CLEAN ECONOMY
WORKING PAPER SERIES

JANUARY 2017/ WP 17-01

ENVIRONMENTAL REGULATIONS AND THE CLEAN-UP OF MANUFACTURING:

PLANT-LEVEL EVIDENCE FROM CANADA

Nouri Najjar
University of British Columbia

Jevan Cherniwchan
University of Alberta

This research project was supported by the Economics and Environmental Policy Research Network (EEPRN)

This project was undertaken with the financial support of:
Ce projet a été réalisé avec l'appui financier de :

Environment and Climate Change Canada
Environnement et Changement climatique Canada

Smart Prosperity Institute

uOttawa
Institut de l'environnement
Institute of the Environment

Environmental Regulations and the Clean-Up of Manufacturing: Plant-Level Evidence from Canada*

Nouri Najjar
University of British Columbia
npmnajjar@gmail.com

Jevan Cherniwchan
University of Alberta
jevan@ualberta.ca

Most Recent Version Here

Abstract

For much of the industrialized world, pollution from manufacturing has been falling despite increased output. In this paper, we provide the first estimates of the extent to which environmental regulations have contributed to this "clean-up" of manufacturing by causing: (i) the adoption of cleaner production processes, (ii) the reallocation of output across producers, and (iii) producer entry and exit. To do this, we examine a major revision to Canadian environmental policy using a novel, confidential dataset containing information on the production decisions and pollution emissions of Canadian manufacturing plants. We find regulation explains, at most, 61% of the Canadian clean-up, but the underlying mechanisms differ strikingly across pollutants. We present a stylized model featuring plant heterogeneity to illustrate how the costs of abating pollution can affect the channels through which regulation causes a clean-up.

JEL: D22, L51, L60, Q52, Q53, Q58.
Keywords: Air Pollution, Manufacturing Clean-Up, Environmental Regulation.

---

[*] We would like to thank Brian Copeland, Carol McAusland, Thomas Lemieux, David Green, Arik Levinson, Joseph Shapiro, and seminar participants at UBC, the CREE annual meeting, the EEPRN symposium, the CEA annual conference, the Young Economists Symposium at Yale, the CU Environmental and Resource Economics Workshop, and the Pacific Northwest Labor Day for helpful comments and suggestions. We are especially indebted to Michael Willox and Danny Leung at Statistics Canada's Economic Analysis Division for their efforts to make this project possible. We would also like to thank members of Environment and Climate Change Canada's (ECCC) Legislative and Regulatory Affairs Division for helpful discussion on the design and implementation of Canadian air quality regulation. Funding from the EEPRN, the Alberta School of Business Nova Faculty Fellowship, and the Centre for Applied Business Research in Energy and the Environment at the University of Alberta is gratefully acknowledged. The data used in this paper was created through a collaboration between ECCC, Statistics Canada, and the Productivity Partnership, which is supported by the Social Sciences and Humanities Research Council of Canada. All results have been reviewed by Statistics Canada to ensure no confidential information is disclosed. The opinions expressed in this paper are the authors' alone.

1 Introduction

The past thirty years have witnessed a marked improvement in manufacturing pollution levels across much of the world despite large increases in manufacturing activity. In the United States, for example, manufacturing emissions of most air pollutants fell by between 52%-69% from 1990 to 2008, while total real shipments from the sector rose by 35% (Levinson, 2015). In Europe, manufacturing air pollution fell by between 23-59% from 1995 to 2008, while real shipments rose by 37% (Brunel, 2016). These patterns appear to extend outside of the United States and Europe; sulphur dioxide emissions from manufacturing have been falling in a number of countries despite increases in shipments (Grether et al., 2009). These broad trends imply that, for much of the world, manufacturing is becoming cleaner.

In this paper, we ask how plant-level responses to environmental regulation have contributed to this "clean-up" of manufacturing. Regulation can cause a clean-up by changing a sector's composition – reallocating output across plants or causing plant entry and exit – or by changing the production processes used by plants. Recent work has shown that environmental policies have affected the composition of the manufacturing sector (Henderson, 1996; Greenstone, 2002) and reduced the sector's aggregate pollution intensity (Shapiro and Walker, 2015). However, there are no estimates of regulation's effect on the processes used by manufacturing plants, or the degree to which each of these channels have contributed to the clean-up. This makes it difficult to determine how the composition of the economy might change in transitioning to a less-polluting economy. Addressing this is important because changes in an economy's composition could potentially disrupt labor and product markets (see, e.g., Foster et al. (2008)), whereas changes in production processes would not. We provide the first estimates of the magnitude of each of these channels of regulation. To do this, we estimate the effects of a major revision to Canadian environmental policy on Canadian manufacturing plant entry and exit, the reallocation in output from regulated to unregulated surviving plants, and the adoption of cleaner production processes by surviving plants.

We start by examining the manufacturing clean-up in Canada using the industry decomposition developed by Levinson (2009, 2015). Based on these results, the clean-up in Canada appears to be very similar to those documented in the United States and Europe. In each case, total manufacturing emissions of most air pollutants have fallen substantially, and this is primarily because of reductions in industry pollution intensity.

The results of our decomposition suggest the clean-up of Canadian manufacturing is primarily due to within-plant changes in emission intensity, within-industry reallocations of activity across plants, and plant entry and exit. In order to understand how environmental regulation has affected these channels, we study the effects of the Canada Wide Standards for Particulate Matter and Ozone (CWS) on the emission intensity and productive activities of plants that emit pollutants directly targeted by the policy.

The CWS was created in the year 2000 as a result of an agreement between the federal government of Canada and the various provincial governments, and was the major air quality regulation in place in Canada over the period 2000-2012. The agreement was designed to ensure each region met a minimum level of air quality by establishing thresholds for the ambient concentrations of two common pollutants: fine scale particulate matter ($PM_{2.5}$) and ground-level ozone (O_3). Regions in which the ambient concentrations of either pollutant exceeded the relevant threshold in a given year were subject to more stringent regulation relative to other regions. In addition, these regulations explicitly targeted plants in a set of "targeted industries". Under the CWS, the annual permits required by plants to operate in each province were used to impose more stringent rules on pollution emissions for plants in targeted industries and regions violating one of the CWS standards. These standards either required plants to adopt technical changes to meet industry best practices, or reduce activities generating the regulated pollutant.[1]

To guide our analysis, we next develop a simple theory based on a closed-economy variant of the Melitz (2003) model in which heterogenous plants face regulatory constraints similar

---

[1] We describe the CWS in more detail in Section 3.

to those imposed under the CWS. This model has three key features. First, it allows for plant productivity differences, which have been highlighted as a key determinant of the effects of environmental regulation in the existing theoretical literature (see, e.g., Konishi and Tarui (2015) or Anoulies (2017)). Second, it allows for endogenous technology adoption by plants to capture the fact that leading technologies were used as a benchmark for the technical changes required under the CWS. Third, it allows for differences across pollutants in the cost of adopting less-polluting production processes, which we call abatement costs. This feature is important because, as we discuss in Section 3, the two pollutants we focus on carry very different abatement costs, at least in Canada: nitrogen oxide[2] (NO_X) abatement can be accomplished at a relatively low cost, while $PM_{2.5}$ abatement typically requires high fixed costs (Canadian Council of Ministers of the Environment, 1998; Environment Canada, 2002).

We then test the predictions generated by this framework. We directly estimate the effects of the CWS on affected plants using a unique confidential dataset containing longitudinal information on the pollution emissions and productive activities of Canadian manufacturing plants over the period 2004-2010. We identify the effects of environmental regulation on manufacturing plants by exploiting variation in the stringency of regulation across regions, time, and industries created by the design of the CWS. As such, we adopt a triple-difference research design that embeds three margins of comparison. First, we examine the effects on plants in regulated regions, by comparing outcomes between plants in regulated and unregulated industries in years in which their region is regulated. Next, we compare outcomes between these plants in years in which their region is unregulated. Lastly, we compare outcomes between plants in regulated and unregulated industries in regions that are never regulated. This approach allows us to flexibly control for factors such as localized recessions or industry demand shocks that would otherwise confound the effects of environmental regulation.

---

[2] Nitrogen oxide is a main contributor to O_3 pollution, and a main target of the CWS.

We start our assessment of the CWS by asking what effect, if any, the policy had on the level of pollution emitted by affected Canadian manufacturing plants. We begin here to establish that the CWS did contribute to the manufacturing clean-up in Canada. We find robust evidence that the CWS reduced pollution emissions from affected manufacturing plants. For the average $\mathrm{PM}_{2.5}$ emitting plant, the CWS is associated with a 15% reduction in $\mathrm{PM}_{2.5}$ emissions. Furthermore, we find that the CWS is associated with a 33% reduction in NO_{X} emissions from the average NO_{X} emitting plant.

Next, we examine the effects of the CWS on the emission intensity of affected plants. Our theory predicts that the effects of the CWS will depend on the fixed costs of abatement. If fixed costs are high, as in the case of $\mathrm{PM}_{2.5}$, only relatively productive plants will adopt cleaner production processes following regulation. As a result, the CWS should have little to no effect on the emission intensity of the average plant. If the fixed costs of abatement are low, as with NO_{X}, then even less productive plants should respond to regulation by adopting cleaner production processes. In this case, the emission intensity of the average plant should fall in response to the CWS. Our empirical estimates support these predictions; we find the CWS did not have a significant effect on the emission intensity of the average affected $\mathrm{PM}_{2.5}$ emitting plant, but is associated with a 29% reduction in the NO_{X} emission intensity of the average affected NO_{X} emitting plant.

We also examine the effects of the CWS on plant output and exit. As predicted by the model, we find that the CWS was associated with a 11% reduction in output from the average affected $\mathrm{PM}_{2.5}$ emitting plant, but had little to no effect on the output of the average plant that emits NO_{X}. In addition, we find that the CWS was associated with a significant reduction in the number of plants that emit $\mathrm{PM}_{2.5}$, but had little to no effect on the entry and exit of plants that emit NO_{X}.

Taken together, these estimates suggest that environmental regulations contributed significantly to the clean-up of the Canadian manufacturing sector. For example, our estimates for the responses of $\mathrm{PM}_{2.5}$ emitters suggest that the effects of the CWS explain close to 21%

of the reduction in the $PM_{2.5}$ intensity of manufacturing. Similarly, our estimates for the responses of NO_X emitters suggest that nearly 61% of the reduction in the manufacturing sector's NO_X emission intensity can be attributed to the effects of the CWS. However, the mechanisms driving these responses vary starkly across pollutants; the $PM_{2.5}$ clean-up was primarily driven by changing the composition of economic activity within industries, whereas the clean-up of NO_X was primarily due to within-plant reductions in emission intensity.

