspected in any of the four years following t, in t, or in the

4 years preceding *t*.[18] The main variable of interest is $D_{0to4,ircnt}$, which is an indicator variable that equals 1 if plant *i* was inspected in year *t* or in any of the preceding 4 years. The coefficient on $D_{0to4,ircnt}$, β_{0to4} ,measures the percentage change in the level of air emissions of a given plant in the five post-inspection years controlling for its *own* emissions in the years surrounding the inspection.

Table 4: Probability of Inspections in Previous and Subsequent Years

Probability of :	Conditional on Having an Inspection in Current Year (1)	Conditional on Not Having an Inspection in Current Year (2)	Difference (3)
Inspection four years after	0.579	0.193	0.386
	(0.494)	(0.395)	(0.003)***
Inspection three years after	0.604	0.183	0.421
	(0.489)	(0.387)	(0.003)***
Inspection two years after	0.641	0.167	0.474
	(0.480)	(0.373)	(0.003)***
Inspection one year after	0.660	0.161	0.499
	(0.474)	(0.368)	(0.003)***
Inspection on current year	1.000	0.000	1.000
	(0.000)	(0.000)	(0.000)***
Inspection one year before	0.703	0.183	0.521
	(0.457)	(0.386)	(0.003)***
Inspection two years before	0.689	0.187	0.503
	(0.463)	(0.390)	(0.002)***
Inspection three years before	0.657	0.205	0.452
	(0.475)	(0.404)	(0.003)***
Inspection four years before	0.637	0.218	0.419
	(0.481)	(0.413)	(0.003)***

Notes:

1. Column 1 reports the probability of inspection in each year, given an inspection in the current year. Column 2 reports this probability conditional on no inspection in the current year. The difference is reported in Column 3.
2. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.
3. The data are drawn from the IDEA database.

[18] We present a log specification so that the coefficients can be interpreted as the percentage change in emissions due to an inspection. The results where the dependent variable is the emissions level are similar to those where the dependent variable is the log of the emissions level. These results can be obtained from the authors upon request.

This coefficient can be interpreted as the average effect over the five years following an inspection plus the inspection year.

We first present the results from the estimation of Equation [1] to provide a check on the exclusion restrictions. Then, we present the results from the estimation of Equation [2] as our main specification to provide easily interpretable estimates of the effect of an actual inspection.

IV. RESULTS

Figure 2 plots the coefficients (and 95-percent confidence intervals) of the nine indicator variables D_{t-k}, from Specification [1]. We restrict the sample to the 17,041 plants (153,371 observations) with four years of leads and lags available. In Panel A, the dependent variable is air emissions, while "emissions other than air" is the dependent variable in Panel B. The vertical line divides the periods before and after the inspection. Note that a response to an inspection could be observed starting on t=0, since on a given year the inspection may occur before the emissions are reported. Hence, we include t=0 in the "after inspection" period. The corresponding regression coefficients for these figures are reported in Appendix Table 1.

As Panel A of Figure 2 illustrates, air emissions decline in the period after an inspection event. The downward trend in emissions occurs only after the inspection event and not before, suggesting that the EPA is not targeting plants that *already* display a negative trend in emissions. The fall in emissions by the third and four year after an inspection falls out of the range of the confidence intervals in the preceding period, suggesting the inspection event had a significant effect on emissions. Panel A of Figure 2 also suggests a causal effect of inspections on emissions since inspections do not follow abnormally high emission years: emissions in all four years preceding an inspection are not significantly different across years, *i.e.* emissions do not "spike" previous to the inspection year. Interestingly, we observe a slight increase in emissions other than air after the inspections event (Panel B). This may suggest a substitution towards other forms of emissions after an inspection event.