Our model suggests these differential responses to regulation are due to differences in the fixed cost of abating across pollutants. While we have focused on this channel given the available evidence documenting the stark differences in the costs of abating $PM_{2.5}$ and NO_X, we do not observe these costs directly. Hence, to provide further evidence that our estimates are consistent with this mechanism, we test the heterogeneity in responses to regulation implied by our model. The effects of regulation should only vary across plants of different productivity levels if the fixed costs of abatement are high. We test this prediction by allowing the estimated effects of the CWS to differ across plants on the basis of their initial labor productivity level. The empirical results match our model's predictions. We find considerable differences across plants in the effects of the CWS on $PM_{2.5}$ emitters, but relatively homogeneous effects on NO_X emitters. Pollution from relatively high-productivity regulated $PM_{2.5}$ plants fell primarily due to reductions in emission intensity, whereas pollution emissions from the least-productive $PM_{2.5}$ plants fell due to a reduction in output. In contrast, NO_X pollution intensity fell for both high and low productivity plants. These results further suggest that fixed abatement costs are a key determinant of how industry cleans up in response to environmental regulation.

Finally, we examine the effects of the CWS on several additional margins via which plants could respond to regulation, including changes in primary factor use, intermediate input use, and productivity. This allows us to test a number of alternative explanations for why we observe different responses to the CWS across pollutants. We find little evidence for these alternative explanations.

Altogether, our findings contribute to a burgeoning literature examining the sources of the clean-up of the manufacturing sector. This research began with the work by Levinson (2009) examining how trade-induced changes in industrial composition have contributed to the clean-up of US manufacturing. Levinson finds that these changes played a small role; the clean-up is primarily due to reductions in industry emission intensity.[3,4] Our work is most closely related to that of Shapiro and Walker (2015), who use a structural model to ask whether the clean-up of the U.S. manufacturing sector has been caused by regulation, trade, productivity growth, or other economic factors. Using this model, Shapiro and Walker conclude the effects of regulation explains most of the reduction in pollution emissions from US manufacturing. By directly estimating the effects of regulation, our analysis complements this work along two key dimensions. First, we provide causal evidence of how environmental regulations contributed to the clean-up of manufacturing by altering the emission intensities of individual plants, reallocating activity across plants, and changing entry and exit decisions. Second, we provide evidence as to the mechanisms via which plants change their emission intensities in response to environmental regulation.[5]

This paper also relates to work examining the effects of air quality regulation on the emissions of manufacturing plants. Fowlie et al. (2012), for example, find Southern California's RECLAIM cap-and-trade program reduced $\mathrm{NO_X}$ emissions from manufacturing plants. In addition, Greenstone (2003) and Gibson (2016) examine the effects of the U.S. Clean Air Act, concluding regulation reduced both the growth (Greenstone, 2003) and level (Gibson, 2016) of air pollutant emissions from manufacturing plants. Our paper complements this work by determining whether changes in plant pollution in response to regulation are due to changes in the level of output produced, or changes in the emission intensity of production.

Lastly, our work also relates to a large literature examining the effects of air quality reg-

---

[3] Brunel (2016) documents a similar pattern using European data.

[4] Others have argued trade may have caused surviving plants to change how they produce their goods (due, for example, to the outsourcing of inputs or the adoption of new abatement technologies), thereby leading to a reduction in plant-level pollution intensity (see Martin (2012) or Cherniwchan (2017)).

[5] Our work is also related to Martin et al. (2014) who tackle a related, but distinct question, in asking how the energy intensity of UK manufacturing plants were affected by a carbon tax.

ulation on the industrial activities of manufacturing plants. This literature has typically focused on the effects of regulation on plant employment, production, capital, and productivity (see, e.g. Berman and Bui (2001), Greenstone (2002), Popp (2003), Shadbegian and Gray (2006), Hanna (2010), Greenstone et al. (2012), and Walker (2013)), primarily in the context of the U.S. Clean Air Act or other U.S. policies. Our results complement these findings in three ways. First, we show that the effects of environmental regulation on industrial activity may vary across plants of different productivity levels. Second, we find that some of the adjustments to policy that have been documented elsewhere in the literature (lower output, as in Greenstone (2002); changes in productivity, as in Greenstone et al. (2012)) extend outside the US. Finally, we show that pollution emissions may fall at regulated plants even if there is no significant reduction in production.

The remainder of this paper proceeds as follows. In Section 2, we document the clean-up of the Canadian manufacturing sector. Section 3 provides a brief overview of the CWS. In Section 4 we outline our theory. Section 5 presents our data, outlines our research design and empirical specification, and presents our empirical results. Finally, Section 6 concludes.

2 The Clean-Up of Canadian Manufacturing

Our goal in this paper is to determine how the effects of environmental regulation on individual plants have contributed to the clean-up of manufacturing. While the clean-up has been documented in several countries, including the United States (e.g. Levinson (2009, 2015), Shapiro and Walker (2015)), and the European Union (e.g. Brunel (2016)), there is no evidence of whether a similar trend appears in Canada. Hence, before we examine the effects of environmental regulation, we first examine whether the changes in the pollution emitted by the Canadian manufacturing sector mirror those that have occurred elsewhere.

These trends, relative to 1992 levels, are illustrated in Figure 1. The figure depicts changes in the aggregate emissions of four common pollutants from the Canadian manufacturing sector, as well as changes in aggregate manufacturing output. As it shows, Canada

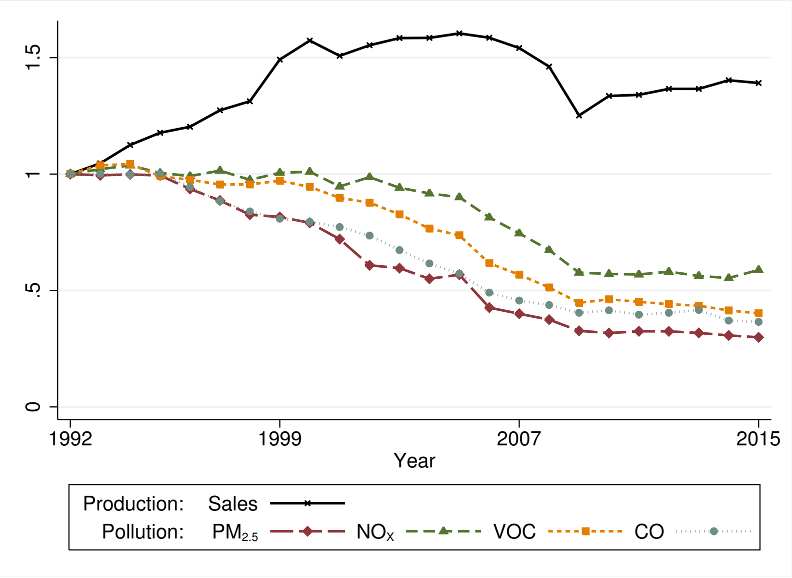

Figure 1: Production and Pollution Emissions from Canadian Manufacturing: 1992-2015

Notes: Figure depicts trends in real manufacturing sales and aggregate emissions of fine scale particulate matter ($PM_{2.5}$), nitrogen oxide (NO_X), volatile organic compounds (VOCs), and carbon monoxide (CO). Aggregate pollution is from Environment and Climate Change Canada's Air Pollutant Emission Inventory. Aggregate output is measured as the real value of manufacturing shipments, constructed by deflating data on industry-level nominal shipment values from Statistics Canada's CANSIM table 304-0014 using the industry price data given in Statistics Canada's CANSIM table 329-0077.

has experienced a large decline in the emission intensity of its manufacturing sector since 1992. Throughout the 1990s, manufacturing output grew substantially, but there was little change in aggregate emissions of NO_X or volatile organic compounds (VOCs) and moderate reductions in particulate matter and carbon monoxide (CO) emissions. Over the early- to mid- 2000s, while the manufacturing sector stagnated, aggregate emissions of all four pollutants fell dramatically. From the late-2000s on, the sector shrank in terms of both output and pollution, with a mild recovery in output following the 2008 financial crisis, and a roughly constant level of emissions from 2010 onwards. Overall, from 1992 to 2015 real manufacturing output rose approximately 39%, while emissions fell by between 41% and 70%, depending on pollutant. These estimates imply that, on average, the emission intensity of the Canadian manufacturing sector fell by 3.5-4.7% annually.

This suggests the Canadian manufacturing clean-up was of a similar magnitude to the clean-ups that occurred in the U.S. and Europe. For example, Levinson (2015) finds the emission intensity of US manufacturing fell by 3.6-4.3% annually from 1990 to 2008. Similarly, Brunel (2016) shows the emission intensity of European manufacturing fell by 4.6-7.4% annually over the period 1995-2008.

While this evidence shows the magnitudes of the clean-ups in Canada, the US, and Europe were similar, it reveals little as to whether the potential sources were the same. As such, we adopt a simple decomposition exercise first used by Levinson (2009) to study the potential sources of the clean-up.[6] This approach is useful because it allows us to determine if the observed reductions in aggregate emission intensity are driven by a "composition effect" created by a reallocation of economic activity from dirty pollution-intensive sectors to clean sectors with relatively low pollution intensities or by a "technique effect" created by reductions in the emission intensity of individual industries.

To make this decomposition explicit, let Z, X, and $E = Z/X$ denote the pollution emissions, output, and pollution intensity of the manufacturing sector, respectively. Let Z_i, X_i, and E_i denote the same for individual manufacturing industries[7], indexed by i. Manufacturing pollution intensity can then be written as $E = \sum_i \theta_i E_i$, where $\theta_i = X_i/X$ denotes industry i's share of output from the manufacturing sector. Totally differentiating yields

$$dE = \sum_i E_i d\theta_i + \sum_i \theta_i dE_i. \tag{1}$$

The first term of equation (1) is the aforementioned composition effect, while the second term is the technique effect.

We follow the approach taken by Levinson (2015) and take equation (1) directly to the data. This gives us estimates of the reduction in manufacturing pollution intensity attributable to both the composition and technique effects for $\mathrm{PM}_{2.5}$, NO_X, VOCs, and CO over the period 1992-2015. These estimates are reported in Table 1. The first two columns report the percentage change in emissions and emission intensity that occurred for the sector

---

[6] This approach traces back to work of Grossman and Krueger (1991) and Copeland and Taylor (1994).

[7] Due to constraints from the pollution data, our industry definitions correspond to either the three- or four-digit NAICS code.

Table 1: Decomposing Emissions from Canadian Manufacturing: 1992-2015

	Δ Aggregate Pollution (1)	Δ Aggregate Pol. Intensity (2)	Technique Effect (3)	Composition Effect (4)	Technique Effect Share (5)
$PM_{2.5}$	-0.70	-0.79	-0.78	-0.01	0.99
NO_X	-0.41	-0.58	-0.52	-0.06	0.90
VOCs	-0.60	-0.71	-0.67	-0.04	0.94
CO	-0.63	-0.74	-0.73	-0.01	0.99

Notes: Table reports estimates from a decomposition of the change in pollution intensity of the Canadian manufacturing sector from 1992 to 2015 into composition and technique effects. Estimates are from a Laspeyre's-type index following Levinson (2015). Each row reports estimates for a different pollutant. The first two columns report the percentage change in total emissions and emission intensity from the manufacturing sector, respectively. The third and fourth columns report the reduction in aggregate emission intensity due the technique and composition effects, respectively. The final column shows the fraction of column (2) attributable to changes in the technique effect, calculated as (column (3)/column (2)).

as a whole. The third and fourth columns report the change in aggregate emission intensity attributable to the technique effect and composition effects, respectively.[8] The final column reports the share of the change in aggregate pollution intensity due to the technique effect.