Next, we estimate Specification [2] to test for the significance of the five-year average effect of inspections given by coefficient $\beta_{0\tau o4}$. In Table 5, we present the results of this estimation. The dependent variable in Panel A is the log of air emissions, while the dependent variable in Panel B is the log of emissions other than air. In Column 1, we present a variation of Specification [2] in which we do not include the plant fixed effects and instead include county fixed effects, g_c:

$$\log(e_{ircnt}) = a + \beta_{-4to4} D_{-4to4,ircnt} + \beta_{0to4} D_{0to4,ircnt} + g_c$$
$$+ u_{nt} + s_{rt} + d_i t + \gamma \hat{z}_{ircnt} + \varepsilon_{icrnt} \quad [3]$$

Figure 2: Event Study Figures

Panel A: Air Emissions

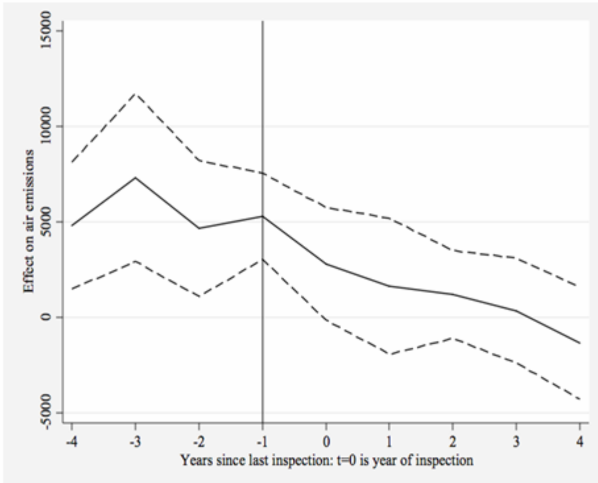

Panel B: All Emissions Other than Air

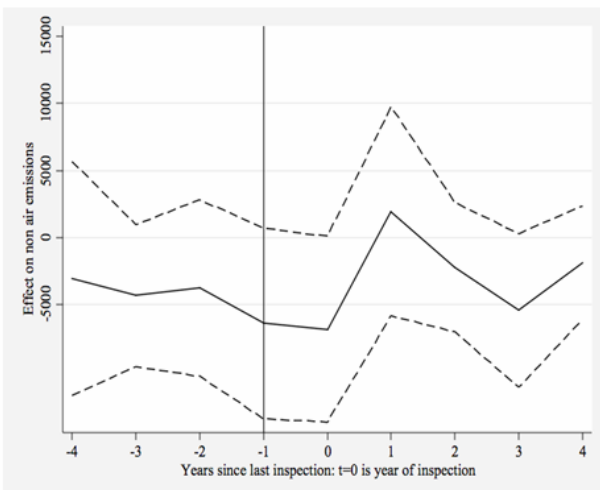

Notes:

1. These figures show the coefficients of the regression of emissions on a set of 9 indicator variables: one for each lag/lead of the inspection event. The sample includes the 153,371 observations with all leads and lags available. In Panel A the dependent variable is air emissions. In Panel B the dependent variable is all other types of emissions other than air emissions. 95 percent confidence intervals are shown in dashed lines above and below the coefficient estimates.

2. All specifications control for year-region fixed effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection.

3. Standard errors are clustered at the plant level and are in parentheses.

4. Inspection and Fine Data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

Table 5: Event Analysis of the Relationship Between Inspections and Air Emissions

	(1)	(2)
A. Air Emissions		
Inspection in any of the four previous years or in the current year	0.462	-0.153
	(0.062)***	(0.048)***
Fitted probability of inspection	0.009	0.002
	(0.002)***	(0.001)
B. All Emissions Other than Air		
Inspection in any of the four previous years or in the current year	0.347	-0.001
	(0.028)***	(0.018)
Fitted probability of inspection	0.003	0.001
	(0.001)***	(0.001)*
County Fixed Effect	X	
Plant Fixed Effect		X

Notes:

1. This table reports the results of regressions of the log of emissions on a constant, a dummy for the nine years around the inspection and a dummy for five years after the inspection including the year of the inspection itself. All specifications control for year-region fixed effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection. The coefficient on the five years after the inspection dummy and the predicted probability of inspection are reported. In Panel A the dependent variable is the log of air emissions, while in Panel B the dependent variable is the log of emissions other than air emissions.
2. Standard errors are clustered at the plant level and are in parentheses.
3. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.
4. Inspection and fine data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

In Column 2, we estimate the model controlling for plant-level heterogeneity (Specification [2]). The table reports the coefficient on the dummy for the four years after the inspection and on the inspection threat. All standard errors are clustered at the plant level.

Even after controlling for county, industrial trends, and regional trends in inspections, plants have significantly higher emissions rates after an inspection (Column 1 of Table 5). However, this model includes plants that have never been inspected, which tend to have lower emissions. Thus, this coefficient may not only capture the effect of an inspection, but also the selection effect of being chosen for an inspection in the first place. Moreover, this coefficient may simply be capturing the fact that, even after controlling for all industrial and regional trends, the EPA may target plants that they expect will have higher pollution levels in the future. Therefore, these results show the importance of controlling for plant level heterogeneity. We turn to these results in Column 2 of Table 5.

We find that, on average, a plant emits 15 percent fewer air toxins in the post-inspection period compared to the pre-inspection period. This effect is significant at the 1 percent level. We cannot reject that the coefficient on the *predicted* probability of an inspection is different from zero.[19] This may be due to the fact that we are "wiping out" a substantial part of the predicted probability through the fixed effects included in the regression. This is confirmed also by a larger coefficient on predicted probability of inspection in Column 1.

As Column 2 of Panel B illustrates, we do not find evidence that plants increase other types of emissions over the five years following the inspection. The small positive effect found in Panel B of Figure 2 seems to fade over time and the fourth year after the inspection sees no difference in air pollution relative to the pre-inspection period. This suggests that, plants may resort to increase other emissions in the short run. However, they seem eventually to choose to dispose of air toxins properly, rather than converting them to wastewater or solids that are then improperly disposed.[20]

The next table explores the effect of a fine following the inspection. After a violator is identified through an unannounced inspection and the violation is reported, the EPA and the State may continue monitoring the plant's emissions until it achieves compliance. Fines are imposed with either severe violations or a failure to comply within the compliance schedule (Consolidated Report on OECA's Oversight of Regional and State Air Inspection Programs, 1998). Given

[19] We do not take into account the estimation error of one of our regressors, the predicted probability of inspection. Generally, the standard errors of a two-step model of this sort should account for the fact that one of the regressors is estimated. However, under the null hypothesis that the coefficient of the predicted regressor is non-significant, the standard errors are consistent. Since our estimation provides evidence that supports this null hypothesis, we can claim that our standard errors are consistent. Wooldridge (2002) shows that a formal derivation of the result that "no adjustment to the standard errors is needed whenever the coefficient of the generated regressor is zero" (page 144).

[20] In Appendix Table 2, we provide the result of a robustness check where provide an alternative specification of the case with "zero" emissions in the dataset (i.e. the plant's use of chemicals fell below the reporting requirement. In Column 1, we replace the "zero" emissions with the minimum non-zero value for the plant. The coefficient estimate falls slightly, but remains relatively large in magnitude and significant. In Column 2, we interpolate the missing years. In this case, the coefficient falls to zero, suggesting that plants may temporarily decrease emissions after an inspection, but may increase them again in later years.

that fines often reflect a failure to meet the compliance schedule, we would expect that plants that are given a fine after an inspection would be less likely to have reduced emissions than those who were not penalized. Column 1 of Table 6 provides the results of a specification where the inspection variables are interacted with indicators for whether or not a plant received a fine after the inspection:

$$\log(e_{ircnt}) = a + \beta_{-4to4C} D_{-4to4,ircnt} \times Fine_i + \beta_{0to4C} D_{0to4,ircnt} \times Fine_i$$
$$+ \beta_{-4to4D} D_{-4to4,ircnt} \times NoFine_i \quad + \beta_{0to4D} D_{0to4,ircnt} \times NoFine_i$$
$$+ w_i + u_{nt} + s_{rt} + j_{rn} + \gamma\hat{z}_{ircnt} + \varepsilon_{ircnt} \quad [4]$$

The specification includes plant fixed effects, with standard errors clustered at the plant level. The results confirm that the plants that responded the most to a single inspection were those that were never fined.