The estimates reported in Table 1 suggest that the clean-up of the Canadian manufacturing sector can primarily be attributed to the technique effect. For example, the estimate reported in the first row indicates that during the 1992-2015 period, changes in industry emission intensity accounted for 99% of the reduction in manufacturing $PM_{2.5}$ intensity. This is further evidence that the Canadian clean-up is similar to those observed elsewhere; as shown by Levinson (2009, 2015) and Brunel (2016), the clean-ups of US and European manufacturing are also primarily due to the technique effect.[9]

2.1 How do Industries Clean-Up?

The evidence presented above indicates that reductions in industry emission intensities are the predominant source of the manufacturing clean-up that has been observed in Canada and elsewhere. This suggests that plant-level responses are a key determinant of the observed

---

[8] The technique effect is calculated by taking the percentage change in a Laspeyre's-type index of $\sum_i \theta_i dE_i$. The composition effect is calculated as the difference between the change in manufacturing emission intensity and the technique effect.

[9] In addition, **??** perform a product-level decomposition, and find the clean-up in the US is primarily due to within-product reductions in pollution intensity.

reductions in manufacturing emission intensity.

To see this, it is useful to extend the logic of the decomposition exercise presented in Levinson (2009) and further decompose the technique effect into plant level responses as in Cherniwchan et al. (2017). Suppose, as above, that pollution intensity of industry i is given by $E_i = Z_i/X_i$. In addition, suppose each industry is composed of a continuum of plants and let $x_i(n)$ and $z_i(n)$ denote output and pollution from plant n, so $e_i(n) = z_i(n)/x_i(n)$ is the emission intensity of plant n and $\lambda_i(n) = x_i(n)/X_i$ is plant n's share of production in industry i. Furthermore, let n_i denote the marginal plant that is endogenously determined by the industry's profitability. In this case, the emission intensity of industry i can be expressed as a weighted average of the emission intensities of each plant in the industry: $E_i = \int_0^{n_i} e_i(n)\lambda_i(n)dn$. The change in emission intensity of any industry i is then

$$dE_i = \int_0^{n_i} de_i(n)\lambda_i(n)dn + \int_0^{n_i} e_i(n)d\lambda_i(n)dn + [e_i(n_i) - E_i]\lambda_i(n_i)dn_i. \tag{2}$$

The first term on the right-hand side of equation (2), which we term the "process effect", captures the changes in industry emission intensity due to plant-level changes in emission intensity created by changes in the processes used in production. As such, this term captures the direct effects of a shock; all else equal, an industry's emission intensity will fall if a shock such as environmental regulation induces plants in the industry to lower their emission intensities. The remaining two terms capture the indirect changes in industry emission intensity created by a shock. The first of these, given in the second term on the right-hand side of equation (2), captures the effects of the shock on the relative size of plants within an industry. This "reallocation effect" would arise if the shock does not affect plants uniformly; if the shock only affects a subset of plants in an industry, as is common with many environmental regulations, the resulting change in output may cause a reduction in the relative output of affected plants leading to a change in industry emission intensity, even in the absence of direct changes in plant emission intensity. Finally, the "selection effect" given by the third term captures the change in emission intensity created by a change in the set of plants operating within the industry owing to plant entry and exit.

Equation (2) shows that regulation may cause an industry's emission intensity to fall by causing plant-level reductions in emission intensity (the process effect), changes in the relative output of dirty and clean plants (the reallocation effect), or a change in the plants that comprise the industry (the selection effect). Yet, the majority of existing studies have examined the effects of regulations on either selection (e.g. Henderson (1996), Becker and Henderson (2000)) or reallocation (e.g. Greenstone (2002)). In what follows, we present direct estimates of the process, reallocation, and selection effects and quantify their contributions to the clean-up of Canadian manufacturing.[10]

3 Air Quality Regulation in Canada

In order to understand how environmental regulations affected pollution from manufacturing plants and quantify their contribution to the clean-up of Canadian manufacturing, we examine the effects of the Canada Wide Standards for Particulate Matter and Ozone (CWS). The CWS was the primary policy targeting particulate matter and ozone pollution throughout Canada over the period 2000-2010.[11] Moreover, the design of the CWS makes it an attractive setting for studying the effects of environmental regulation.

First signed in 2000, the CWS was an agreement between the federal government of Canada and the various provincial environment ministries. The intent of the CWS was to improve air quality across the country by the end of 2010 by implementing two air quality standards – one for $PM_{2.5}$ and one for O_3 – that applied to each major town or city in Canada (we call these Census Metropolitan Areas or CMAs) (Canadian Council of Ministers

---

[10] Holladay and LaPlue (2017) present the results of a plant-level decomposition exercise similar to Equation (2) of toxic chemical emissions from the U.S. manufacturing. They show that toxic emissions fell throughout the 1990s primarily due to plant-level changes in pollution intensity, although cannot explain why pollution intensity fell.

[11] It was subsequently replaced with the Canadian Ambient Air Quality Standards for Fine Particulate Matter and Ozone in 2012. We end our study period in 2010 to avoid any potential contamination by this regulatory change, as the planning for this transition began in 2011.

of the Environment, 2000a).[12] Much like the National Ambient Air Quality Standards at the centre of the Clean Air Act Amendments (CAAAs) in the United States, these standards created a target level of air quality that needed to be achieved by each CMA in Canada. These standards were common across all CMAs, and each CMA was required to meet the standards by the end of 2010.[13] To that end, plants in CMAs with ambient concentrations of either $PM_{2.5}$ or O_3 in excess of the relevant standard's threshold were subject to more stringent environmental regulation than plants in relatively clean CMAs.

In addition to differentiating between regions on the basis of air quality, the CWS explicitly designated a set of "targeted industries" that were to be the focus of more stringent regulation.[14] These industries were chosen because they were viewed as major contributors to the air quality problems that motivated the CWS, and were common across all CMAs.

While on its surface the CWS was a relatively simple policy (manufacturing plants in targeted industries in dirty CMAs were more stringently regulated than other plants), the implementation was anything but simple. This is because regulatory authority over environmental issues in Canada is shared between the local, provincial, and federal governments (Boyd, 2003). Below, we briefly describe the broad framework used across the country, so as to clarify the core regulatory mechanisms through which the CWS was implemented.[15]

Broadly, the CWS was a tiered regulatory approach in which the federal and provincial

---

[12] The agreement defines a major town or city as a Census Agglomeration (CA) or Census Metropolitan Area (CMA). A CMA must have a total population of at least 100,000 with a population in the center of the region of 50,000. A CA must have a core population of at least 10,000 (and generally have a total population less than 100,000). For detailed CA and CMA definitions, see Statistics Canada (2015). For convenience, we use the term CMA to refer to both CAs and CMAs.

[13] The standard for particulate matter required each CMA's $PM_{2.5}$ concentration lie below $30\mu g/m^3$ over each 24-hour period (that is, every day from midnight to midnight). Achievement of the $PM_{2.5}$ standard was based on the 98^{th} percentile of each region's 24-hour ambient concentration in a given year. The O_3 standard was applied as an 8-hour standard that required each CMA's O_3 concentration lie below 65 parts per billion (ppb). Achievement of the O_3 standard was based on the 4^{th} highest 8-hour concentration reported in a given year. In comparison, the National Ambient Air Quality Standards in the United States currently contain a 24-hour $PM_{2.5}$ standard set at $35\mu g/m^3$, and an 8-hour O_3 standard set at 70 pbb. For details on the US standards, see Environmental Protection Agency (2016).

[14] The targeted industries were pulp and paper, lumber and wood product manufacturing, electric power generation, iron and steel manufacturing, base metal smelting, and the concrete and asphalt industries (Canadian Council of Ministers of the Environment, 2000b).

[15] For a detailed overview of air quality regulation in Canada, see Taylor and McMillan (2014).

governments agreed on local air quality targets (the CWS standards), the federal government developed best practice and guidance documents for targeted industries (called Multi-pollutant Emission Reduction Strategies, or MERS) to provide management tools to the provincial governments (Government of Canada, 2003), and the provincial governments used new and existing regulations to meet these standards.[16] The main provincial mechanisms were annual provincial operation permits and licenses, which plants required to operate. Plants had to submit applications for these permits to prove compliance with certain environmental regulations[17] and to have any facility changes approved (see, e.g. Angle (2014, p. 290), Environment Canada (2004, p. 11), and Environment Canada (2002, p. xi)).[18] In most instances, facilities could effectively follow one of two paths to meet the permitting requirements: they could either adopt one or more technical changes recommended in their industry's MERS, or reduce activities contributing to the problematic pollutant.[19] When local air quality was relatively clean (i.e. a region was in compliance with the CWS), the permitting constraints were laxer than when air quality was poor. This means the regulatory stringency facing a plant varied over time according to their region's air quality. Indeed, there are examples of plants replacing relatively clean methods with less expensive, more polluting, options that were approved because local air quality had improved.[20]

There is reason to believe the CWS caused substantial variation in the stringency of environmental regulation. Nearly 40% of the most populated regions of the country had

---

[16] Provincial regulations played a similar role to the State Implementation Plans submitted by each US state to the EPA under the Clean Air Act. However, the CWS allowed provinces to adopt a wide number of instruments to achieve each standard, which we discuss in the online appendix (see Section C).

[17] In Section C, we provide a full overview of the relevant provincial policies in place.

[18] For example, in Ontario facilities wishing to make even small changes, such as the replacement of natural gas burners or the installation of new generators, had to show they could do so without violating the CWS and provincial air quality regulations. All facilities, regardless of whether changes were being made, also had to show annual compliance with provincial air quality regulation. From 2000 to 2010, at least 2,800 permit applications cited particulate matter or nitrogen oxide emissions.

[19] For an example of a permit requiring a production limit, see `https://www.accessenvironment.ene.gov.on.ca/instruments/2000-9GFRTP-14.pdf`. For an example of a permit stipulating a particulate matter abatement device (a baghouse), see `https://www.accessenvironment.ene.gov.on.ca/instruments/7222-97RNJW-14.pdf`.

[20] For an example in which a low nitrogen oxide natural gas burner was replaced with a higher polluting option see `https://www.ebr.gov.on.ca/ERS-WEB-External/displaynoticecontent.do?noticeId=MTA0NDQ1&statusId=MTU2OTQ0&language=en`.

their regulatory status change during the period over which the CWS was in force. The regulations also appear to be associated with large changes in air quality. On average, $PM_{2.5}$ concentrations in affected regions fell by approximately 20%, while O_3 concentrations in affected regions fell by approximately 9%.[21]

The CWS also targeted two different pollutants, each facing different technical constraints on the abatement options available to facilities. The first pollutant, NO_x, is primarily caused by the combustion of fossil fuels. Facilities can reduce NO_X emissions at a relatively low cost by adopting efficient combustion processes[22] or by adopting relatively low-cost low-NO_x emissions burners (see, e.g. Environment Canada (2002, p.xiii), Canadian Council of Ministers of the Environment (1998), or Environmental Protection Agency (1999a)). Indeed, Canada's NO_X emissions guidelines for industrial boilers and heaters "are based on proven compatibility with efficient combustion operation and the use of cost-effective technology such as low NO_X burners" (Canadian Council of Ministers of the Environment, 1998, p. 3). In contrast, the second pollutant, $PM_{2.5}$ is caused by the combustion of fossil fuels, chemical reactions, wear and tear on machinery, and the processing of lumber. Reducing $PM_{2.5}$ emissions typically requires installing a large filtration system, such as a baghouse or electrostatic precipitator, that carries a large fixed cost (see, e.g., Environment Canada (2002, p.xiii and xvii) and Environmental Protection Agency (1998, 2002, 1999b)).[23]

4 Theory

As we described above, the CWS regulations required plants to either adopt technical changes to meet industry best practices, or reduce the activity responsible for generating the regulated pollutant. Given this regulatory structure, we develop a simple theoretical framework featuring process, reallocation, and selection effects in response to environmental regulation to help guide our empirical analysis.