On the other hand, plants that were fined showed no change in emissions after an inspection. It is likely that for these plants, the cost of cleanup was greater than the potential fines that the plants faced, so they had no incentive to clean up their emissions.

An interesting extension of this analysis is to ask whether plants that belong to industries that tend to be fined are more likely to clean up. In Column 2 of Table 6, we estimate Specification [4], replacing the actual fine variable with an indicator variable for whether the plant belonged to an industry with high fine rates in the first four years of our study. We find that plants that belong to industries that have a low probability of fine reduce emissions by about 23 percent after an inspection event. Those that belong to industries that tend to be fined also reduce emissions by about eight percent, but not significantly so. This finding has several explanations. For example, those industries with high fine levels also happen to be the ones with high abatement costs. Thus, the fine variable could be capturing other characteristics other than the fine.

Taken together, the findings reported above indicate that plants do react to an inspection event by reducing emissions. This naturally raises the question as to which types of plants respond the most to an inspection event. In Table 7, we shed some light on this question. We first estimate a specification where the inspection variables are interacted with indicator variables for whether a plant belongs to a clean or dirty industry:

$$\log(e_{ircnt}) = a + \beta_{-4to4C} D_{-4to4,ircnt} \times Clean_n + \beta_{0to4C} D_{0to4,ircnt} \times Clean_n$$
$$+ \beta_{-4to4D} D_{-4to4,ircnt} \times Dirty_n + \beta_{0to4D} D_{0to4,ircnt} \times Dirty_n + w_i + u_{nt} + s_{rt} +$$
$$j_{rn} + \gamma\hat{z}_{ircnt} \times Clean_n + \lambda\hat{z}_{ircnt} \times Dirty_n + \varepsilon_{ircnt} \quad [5]$$

The results are reported in Column 1. In Column 2, we report results from a similar specification where we interact the inspection variables with a dummy

Table 6: The Effect of Emissions from an Inspection Interacted with a Fine

	(1)	(2)
Inspection Event × Fine	0.039 (0.024)	
Inspection Event × No Fine	-0.168 (0.048)***	
Inspection Event × High Probability of Fine		-0.081 (0.071)
Inspection Event × Low Probability of Fine		-0.231 (0.064)***
Plant Fixed Effect	X	X

Notes:

1. In Column 1, we report the results of regressions of the log of emissions on a constant, and the inspection dummies (one for the nine years around the inspection and one for the five years after the inspection) interacted with an indicator variable for whether a plant received a fine in the five years following an inspection or not. We report the results on the interactions between the post-inspection dummy and the fine status of the plant. In Column 2, we replicate this analysis for whether the plant has a high or low probability of fine, based on fine levels at the state level from 1985 to 1988.
2. All specifications control for year-region effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection from Column 1 of Table 3.
3. Standard errors are clustered at the plant level and are in parentheses.
4. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.
5. Inspection and fine data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

variable for whether a plant belongs to a low abatement cost or high abatement cost industry. Finally, Column 3 reports whether plants that belong to industries that were less likely to be inspected in 1985-1988 react differently than those that belong to industries that are frequently inspected. All regressions include plant fixed effects, and thus correspond to the regression in Table 5, Column 2. All standard errors are clustered at the plant-level.

We find that the effect of an inspection on emissions varies by some, but not all, industry characteristics. Column 1 of Table 7 shows that clean and dirty plants tend to have similar reactions to an inspection: both have about a 16 percent decline in emissions after an inspection event. On the other hand, Column 2 of Table 7 shows a large impact of the inspections on plants in industries with low abatement costs and no effect on plants in industries with high abatement costs. This finding provides additional evidence that plants react most strongly to an inspection when it is relatively cheap for them to do so (when the cost of cleanup may be less than a potential fine).[21] Finally, we find evidence that plants belonging to industries that are less frequently inspected are more likely to reduce emission after an inspection event (Column 3). Specifically, plants belonging to industries that are less frequently inspected reduce emissions by about 26 percent after an actual inspection.

In sum, we find that air inspections reduce air emissions by about 15 percent, with no significant resulting increase in emissions through other media. Given that inspections have been shown to reduce emissions, the next question is whether the inspection is cost-effective or whether emissions can be cut more effectively by other means. We attempt to provide evidence on this question in the next section.

V. INTERPRETATION

The "right" level of inspections depends, of course, on the benefits of the inspections—reduced emissions, which helps to reduce damage to ozone layers, potential improvements in health, etc.—versus the costs of an inspection. While measuring the costs of an inspection is relatively straightforward, measuring the benefits is much more challenging. Many assumptions need to be made to assess benefits: What is the relationship between a reduction in emissions and human health? What value do we place on human health? And so forth. With this in mind, we attempt to put our results into context.

In our analysis, we find that plants reduce emissions in the period immediately after an inspection by about fifteen percent. Each inspection can be thought of as reducing the number of passenger cars on the road by 50 cars per year for five years. Given that a typical inspection costs about $2,000 - $5,000,

[21] In Appendix Table 3, we report the results of specification with continuous measures of the pollution level and abatement costs for the industries that the plant belongs to. As in Table 7, we do not observe a difference in reaction to the inspection event between those in cleaner or dirtier industries (Column 1). Moreover, while the direction of the coefficient is positive (as expected), we do not observe a significant difference in the reaction to the inspection event by abatement cost (Column 2). This suggests that while we see broad difference across those with high and low abatement costs, local changes in abatement costs may not be large enough for us to measure.

Table 7: The Effect of an Inspection on Emissions, by Industry Type of the Plant

	(1)	(2)	(3)
Clean * Inspection Event	-0.163 (0.068)**		
Dirty * Inspection Event	-0.166 (0.066)**		
Low Abatement Cost * Inspection Event		-0.335 (0.089)***	
High Abatement Cost * Inspection Event		-0.083 (0.055)	
Industry with Low Frequency Inspections * Inspection Event			-0.258 (0.069)***
Industry with High Frequency Inspections * Inspection Event			-0.098 (0.064)
Plant Fixed Effects	X	X	X

Notes:

1. Column 1 reports the results of regressions of the log of emissions on a constant, and the inspection dummies (one for the nine years around the inspection and one for the five years after the inspection) interacted with an indicator variable for whether a plant belongs to the cleanest or dirtiest half of industries. Column 2 reports the results of regressions of the log of emissions on a constant, and the inspection dummies interacted with an indicator variable for whether a plant has high or low emission abatement costs according to its industry. Column 3 reports the results of regressions of the log of emissions on a constant and the inspection dummies interacted with an indicator variable for whether a plant has high or low probability of inspection according to its industry. We report the results on the interactions between the post-inspection dummy and the

2. All specifications control for year-region effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection.

3. Standard errors are clustered at the plant level and are in parentheses.

4. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.

5. Inspection and fine data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

one "inspection" dollar results in a reduction of 1.8 to 4.5 pounds of air emissions per year.

It is important to note two key caveats regarding this estimate of the cost effectiveness. First, our measure only captures the emissions averted as a result of the actual cost of an inspection—the estimate does not include the potential economic costs if the inspections hinder firm activity. While we cannot measure these economic effects within the scope of our study, the literature suggests that firms do make location decisions based on the stringency of environmental conditions (see for example, Hanna, 2009; Levinson and Keller, 2002; Kahn, 1997; Greenstone, 2002). Therefore, it is possible that our measure underestimates the total costs of inspection.

Second, our measure does not quantify the benefits of the inspection program. Reductions in emissions can translate into significant health gains. A large share of the emissions studied in this paper falls under the VOC/ozone category, with a well-documented connection between ambient VOC concentrations and health. For example, Levy, et al (2001) showed a 0.5 percent increase in premature deaths per 10 μg/m^3 increase in average ozone concentrations, and also showed a 0.1 percent increase in Minor Restricted Activity Days per μg/m^3 increase in two-week average high-hour ozone concentrations. Borja-Aburto, et al (1997) found a 0.4 percent increase in premature deaths per 10 μg/m^3 increase in average ozone. However, to fully quantify the impact of the inspections policy on health outcomes would require additional assumptions as to how the emissions translate to ambient concentrations. While this is outside the scope of this study, it is an important component of our future research.