---

[21] In the online appendix (Section B.3) we present descriptive evidence of this improvement in air quality.

[22] This may entail either changing the temperature or the fuel-oxygen ratio at which combustion occurs.

[23] As a reference, engineering abatement cost estimates are between \$1,000 to \$20,000 per ton of $PM_{2.5}$ using an electrostatic precipitator, between \$2,000 to \$100,000 per ton of $PM_{2.5}$ using a baghouse, and between \$200 to \$1,000 per ton NO_X (low-NO_X burner) (Environmental Protection Agency, 2006).

Our model is based on a closed economy version of the Melitz (2003) model in which heterogeneous plants emit pollution as a byproduct of production, and face regulations similar to those imposed under the CWS. This simple, stylized structure is intended to capture three key features of our regulatory setting. First, it features differences in productivity across plants, as recent theoretical work has emphasized the importance of these differences in determining the effects of environmental regulation (e.g. Konishi and Tarui (2015) or Anoulies (2017)). Second, it allows for endogenous technology adoption by plants following the approach developed by Bustos (2011).[24] This means plants differ in their technology choices even in the absence of regulation, allowing us to capture the fact that pre-regulation technology differences were used to inform the technical changes required under the CWS. Finally, our model allows for differences in abatement costs across pollutants, meaning we are able to examine how these costs affect a plant's response to regulation.

Below, we outline the features of the model, describe its solution and highlight its key empirical predictions. For the sake of brevity, we relegate the details of the model's solution and the derivations of the comparative statics underlying the empirical predictions to the online appendix (see Section A).

4.1 Setup

We consider an economy comprised of L identical consumers, each endowed with a single unit of labor. Labor is supplied inelastically and used to produce differentiated products in a single industry. Production also creates pollution as a byproduct, and this harms consumers, lowering their utility. For convenience, in what follows, we let wages be the numeraire.

The representative consumer derives utility from the consumption of goods and disutility from aggregate pollution according to $U = [\int_0^M q(\omega)^\rho d\omega]^{1/\rho} - h(Z)$, where $q(\omega)$ denotes consumption of good ω, and M denotes the measure of varieties available in the economy. It is assumed consumers ignore pollution when making their consumption decisions. As a

---

[24] This approach has been used previously to link technological upgrading to changes in pollution levels. See, for example, Batrakova and Davies (2012), Forslid et al. (2014) or Cherniwchan et al. (2017).

result, the demand for variety ω is given by $q(\omega) = IP^{\sigma-1}p(\omega)^{-\sigma}$, where I denotes consumer income, $P = [\int_0^M p(\omega)^{1-\sigma}d\omega]^{1/(1-\sigma)}$ is the economy's price index, and $\sigma = 1/[1-\rho] > 1$ is the elasticity of substitution between goods.

The supply side of the economy features monopolistic competition and free entry, meaning each firm in the economy produces a unique variety. To enter, firms pay a fixed entry cost f_ϵ, and upon entry, draw a productivity level φ from a common distribution $G(\varphi)$.[25] Based on the realization of φ, firms decide whether to exit or stay in the market, and conditional on staying, how much to produce and what technology to use in production.

Upon entering, firms are able to produce output x using a business-as-usual technology (labeled b) that features increasing returns to scale. With this technology, the total costs of production are given by $C_b = c_b^l(\varphi)x + f$, where $c_b^l(\varphi)$ is the marginal cost of producing x with technology b under regulatory regime l, which we describe further below. Moreover, the business-as-usual technology has an emission intensity of $e_b = \kappa/\varphi$, meaning the production of x creates $z_b(\varphi) = [\kappa x]/\varphi$ units of pollution.

While firms are endowed with the business-as-usual technology, they can choose to upgrade their technology along one of two dimensions. First, they can adopt a state-of-the-art technology (labeled s) that boosts labor productivity[26], lowering marginal costs by a factor $1/\alpha$. The state-of-the-art technology also produces fewer emissions per unit of output. In this case, the emission intensity of production is given by $e_s = \kappa/[\gamma\varphi]$, where $\gamma > 1$, so total pollution from production is $z_s(\varphi) = [\kappa x]/[\gamma\varphi]$. Adopting the state-of-the-art technology requires that firms pay an additional fixed cost f_s, meaning total production costs with the state-of-the-art technology are given by $C_s = c_s^l(\varphi)x + f + f_s$, where $c_s^l(\varphi)$ is the marginal cost of producing x with the state-of-the-art technology in regime l.

The second option for firms is to retrofit their business-as-usual technology so that it has the same emission intensity as the state-of-the-art technology. As such, the emission intensity of a retrofitted plant (e_r) is also $\kappa/[\gamma\varphi]$, meaning the total level of pollution generated

---

[25] For simplicity, $G(\varphi)$ is assumed to be a type-I Pareto distribution such that $G(\varphi) = 1 - \varphi^{-k}$.

[26] For example, by increasing fuel efficiency, as is the case with low-NO_X burners.

by production is $z_r(\varphi) = [\kappa x]/[\gamma\varphi]$. Retrofitting also requires firms to pay a fixed cost (f_r). However, retrofitting does not affect labor productivity, meaning it is less costly than adopting the state-of-the-art technology, so $f_r < f_s$. The total costs of production for a retrofitted plant are given by $C_r = c_r^l(\varphi)x + f + f_r$, where $c_r^l(\varphi)$ is the marginal cost of producing x in regulatory regime l with the retrofitted technology.

4.2 The No-Regulation Equilibrium

Our interest is in understanding the effects of imposing the regulatory structure created by the CWS. Hence, we begin by considering a no regulation regime (labeled n) in which pollution is not regulated. This means labor costs are the only variable costs of production, so $c_b^n(\varphi) = c_r^n(\varphi) = 1/\varphi$ and $c_s^n(\varphi) = 1/[\alpha\varphi]$.

A firm that has drawn a productivity level maximizes profits by deciding whether to stay in the market, and if they stay, choosing how much to produce and what technology to use. Given the structure of consumer preferences, this implies that producing firms set prices at a constant mark-up over marginal costs. Hence, in the absence of regulation, firms that employ business-as-usual and retrofitted technologies charge the same price: $p_b^n(\varphi) = p_r^n(\varphi) = 1/[\rho\varphi]$. If, instead, a firm employs the state-of-the-art technology, it charges $p_s^n(\varphi) = 1/[\rho\alpha\varphi]$.

Firms choose between the three available technologies to maximize profits. If firms employ the business-as-usual technology, profits are given by $\pi_b^n = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} - f$. Profits from employing the retrofitted technology are $\pi_r^n = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} - [f + f_r]$. Finally, profits from choosing the state-of-the-art technology are given by $\pi_s^n = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} \alpha^{\sigma-1} - [f + f_s]$.

Lastly, note that conditional on entering and paying the fixed cost f_ϵ, in expectation firms must earn zero profits. Using this condition solves for all endogenous variables, including industry prices, firm exit, and technology choices. The model's full solution is presented in an online appendix (see Section A).

The exit and technology choices made by firms are highlighted in Figure 2, which depicts the profits associated with adopting each technology as a function of firm productivity[27].

---

[27] To linearize this figure, we show profits as a function of $\varphi^{\sigma-1}$, not φ.

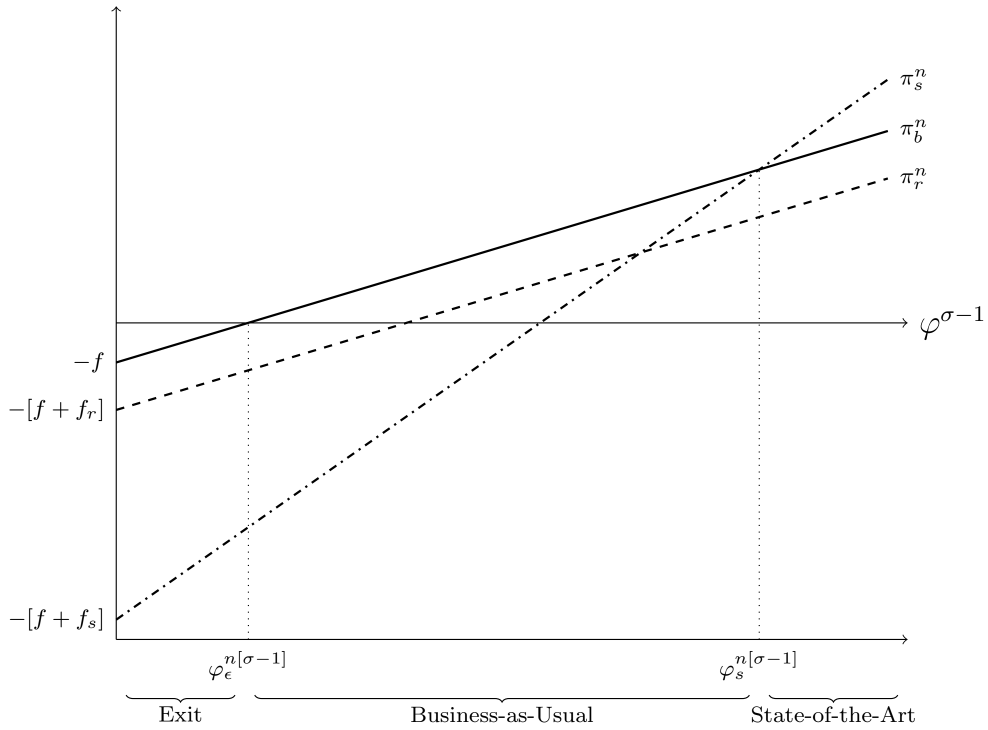

Figure 2: Technology Choices without Environmental Regulation

As the figure shows, for productivity levels below φ_ϵ^n it is unprofitable for a firm to operate using any technology. Hence if a firm has a φ less than φ_ϵ^n, it exits the market. If firms stay in the market, they choose the technology that yields the highest profit. This means that if a firm has a productivity level $\varphi \in \{\varphi_\epsilon^n, \varphi_s^n\}$, then it will produce using the business-as-usual technology. However, if a firm has a productivity level $\varphi > \varphi_s^n$, then the reduction in variable cost created by adopting the state-of-the-art technology is great enough to justify the fixed cost of adoption, meaning that these firms adopt the state-of-the-art technology.

It is also worth noting that, as Figure 2 shows, firms never choose the retrofit technology in the absence of regulation. If firms adopt the retrofit technology, the emission intensity of production falls, but this has no effect on the variable costs of production because pollution is not costly to the firm if it is not regulated. As a result, retrofitting simply lowers firm profits below what can be obtained using the business-as-usual technology by increasing the average costs of production.

4.3 The Effects of Environmental Regulation

To show the intuition behind the process, reallocation, and selection effects caused by a CWS-style regulation, we now consider the effects of adopting environmental regulations similar to those imposed under the CWS. We do so in a partial equilibrium context in which industry prices are held fixed at the no-regulation level. In the online appendix we solve for the full effects of regulation (see Section A).

In this regime (labeled *cws*), the government regulates pollution using a two-part regulatory rule. If a firm uses a clean production process (either the state-of-the-art technology or the retrofitted technology), it is not subject to regulation because it is operating with the lowest emission intensity currently available. As a result, the marginal costs of production for these firms are unaffected by regulation, meaning $c_r^{cws}(\varphi) = c_r^n(\varphi) = 1/\varphi$ and $c_s^{cws}(\varphi) = c_s^n(\varphi) = 1/[\alpha\varphi]$. In contrast, a firm that employs a dirty production process (the business-as-usual technology) is subject to a regulatory constraint in the form of a tax τ on each unit of pollution emitted.[28] Hence, regulation raises the marginal costs of production for these firms, meaning $c_b^n(\varphi) < c_b^{cws}(\varphi) = [1 + \kappa\tau]/\varphi$.

Given that firm prices feature a constant markup, this increase in marginal costs raises the price of output for firms producing with the business-as-usual technology. That is, $p_b^n(\varphi) < p_b^{cws}(\varphi) = [1 + \kappa\tau]/[\rho\varphi]$, and profits are $\pi_b^{cws} = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} [\frac{1}{1+\kappa\tau}]^{\sigma-1} - f$. This means, holding industry prices fixed, the profit from using the business-as-usual technology falls for any level of productivity φ.

This partial equilibrium outcome is depicted in Figure 3, which displays the technological choices made by firms when faced with CWS-type regulation holding industry prices (P) fixed. As the figure shows, a reduction in the profitability of using the business-as-usual technology increases the productivity level for which it is unprofitable to enter the market from φ_ϵ^n to φ_ϵ^{cws}. As such, firms with $\varphi \in \{\varphi_\epsilon^n, \varphi_\epsilon^{cws}\}$ exit in response to regulation. Moreover,

---

[28] Alternatively, we could impose a more realistic regulatory constraint, such as a production cap, without substantively affecting the results. We use a tax for analytical tractability.

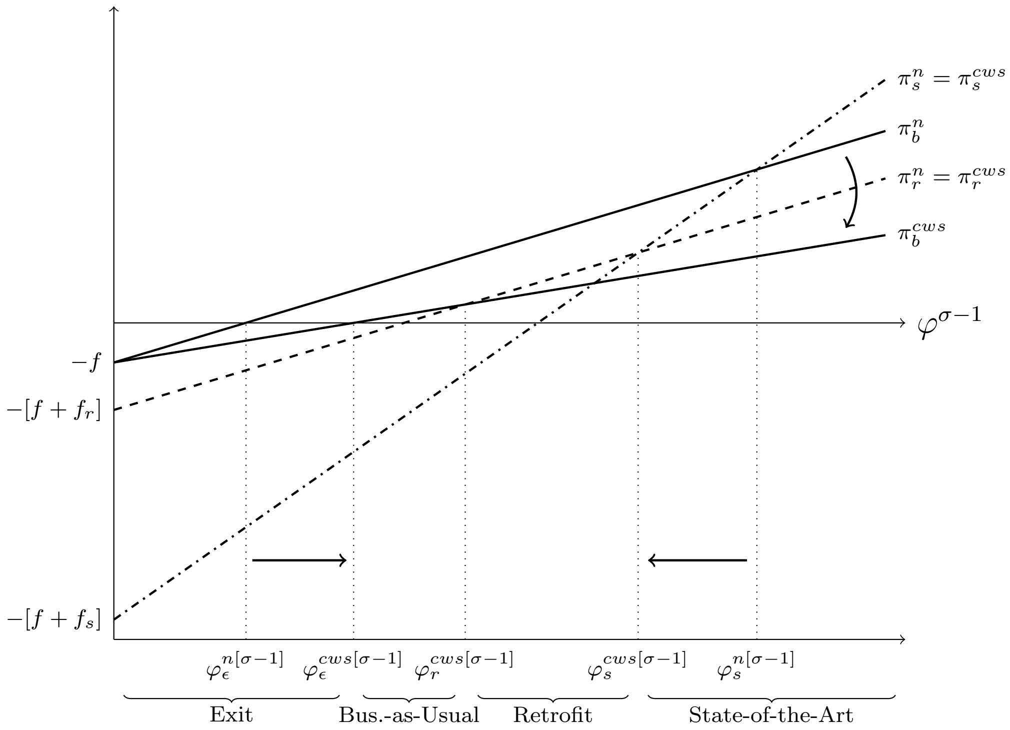

Figure 3: Technology Choices with CWS-Type Environmental Regulation

given the design of regulation, profits from using the retrofitted or state-of-the-art technology do not change. This means the increase in the variable cost of the business-as-usual technology makes technology upgrading a profitable alternative for some firms. As depicted, it is profit maximizing for firms with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^n\}$ to upgrade their technology in response to regulation. For these firms, the benefit of avoided tax payments outweighs the increase in fixed production costs. Similarly, firms with productivity $\varphi \in \{\varphi_c^{cws}, \varphi_s^n\}$ adopt the state-of-the-art technology in response to regulation because it is now profit maximizing to do so.

While Figure 3 clearly highlights how environmental regulations create selection effects by causing firms to exit in response to regulation, the reallocation and process effects are not readily apparent from the figure. As such, we further explore how regulations affect firm revenues and emission intensities to make these additional effects clear.

These effects for firms that survive regulation (those with $\varphi > \varphi_\epsilon^{cws}$) are displayed in Figure 4. This figure depicts the effects of environmental regulation on firm revenues (Panel

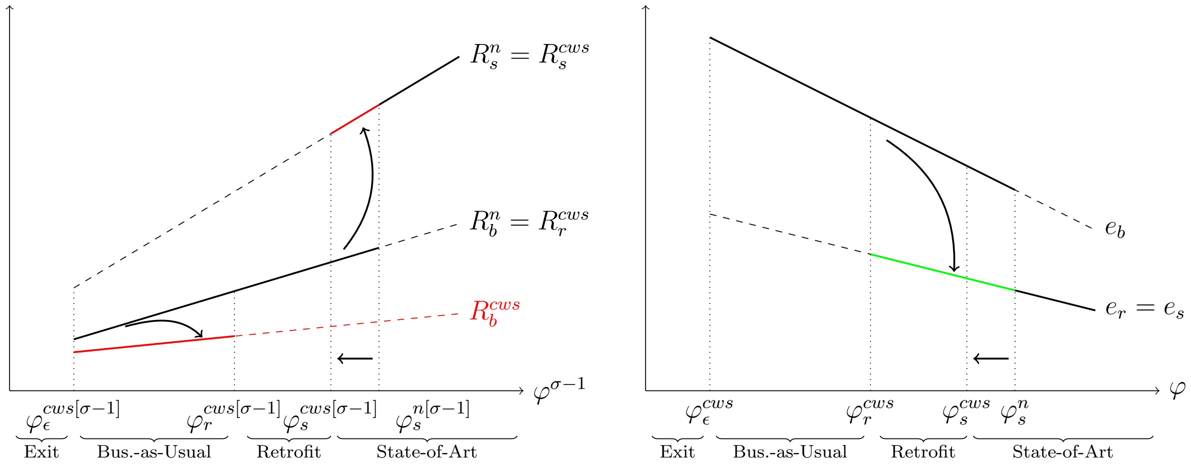

(a) Facility Revenues

(b) Facility Pollution Intensity

Figure 4: Revenues and Pollution Intensity for Surviving Firms with CWS-Type Environmental Regulation

(a)) and emission intensity (Panel (b)) holding industry prices (P) fixed. Both panels show that the most productive firms, with productivity $\varphi > \varphi_s^n$, are unaffected by regulation, as they use a clean production technology in either regime. In contrast, regulation causes the least productive firms, with productivity $\varphi < \varphi_r^{cws}$, to produce less, but with the same pollution intensity. This is because they use the business-as-usual technology under either regime, and variable costs rise under regulation. Lastly, pollution intensity falls for the firms in the middle of the productivity distribution, with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^n\}$. This occurs because they either retrofit or adopt state-of-the-art technology. The retrofitting firms experience no change in output, as their variable costs do not change relative to business-as-usual. However, output increases for the new state-of-the-art adopters, as both their pollution intensity and variable costs fall.

In the online appendix (see Section A) we show the full effects of a CWS-type regulation in which we allow industry prices to respond. In particular, it can be shown that regulation causes the following effects

1. If $f_r > 0$, then some firms exit and average revenues fall.

2. Revenues fall for firms using the business-as-usual technology. These are the least

productive surviving firms.

3. There is an ambiguous effect on revenues for all other surviving firms, but this change in revenues is bounded from below by the reduction in revenues for the business-as-usual technology firms.

4. Pollution intensity falls for all firms that retrofit or upgrade to state-of-the-art technology. These firms are in the middle of the productivity distribution.

5. There is no change in pollution intensity for firms that do not change their technology. These are at the top and bottom of the firm productivity distribution.

In addition to the above results, it can be shown that the fixed cost to retrofit (f_r) plays an important role in determining the channels through which regulation causes an industry to clean-up. In the online appendix, we show decreasing f_r increases the measure of firms that abate in response to regulation and reduces the measure of firms that exit in response to regulation[29]. Thus, when f_r is very small, regulation should primarily cause an industry to clean-up through process effects. Otherwise, reallocation and selection effects will play an important role in an industry's clean-up.

5 Empirics

Our theoretical model provides a number of clear predictions as to how facilities would respond to the CWS. Taken together, these results imply that when the fixed costs of abatement are high, environmental regulations should primarily reduce industry emission intensity via reallocation and selection effects. In contrast, when the fixed costs of abatement are low, the industry clean-up should be driven by process effects. In this section, we explore those plant-level predictions empirically by estimating the CWS' effect on plant pollution intensity, production, and exit. We use the resulting estimates to determine how the process, reallocation, and selection effects created by the CWS have contributed to the clean-up of

---

[29] This last result requires restricting the size of f_s.

Canadian manufacturing.

5.1 Research Design

Given that certain industries and regions were the primary focus of regulation, we identify the causal effects of the CWS by measuring its effects on manufacturing plants that were both located in dirty CMAs and operating in a targeted industry. We do so by using a triple-difference research design that exploits the variation in CWS regulation across time, industries and regions.[30]

Our design begins by comparing the average outcomes of plants in regulated CMAs while regulated (i.e. while violating one of the standards) to their average outcomes while unregulated. This allows us to control for any unobserved time-invariant industry, CMA or plant characteristics that would affect plant pollution emissions. Moreover, in the absence of any other shocks, this comparison would identify the average causal effect of the CWS on pollution emissions. Yet, such absence is unlikely; there is strong reason to believe that a simple before-and-after comparison of affected plants could also capture the effects of regional, industry, or aggregate economic shocks.[31] We discuss each in turn.