VI. CONCLUSION

This study estimates the impact of inspections under the Clean Air Act on individual plant emissions. Using an event study methodology, we find that plants significantly reduce their air emissions in the years following an inspection. In particular, an actual inspection reduces the air emissions of a plant by about 15 percent over the five years that follow an inspection, with little to no resulting increase in emissions from other media. The effect is concentrated among plants with low pollution abatement costs and plants that belong to industries that are infrequently inspected.

These results suggest several interesting directions for future work. What can be done to improve the efficiency and impact of environmental inspections? What frequency and duration of inspection events will maximize benefits relative to costs? How do variations in the level of fines per violation affect the behavior of plants? Is increasing the fine per violation a more effective policy tool than conducting more frequent inspections? In addition to shedding more light on these questions, future research should also aim to provide further insight into the social benefits and costs of inspection programs, and to compare these programs to other approaches that are designed to combat environmental degradation.

Appendix Table 1: Event Analysis of the Relationship Between Inspections and Emission Levels

	Number of years since inspection									
Specification	Inspection in the fourth following year (1)	Inspection in the third following year (2)	Inspection in the second following year (3)	Inspection in the following year (4)	Inspection in the current year (5)	Inspection in the previous year (6)	Inspection two years before (7)	Inspection three years before (8)	Inspection four years before (9)	Predicted probability of inspection (10)
Air Emissions	4,805.01	7,319.80	4,645.46	5,288.56	2,795.05	1,620.07	1,187.72	344.782	-1,355.01	-84.231
	(1,663.687)***	(2,200.032)***	(1,785.354)***	(1,126.073)***	(1,477.168)*	(1,783.423)	(1,150.256)	(1,374.634)	(1,471.314)	(73.942)
Other Emissions	-3,042.71	-4,325.87	-3,763.75	-6,406.19	-6,823.80	1,956.46	-2,208.40	-5,431.20	-1,858.01	-91.486
	(4,359.124)	(2,644.366)	(3,292.446)	(3,550.972)*	(3,462.596)**	(3,895.930)	(2,409.872)	(2,866.015)*	(2,114.417)	(131.185)

Notes:

1. Each row reports the coefficient estimates of emissions on a set of 9 indicator variables, one for each lag/lead of the inspection event. The results were used to create Figure
2. All specifications control for year-region effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection.
3. Standard errors are clustered at the plant level and are in parentheses.
4. * means significant at 5%; ** means significant at 1%.
5. Inspection and fine data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

Appendix Table 2: Event Analysis of the Relationship Between Inspections and Air Emissions, Robustness Checks

	Minimum Value Per Firm (1)	Interpolation (2)
Inspection in any of the four previous years or in the current year	-0.105	0.025
	(0.037)***	(0.044)
Fitted probability of inspection	0.001	0.001
	(0.001)	(0.001)
Plant Fixed Effect	X	X

Notes:

1. This table reports the results of regressions of the log of emissions on a constant, a dummy for the nine years around the inspection and a dummy for five years after the inspection including the year of the inspection itself. All specifications control for year-region fixed effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection. The coefficient on the five years after the inspection dummy is reported. In Column 1, missing emission values are substituted by the plant specific minimum emissions levels. In Column 2, missing values are interpolated.
2. Standard errors are clustered at the plant level and are in parentheses.
3. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.
4. Inspection and Fine Data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

Appendix Table 3: The Effect of an Inspection on Emissions, by Industry Characteristics

	(1)	(2)
Inspection Event	-0.187	-0.196
	(0.080)**	(0.059)***
Average Industry Per-Plant Air Pollution * Inspection Event	0.00018	
	(0.00054)	
Abatement Cost ($1,000) * Inspection Event		0.00039
		(0.00032)
Plant Fixed Effect	X	X

Notes:

1. Column 1 reports the results of regressions of the log of emissions on a constant, and the inspection dummies (one for the nine years around the inspection and one for the five years after the inspection) interacted with a variable average per plant pollution by industry and year. Column 2 reports the results of regressions of the log of emissions on a constant, and the inspection dummies (one for the nine years around the inspection and one for the five years after the inspection) interacted with estimated operating abatement cost at the
2. All specifications control for year-region effects, industry-year interactions, region-year interactions, plant fixed effects and the predicted probability of inspection.
3. Standard errors are clustered at the plant level and are in parentheses.
4. * means significant at 10%; ** means significant at 5%; *** means significant at 1%.
5. Inspection and Fine Data are drawn from the IDEA database. Emissions data are drawn from the Toxic Release Inventory.