To address possible confounding regional shocks, we exploit the fact that each CMA contains manufacturing plants in both regulated and unregulated industries. This allows us to utilize the unregulated plants in a given CMA as a counterfactual for regulated plants in the same location. This will capture the effects of any unobserved time-varying provincial or CMA-level heterogeneity, such as changes in regional economic conditions or concurrent

---

[30] It is worth mentioning that, while plants in dirty CMAs that were operating in a targeted industry were subject to more strict regulation and enforcement, it is possible that other plants in the country were regulated to some degree as a result of the CWS. If this is the case, then our research design produces estimates that give a lower bound on the CWS' effects on the manufacturing sector.

[31] Note that this raises an issue with identifying the effects of any provincial environmental regulation in Canada: who gets regulated and when are unlikely to be randomly assigned if left entirely up to regional authorities. The CWS allows us to overcome this concern by providing within-province variation in regulatory stringency. As a result, the CWS can be thought of as an instrument that allows us to identify the effects of environmental regulation on a select group of plants: those that are regulated because they are in a CMA with air quality above one of the CWS standards. Adopting the language used in the treatment effect literature, these plants are called compliers, and the CWS provides a local average treatment effect of environmental regulation for these plants.

changes in provincial policy that would otherwise confound the effects of the CWS.

The simple before-and-after comparison could also be contaminated as a result of economic shocks that affect individual industries, which could arise due to the effects of increased foreign competition created by international trade, or by revisions to federal policies that target certain sectors. To address these issues we exploit cross-CMA variation in regulation, and utilize the fact that in any particular industry, only plants in areas with poor air quality were subject to stringent environmental policy. This allows us to use the average outcomes from plants in a targeted industry in an unregulated CMA as a counterfactual for the average outcomes of plants from that industry that are located in a regulated CMA. This captures the effects of any industry specific shocks.

The cross-industry and cross-CMA variation in the stringency of environmental regulation also allows us to compare the average outcomes from regulated plants with the average outcomes from plants in non-targeted industries located in unregulated CMAs. These non-targeted plants in unregulated CMAs are not regulated under the CWS, and as such, capture the underlying aggregate trend in pollution emissions. This allows us to control for country-wide shocks, such as aggregate technological change, changes in national policy, or changes in aggregate expenditure due to the 2008 recession.

We estimate the effect of regulation on plant outcomes using the following equation:

$$y_{pict} = \beta_{PM} T_{ict}^{PM} + \beta_{O3} T_{ict}^{O3} + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \tag{3}$$

where y_{pict} is the natural log of the dependent variable of interest (pollution, sales, etc), at plant p, in industry i, located in CMA c, at time t.[32] T_{ict}^{j} is an indicator of treatment for standard j, and takes a value of one for plants that are in industries targeted by the CWS for years in which their CMA exceeds threshold j.

Equation (3) also includes plant (ρ_p), CMA-year (ξ_{ct}), industry-year[33] (λ_{it}) fixed effects

---

[32] We employ the natural log transformation to address the skewness in the distribution of each variable.

[33] The CWS defined the targeted industries at the 3- or 4-digit North American Industry Classification System (NAICS) level. We create an industry indicator that corresponds to either the 3- or 4-digit NAICS

and an error term (ε_{pict}). The plant fixed effects account for any unobserved plant-specific heterogeneity, as well as time-invariant industry and CMA characteristics. The CMA-year fixed effects capture any region specific shocks. The industry-year fixed effects account for any industry-wide events. Finally, the error term captures idiosyncratic changes in outcomes across plants.

The coefficients of interest in Equation (3) are β_{PM} and β_{O3}. β_{PM} measures the average percentage change in outcomes for plants affected by the particulate matter standard relative to those that are not. Similarly, β_{O3} measures the average percentage change in outcomes for plants affected by the ozone standard relative to those that are not. These coefficients are identified from within plant comparisons over time.[34,35]

Changes in plant regulatory status must be plausibly exogenous for this research design to credibly identify the effects of the CWS. There is strong reason to believe this is the case, as variation in regional air quality determines assignment to treatment. As with the CAAAs in the US, regulations are determined by a nationally set air quality threshold, meaning that they are unrelated to differences in local tastes, characteristics or economic conditions (Greenstone, 2002). Moreover, $\mathrm{PM}_{2.5}$ and O_3 are capable of being transported long distances by prevailing wind patterns, meaning that ambient pollution levels in Canada do not solely reflect local economic activity.[36] Indeed, transboundary pollution from the US appears to have been a concern to the federal government over this period. Shortly after the CWS was developed, Canada and the US signed an air quality agreement to address

---

level. All 3-digit industries that contain targeted industries defined at the 4-digit level are grouped at the 4-digit level. The remaining industries are grouped at the 3-digit level.

[34] It is worth noting that regulatory enforcement is applied more stringently to plants that are in regions that currently violate a standard, and that if a region's air quality improves sufficiently, regulation will become less strict. As a result, the variation we are using is from plants in regions that cross one of the CWS thresholds over our sample period. Over our sample, some of these plants move from regulated to unregulated status. This means if plants make changes to production processes that result in permanently lower emissions, then our research design will underestimate the effects of the CWS. As our goal is to be conservative in assessing the effects of the CWS, we view this as an acceptable trade-off.

[35] We are able to separately estimate the effect of both standards because there are cities that exceed one, both, or none of the standards. Of all treated CMA-years in our sample, approximately 80% violated one (and only one) standard, while the remaining 20% violated both standards.

[36] For evidence of how wind patterns shape ambient pollution concentrations in Canada, see, for example, Brankov et al. (2003) or Johnson et al. (2007).

transboundary pollution, Canada's contribution to which involved ensuring the CWS was met (International Joint Commission, 2002). The transboundary nature of pollution means it is unlikely a single plant can directly manipulate their treatment status.

5.2 Data and Measurement

Our analysis relies on a unique confidential micro-dataset that contains information on the $PM_{2.5}$ and NO_X emission intensity of Canadian manufacturing plants. This dataset was created by merging data from two existing sources: the National Pollutant Release Inventory (NPRI) and the Annual Survey of Manufactures (ASM).[37] The NPRI contains information on the emissions of various pollutants from Canadian manufacturing plants. By law, any facility that emits one of the covered pollutants above a minimum threshold must report to the NPRI. The ASM was used as Statistics Canada's manufacturing census until 2012, and it provides longitudinal information on plant sales, production costs, employment, and other plant characteristics for the majority of manufacturing plants in Canada.[38] Plants in these two datasets were linked by Statistics Canada, allowing us to create a longitudinal dataset containing information on $PM_{2.5}$ and NO_X emission intensity as well as other plant characteristics over the period 2004-2010. Additional details on each data source and the construction of the dataset used in our analysis are given in the online appendix (see Section B.1).

Descriptive statistics for the key variables that we employ are reported in Table 2. Each column in Table 2 presents averages and standard deviations for a different sample corresponding to emitters of each pollutant. The first column corresponds to the set of plants that emit $PM_{2.5}$, the second column shows statistics for plants that emit NO_X, and the final column of the table reports summary statistics for the entire sample of plants in the ASM. The statistics in columns one and two are weighted to account for potential sample bias induced by the linking procedure used to match plants across datasets (see Section B.1 of

---

[37] This dataset was created through a collaboration between the Economics and Environmental Policy Research Network, Environment and Climate Change Canada, and Statistics Canada.

[38] The ASM was discontinued in 2012 and was replaced with a repeated cross-section survey.

Table 2: Summary Statistics

	$PM_{2.5}$ (1)	NO_X (2)	Full ASM (3)
Emissions (tonnes)	25.83	262.14	
	(103.43)	(646.14)	
Sales ($1 mill.)	194.62	342.15	11.12
	(890.55)	(1,305.95)	(123.56)
Value Added ($1 mill.)	62.46	102.11	4.29
	(241.82)	(346.27)	(34.34)
Employment	280.11	382.03	35.69
	(634.85)	(868.68)	(125.27)
VA/Worker ($1,000)	200.18	265.41	84.78
	(243.63)	(297.06)	(166.11)
N	6501	3012	309541

Notes: Table reports averages and standard deviations of key variables examined in the main analysis. Each column reports the summary statistics for a different sample. Column (1) is the sample of $PM_{2.5}$ polluters, column (2) is the sample of NO_X polluters, and the final column reports plant characteristics for the entire manufacturing sector. Statistics in columns 1 and 2 are weighted to account for potential sample bias induced by the match of the NPRI and ASM. All monetary values are reported in 2007 Canadian dollars.

the online appendix for further details). Each sample is an unbalanced panel; the sample for $PM_{2.5}$ contains 6501 plant-year observations and the sample for NO_X contains 3012 plant-year observations. For comparison, the final column of the table reports summary statistics for the entire sample of plants in the ASM.

The summary statistics reported in Table 2 suggests that there are systematic differences in plants that emit different types of pollutants. For example, on average, the NO_X sample emitted more pollution, produced more output, had higher employment levels, and had higher labour-productivity levels than the $PM_{2.5}$ sample. This potentially reflects substantial differences in how pollution is produced and abated, given that pollutants are typically produced by a few industries (Greenstone, 2002), and there are substantial differences in the fixed costs of abatement across pollutants (Canadian Council of Ministers of the Environment, 1998; Environment Canada, 2002).

Table 2 also shows that polluters represent the largest plants in the manufacturing sector. Relative to the full manufacturing sector, the sample of plants that emit either $PM_{2.5}$ or NO_X

sell more goods (15 to 30 times on average), employ more workers (7 to 10 times), and have higher value added per worker (2 to 3 times) than the average manufacturing plant.[39] This is, in part, due to the reporting requirements for the NPRI; by law, plants only report if they emit at least one covered pollutant above a minimum threshold level and employ at least 10 individuals or operate an on-site generator (Environment and Climate Change Canada, 2016c). While this means we systematically exclude small facilities, our analysis covers plants that account for the majority of manufacturing pollution in Canada.[40]

5.2.1 Determining Regulatory Status under the CWS

Our analysis also requires determining which CMAs were affected by the CWS. To do so, we use local air quality information from Environment and Climate Change Canada's National Air Pollution Surveillance Program (NAPS), which provides data on hourly monitor-level $PM_{2.5}$ and O_3 concentrations. We use this data to construct CMA-level pollution concentration measures for each year in our sample, where the measures computed are those associated with each standard.[41]

The variation in regulatory status created by changes in ambient air quality is illustrated in Figure 5, which shows the CMAs that changed regulatory status for the $PM_{2.5}$ and O_3 standards. In Figure 5, the red CMAs changed status under both the $PM_{2.5}$ and O_3 standards, the orange CMAs only changed status for the $PM_{2.5}$ standard, the yellow CMAs only changed status for the O_3 standard, and the green CMAs didn't change status under either standard. As the figure shows, there was substantial variation in which CMAs changed their regulatory status over the 2000-2010 period. Of the 149 CMAs in our sample, 23% changed status under the $PM_{2.5}$ standard, 26% changed status under the O_3 standard, 11% changed status under both standards, and 60% never changed regulatory status.

---

[39] This is still true when we consider medians instead of averages.

[40] In addition, the majority of $PM_{2.5}$ and NO_X emitters use an on-site generator or boiler, which means the the employment thresholds are likely not relevant for most of these plants.

[41] For more details on the construction of the hourly pollution concentration measures, see Section B.1 in the online appendix.

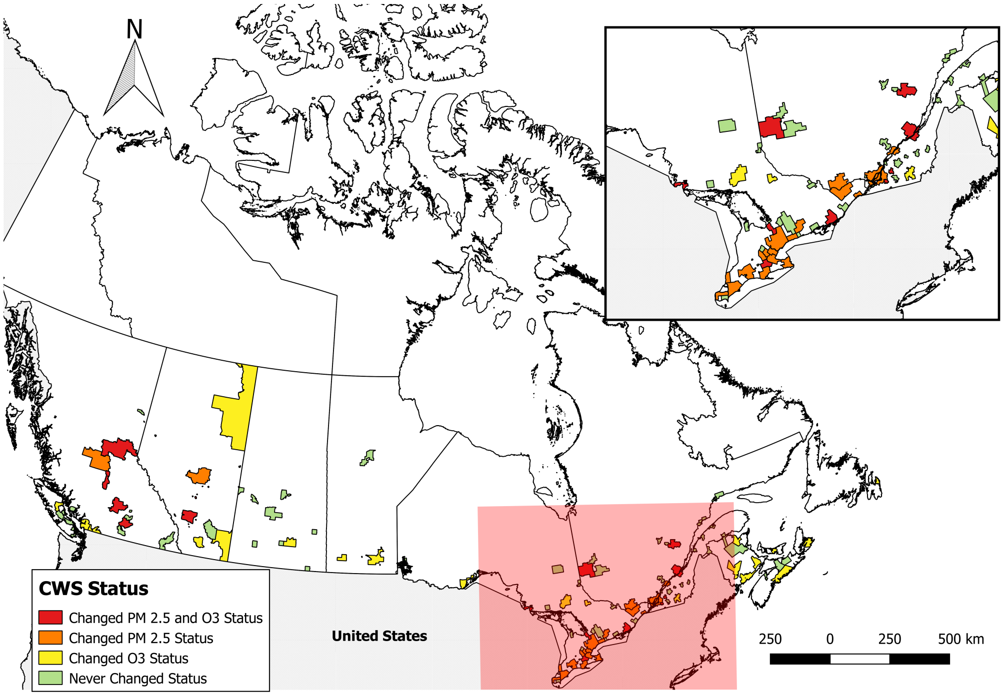

Figure 5: Regulatory Status Changes under the CWS

Notes: Figure depicts $PM_{2.5}$ and O_3 standard status changes for each CMA from 2000 to 2010. Red CMAs changed status under both the $PM_{2.5}$ and O_3 standards. Orange CMAs only changed status for the $PM_{2.5}$ standard. Yellow CMAs only changed status for the O_3 standard. Green CMAs didn't change status under either standard. The mainland United States is shown in light gray. Part of the northern Canadian Territories are trimmed for scale. The inset shows detail on the most densely populated area of Canada, colored in light red on the main map.

5.3 Empirical Results

5.3.1 The CWS and Plant Pollution Emissions

We begin our analysis by estimating the effects of the CWS on the level of pollution emitted by affected Canadian manufacturing plants.[42] We start here for two reasons. First, it provides some indication as to the effectiveness of the CWS; if the regulations were responsible for the reduction in pollution levels documented in Section 2, then we should observe reductions in the emissions of targeted pollutants as a result of the CWS. Second, this also provides us with a means to assess the external validity of our results. As we discussed above, there is little evidence as to the effects of environmental regulation on the emission

---

[42] A related, but distinct question, is to ask what the CWS did to regional air quality. While this is beyond the scope of this paper, in the online appendix we provide descriptive evidence that air quality improved in Canada over this period (see Section B.3).

intensity of manufacturing plants. Focusing on pollution levels allows us to directly compare the effects of the CWS with the effects of other environmental policies.

Table 3 reports our estimates of the effects of the CWS on plant pollution emissions. We estimate Equation (3) for two samples of plants. The first sample (in Panel A) are plants that emit $\mathrm{PM}_{2.5}$, which is the main contributor to $\mathrm{PM}_{2.5}$ pollution. The second sample (in Panel B) are plants that emit NO_{X}, which is the main contributor to O_3 pollution.[43] The first column of each panel reports estimates from a version of Equation (3) that only includes the particulate matter standard. Similarly, the second column reports estimates from a specification that only includes the ozone standard. Finally, column (3) in each panel reports estimates from the specification given in Equation (3). The first row in each panel reports the effect of the $\mathrm{PM}_{2.5}$ standard (β_{PM} in Equation (3)); the second row shows the effect of the O_3 standard (β_{O3} in Equation (3)). The dependent variable in each of these regressions is the natural log of plant pollution emissions for the relevant pollutant. Each regression is weighted to correct for potential sample bias introduced by the procedure used to match plants in the NPRI with plants in the ASM.[44] In all cases, standard errors clustered at the CMA-industry level are reported in parentheses.

The estimates reported in Panel A of Table 3 indicate that the CWS particulate matter regulations led to statistically significant reductions in the emissions of both particulate matter and nitrogen oxide from affected plants. Our baseline estimates for $\mathrm{PM}_{2.5}$, reported in column (3) of Panel A, indicate that the CWS particulate matter regulations are associated with a 15.1% reduction in emissions from affected plants. Our baseline estimates for NO_{X}, reported in column (6) of Panel B indicate that the ozone regulations are associated with 32.5% decrease in emissions from affected plants. The estimates reported in Panels

---

[43] There are other pollutants that may also contribute to $\mathrm{PM}_{2.5}$ and O_3 pollution, including volatile organic compounds and carbon monoxide. In the online appendix we examine the CWS' effects on the emissions of a number of other pollutants (see Section B.3).

[44] In brief, the potential bias happens because the probability of a successful match is positively correlated with a plant's size. If the effects of the CWS vary by plant-size, then relying on the matched data would produce bias estimates. Details on the weighting procedure used to address this can be found in Section B.1 of the online appendix.

Table 3: The Effects of the CWS on Plant Pollution Emissions

	Panel A: $PM_{2.5}$			Panel B: NO_X		
	(1)	(2)	(3)	(4)	(5)	(6)
$PM_{2.5}$ Standard	-0.149**		-0.151**	0.107		0.106
	(0.076)		(0.076)	(0.070)		(0.069)
O_3 Standard		-0.105	-0.113		-0.327*	-0.325*
		(0.164)	(0.164)		(0.183)	(0.179)
R^2	0.175	0.175	0.175	0.310	0.311	0.311
N	6501	6501	6501	3012	3012	3012

Notes: Table reports estimates of the effects of the CWS on plant pollution emissions. Each panel reports results for a different sample of emitters. Each column displays estimates from a different regression. In all cases, the dependent variable is the natural log of pollution emissions. The first row reports the effects of the $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. All regressions include plant, industry-year and CMA-year fixed effects, and are weighted by the inverse of the match probability to control for potential match-induced sample bias. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

A and B also show no statistically significant cross-effects of either standard. That is, O_3 regulation did not significantly affect particulate matter emissions and PM regulation didn't significantly affect NO_X emissions.

We view the results in Table 3 as an exploratory analysis of the CWS effects on plants. While the effect of O_3 on NO_X emitters is only marginally significant, we call attention to these estimates because, as we show later in this section, the average effects of the CWS mask considerable heterogeneity across plants (see Section 5.3.3). Taken together, our evidence suggests O_3 regulation had a meaningful effect on manufacturing plants. Moreover, our theory suggests an average treatment effect is not very illustrative of how plants respond to a policy such as the CWS.

These results are consistent with the few existing estimates of the effects of air quality regulation on pollution emissions from manufacturing plants. For example, Fowlie et al. (2012) find California's NO_X trading program reduced NO_X emissions from regulated plants by between 10% and 30% over the period 1990-2005. Similarly, Gibson (2016) finds that Clean Air Act regulation reduced PM emissions from regulated plants by 38% between 1987 and 2014.[45] This suggests that the CWS had similar effects on pollution levels as the

---

[45] Greenstone (2003) also finds the US Clean Air Act regulation reduced the growth of particulate matter, lead, and VOC emissions from regulated plants by between 4% and 7% over the period 1987-1997.

environmental policies enacted elsewhere.

It is also worth noting that the estimates reported in Table 3 are not simply capturing pre-existing differences in trends across plants or the effects of a negative relationship between a CMA's air quality and the production choices of the plants therein. Moreover, the estimates are robust to accounting for preemptive changes by regulated plants to avoid regulation, plants that account for a significant fraction of their CMA's air pollution, differential trends across large and small emitters, and firm ownership. For the sake of brevity, these results are presented in Section B.2 of the online appendix.

5.3.2 The CWS and the Clean-up of Manufacturing

Having determined the CWS significantly affected plant pollution levels, we now turn to estimating the process, reallocation, and selection effects caused by the CWS. To do this, we start by estimating the effect of the CWS on the emission intensity, output, and exit of affected manufacturing plants. We then use these estimates to determine the implied contribution of the CWS to the clean-up of Canadian manufacturing.

Plant-Level Estimates

In Table 4 we report our estimates of the CWS' effect on the emission intensity of manufacturing plants. As in Table 3, panel A shows estimates of Equation (3) for the sample of plants that emit $PM_{2.5}$ and panel B shows estimates for the NO_X emitters. In each panel, we report estimates from two separate regressions each with a different measure of emission intensity, as well as reproducing our baseline estimates of the CWS' effects on plant pollution levels. The first column shows the CWS effect on pollution levels. In the second column, we show the CWS' effects on emission intensity, measured as the ratio of emissions to total plant shipments (sales), given this is the measure of output used previously in the literature documenting the manufacturing clean-up. In the third column, we measure emission intensity as the ratio of emissions to value-added. Value added may provide a more accurate reflection of the level of productive activity that occurs in each plant (Cherniwchan et al., 2017). However, we focus on the estimates in the second column of each panel, as

Table 4: The Effects of the CWS on Plant Emission Intensity

	Panel A: $PM_{2.5}$			Panel B: NO_X		
	(1)	(2)	(3)	(4)	(5)	(6)
	$PM_{2.5}$	$PM_{2.5}$/Sales	$PM_{2.5}$/VA	NO_X	NO_X/Sales	NO_X/VA
$PM_{2.5}$ Std.	-0.151**	-0.043	-0.013	0.106	0.127	0.333***
	(0.076)	(0.096)	(0.110)	(0.069)	(0.080)	(0.098)
O_3 Std.	-0.113	-0.169	-0.224	-0.325*	-0.286*	-0.200
	(0.164)	(0.169)	(0.189)	(0.179)	(0.153)	(0.157)
R^2	0.175	0.161	0.156	0.311	0.281	0.260
N	6501	6501	6501	3012	3012	3012

Notes: Table reports estimates of the effects of the CWS on plant emission intensity for $PM_{2.5}$ (panel A) and NO_X (panel B) emitting plants. For each group of emitters, the first column reports estimates from a regression of the CWS regulations on the natural log of plant emissions. The second column shows the CWS' effects on the plant emissions-sales ratio, while the third reports estimates from a regression of the regulations on the natural log of the emissions-value added ratio. In all cases, the first row reports the effects of $PM_{2.5}$ regulations, and the second row reports the effects of the O_3 regulations. All regressions include plant, industry-year, and CMA-year fixed effects, and are weighted by the inverse of the NPRI-ASM match probability to control for potential sample bias. Standard errors clustered by CMA-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

our goal is to contribute to a literature that uses shipments as its measure of output.[46] In both cases, the dependent variable is the natural log of emission intensity. The first row in each panel reports the effect of the $PM_{2.5}$ regulation (β_{PM} in Equation (3)) and the second row reports the effect of the O_3 regulation (β_{O3} in Equation (3)). As before, each regression is weighted to correct for potential bias from matching the NPRI and ASM, while standard errors clustered by CMA-industry are reported in parentheses.