References

Ashenfelter, Orley and David Card (1985), "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs," *The Review of Economics and Statistics*, 67(4), pp. 648-660.

Becker, Randy A., (2005), "Air Pollution Abatement Costs under the Clean Air Act: Evidence from the PACE Survey," *Journal of Environmental Economics and Management*, 50(1), pp. 144-169.

Borja-Aburto, BV.H., D.P. Loomis, S.I. Bangdiwala, C.M. Shy and R.A. Rascón-Pacheco, 1997. Ozone, Suspended Particulates, and Daily Mortality in Mexico City, *American Journal of Epidometry,* 145(3), pp. 258-268.

Duflo, Esther, Rema Hanna, and Stephan Ryan (2007), "Monitoring Works: Getting Teachers to Come to School," Mimeo, Harvard University, Cambridge, MA.

Eckert, Heather (2004), "Inspections, Warnings, and Compliance: the Case of Petroleum Storage Regulation," *Journal of Environmental Economics and Management*, 47(2), pp. 232-259.

Gray, Wayne B. and Ronald J. Shadbegian (2004), "'Optimal' Pollution Abatement--Whose Benefits Matter, and How Much?" *Journal of Environmental Economics and Management*, 47(3), pp. 510-34.

Gray, Wayne B. and Ronald J. Shadbegian, (2005), "When and Why do Plants Comply? Paper Mills in the 1980s," *Law & Policy*, 27(2), pp. 238-261.

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

Hanna, Rema (2009). "U.S. Environmental Regulation and FDI: Evidence from a Panel of U.S. Based Multinational Firms," Forthcoming in the *Applied Economic Journal: Applied Economics*.

Jacobson, Louis S., Robert LaLonde and Daniel Sullivan (1993), "Earnings Losses of Displaced Workers," *The American Economic Review,* 83(4), pp. 685-709.

Keller, Wolfgang, and Arik Levinson (2002). "Environmental Regulations and FDI Inflows to U.S. States," *Review of Economics and Statistics*, 84(4), pp. 691 -703.

Keohane, Nathaniel O., Erin T. Mansur, and Andrey Voynov, (2009), "Averting Enforcement: Strategic Response to the Threat of Environmental Regulation," *Journal of Economics and Management Strategy*, 18(1), pp. 75-104.

Kahn, Matthew E. (1997). "The Particulate Trends in the United States," *Regional Science and Urban Economics*, 27(1), pp. 87-107.

Laplante, B. and P. Rilstone (1996), "Environmental Inspections and Emissions of the Pulp and Paper Industry in Quebec," *Journal of Environmental Economics and Management*, 31(1), pp. 19-36.

Levy, J.I., T.J. Carrothers, J. Tuomisto, J.K. Hammitt and J.S. Evans (2001). "Assessing the public health benefits of reduced ozone concentrations," *Environmental Health Perspectives* (in press).

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

Magat, Wesley and Kip Viscusi (1990), "Effectiveness of the EPA's Regulatory Enforcement: The Case of Industrial Effluent Standards," *Journal of Law and Economics*, 33(2), pp. 331-360.

Mid-Atlantic Audit Division Philadelphia (1998), "Consolidated Report on OECA's Oversight of Regional and State Air Enforcement Programs," EPA Reports.

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

Shimshack, Jay and Michael Ward (2008), "Enforcement and Overcompliance," *Journal of Environmental Economics and Management,* 55(1), pp. 90-105.

Wooldridge, Jeffrey M. (2002), *Econometric Analysis of Cross Section and Panel Data*, The MIT Press, pp. 116 and 355.