The estimates reported in Panel A of Table 4 indicate $PM_{2.5}$ regulation had little-to-no-effect on the emission intensity of plants that emitted $PM_{2.5}$, with an estimated coefficient in column (2) that is relatively small and statistically insignificant. In contrast, the CWS O_3 regulations appear to have caused a significant reduction in NO_X pollution intensity. The estimate reported in column (5) of Table 4 indicate that the CWS ozone regulations are associated with a 28.6% decrease in the level of NO_X emitted per unit of output.[47]

In addition, $PM_{2.5}$ regulation caused a significant increase in NO_X intensity measured

---

[46] In addition, value added may be less precisely reported in our context. This occurs because Statistics Canada is able to use corporate tax filings to check annual shipment amounts reported by plants, but cannot do so for value added.

[47] Though there are no existing estimates to which we can directly compare, Martin et al. (2014) show a carbon tax levied in the United Kingdom led to an 18% drop in energy intensity at affected manufacturing plants.

in value added terms. These results are driven by a very small number of plants that are regulated by the $PM_{2.5}$ standard and emit NO_X, but not $PM_{2.5}$. For these plants, $PM_{2.5}$ regulation caused a large increase in NO_X emissions and decrease in value added. We do not probe these findings further, as they are driven by fewer than ten plants.[48]

$PM_{2.5}$ regulation caused a sizable reduction in plant $PM_{2.5}$ emissions, but had no significant effect on plant emission intensities. On the other hand, the O_3 standard caused a large reduction in NO_X emissions in both levels and pollution intensity. This implies the $PM_{2.5}$ standard must have led to large decreases in output from affected plants, whereas the ozone standard had relatively minor affects on output. We confirm these conclusions by directly estimating Equation (3) on both sales and value added.

Estimates of the effects of the CWS on plant output are given in Table 5 for both $PM_{2.5}$ (Panel A) and NO_X (Panel B) emitters, with each panel reporting estimates from two separate regressions. In the first, we measure output as the value of total plant shipments (sales), and in the second as value added. In both cases, the dependent variable is the natural log of output. The first row in each panel reports the effect of the $PM_{2.5}$ regulation (β_{PM} in Equation (3)) and the second row reports the effect of the O_3 regulation (β_{O3} in Equation (3)). As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure, and standard errors clustered by CMA-industry are reported in parentheses.

The estimates reported in Panel A of Table 5 confirm the $PM_{2.5}$ standard led to a large decrease in output from affected plants that emitted particulate matter. The estimate in column (1) of Panel A indicate the CWS particulate matter regulation is associated with a 10.8% decrease in sales from plants that emitted $PM_{2.5}$. Conversely, the estimates in panel B show the O_3 standard had no statistically significant effects on output.[49]

Lastly, we estimate a variant of our main specification (Equation (3)) in which we compare

---

[48] Dropping these plants yields a point estimate of the $PM_{2.5}$ regulation's effect on NO_X emissions of 0.052 with a standard error of 0.073.

[49] Note that $PM_{2.5}$ regulation also caused a significant reduction in value-added from affected NO_X emitters. As we discuss above, this is driven by a very small number of plants. Thus, we pay little attention to this result.

Table 5: The Effects of the CWS on Plant Output

	Panel A: $PM_{2.5}$		Panel B: NO_X	
	(1) Sales	(2) Value Added	(3) Sales	(4) Value Added
$PM_{2.5}$ Standard	-0.108**	-0.138**	-0.022	-0.227***
	(0.050)	(0.065)	(0.059)	(0.083)
O_3 Standard	0.056	0.111	-0.039	-0.125
	(0.060)	(0.070)	(0.161)	(0.188)
R^2	0.224	0.221	0.265	0.294
N	6501	6501	3012	3012

Notes: Table reports estimates of the effects of the CWS on plant output for $PM_{2.5}$ and NO_X emitting plants. For each panel, each column reports the results of a different regression. In the first column, the dependent variable is the natural log of plant sales. In the second, the dependent variable is the natural log of plant value added. In each panel, the first row reports the the effects of $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. All regressions include plant, industry-year and CMA-year fixed effects are weighted by the inverse of the NPRI-ASM match probability to control for potential sample bias. Standard errors clustered by city-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

the number of plants operating in a treated industry-CMA-year cell to the number operating in an untreated industry-CMA-year cell. That is, we estimate the following regression

$$N_{ict} = \beta_{PM} T_{ict}^{PM} + \beta_{O3} T_{ict}^{O3} + \alpha I(\text{CWS})_{ic} + \xi_{ct} + \lambda_{it} + \varepsilon_{ict}, \tag{4}$$

where N_{ict} is the number of active plants in industry i in CMA c, T_{ict}^{j} is the treatment indicator for standard j (which takes a value of one for industries targeted by the CWS for years in which their CMA exceeds threshold j), $I(\text{CWS})_{ic}$ is an indicator for whether the industry-CMA was ever regulated by the CWS, λ_{it} are industry-year fixed effects, ξ_{ct} are CMA-year fixed effects, and ε_{ict} is an error term that captures idiosyncratic changes in outcomes across industry-regions. The main coefficients of interest (β_{PM} and β_{O3}) show the number of plants that exit an industry-CMA due to the CWS.

As the dependent variable is a count variable, we estimate Equation (4) using both ordinary least squares and Poisson regression. As above, we report estimates for two groups of plants: those that emit $PM_{2.5}$ (Panel A) and those that emit NO_X (Panel B). These results are presented in Table 6, which includes standard errors clustered by CMA in parentheses.

Table 6: The Effects of the CWS on Plant Exit

	Panel A: Emit PM		Panel B: Emit NO_X	
	OLS	Poisson	OLS	Poisson
$PM_{2.5}$ Std.	-1.134**	-0.347**	-0.188	-0.031
	(0.626)	(0.169)	(0.293)	(0.119)
O_3 Std.	0.726	0.142	-0.457	-0.135
	(0.547)	(0.147)	(0.489)	(0.221)
R^2	0.481	0.365	0.443	0.207
N	2776	3023	1252	1582

Notes: Table reports estimates of the effects of the CWS on the number of plants operating in an industry-CMA-year. Panel A shows estimates using plants that emit particulate matter only, and Panel B shows estimates using plants that emit nitrogen oxide only. In each panel, the first column shows the results using OLS estimation and the second column shows results using Poisson estimation. In all cases, the first row reports the effects of $PM_{2.5}$ regulations, and the second row reports the effects of the O_3 regulations. All regressions include industry-year and CMA-year fixed effects. Standard errors clustered by CMA are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

We find a significant reduction in the number of plants operating in an industry-region in response to particulate matter regulation. For example, the estimates in column (1) of Panel A show that $PM_{2.5}$ regulation reduced the number of operating plants in the average affected industry-CMA by 1.134 plants. In contrast, O_3 regulation had no significant effect on plant exit. This is consistent with the predictions of our model, as abatement caries a high fixed cost for $PM_{2.5}$ and a low fixed cost for NO_X.

Aggregate Implications

The implication of the results presented in Table 4, Table 5, and Table 6 is that the CWS contributed to the manufacturing clean-up through different channels for different pollutants. The particulate matter standard primarily caused a reduction in output at regulated plants and plants to exit. In contrast, the ozone standard caused regulated plants to adopt cleaner processes. To quantify the total contribution of the CWS to the manufacturing clean-up we present a simple counterfactual exercise in which we ask how much of the clean-up can be attributed to the process, reallocation and selection effects induced by the CWS. We do this by using our estimates, paired with an empirical analogue of the industry decomposition given by Equation (2), to compute the implied change in manufacturing pollution intensity over our sample that occurred because of each of the CWS channels. We then compare these estimates to the observed change in manufacturing pollution intensity.[50]

To develop an empirical analogue to Equation (2), we follow an approach used in much of the labor literature and consider total changes in emission intensity over time (for a relevant review, see Foster et al. (2001)). To start, add t as an index for time in the decomposition presented in Section 2.1, such that industry i's pollution intensity at time t is $E_{it} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn$, where $e_{it}(n)$ is a plant's pollution intensity, $\lambda_{it}(n)$ is a plant's share of industry output, and n_{it} is the marginal surviving plant. Assuming, for convenience, that plants only exit the industry over time and never enter, then the change in an industry's emission intensity from $t-1$ to t is given by

$$E_{it} - E_{it-1} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn - \int_0^{n_{it}} e_{it-1}(n)\lambda_{it-1}(n)dn - \int_{n_{it}}^{n_{it-1}} e_{it-1}(n)\lambda_{it-1}(n)dn.$$

In the online appendix (see Section B.4), we show that the percentage change in an industry's emission intensity, $\dot{E}_{it} = \frac{E_{it}-E_{it-1}}{E_{it-1}}$, can then be expressed as

$$\begin{aligned}\dot{E}_{it} = &\int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{\lambda}_{it}(n)dn \\ &- \int_{n_{it}}^{n_{it-1}} s_{zit-1}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)\dot{\lambda}_{it}(n)dn,\end{aligned} \tag{5}$$

where $s_{zit-1}(n)$ is plant n's share of industry i's pollution at time $t-1$, and dot notation is used to denote percentage changes. The first three terms of Equation (5) are the process, reallocation, and selection effects that we discussed previously in Section 2.1. The final term is an interaction effect created by the interaction between the process and reallocation effects, and can be thought of as the approximation error in Equation (2) caused by focusing on small, rather than potentially large, changes.

We use our estimates presented above in Section 5.3.2 to construct the four terms on the left-hand side of Equation (5). As such, let $\hat{\beta}_e$, $\hat{\beta}_x$, and $\hat{\beta}_n$ denote our estimates of the

---

[50] For simplicity, we will focus on the direct effects of each standard and ignore any cross-pollutant effects. That is, we ignore the PM standard's effect on NO_X emitters and the O_3 standard's effect on PM emitters.

effects of the CWS on plant pollution intensity (from Table 4), plant output (from Table 5), and selection (from Table 6), respectively. Moreover, recall that, given our identification assumptions, $\hat{\beta}_e$ captures the average change in emission intensity due to the CWS, meaning that we can write

$$\dot{e}_{it}(n) = \begin{cases} \hat{\beta}_e, & \text{if } n \text{ is treated} \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

In addition, an estimate of $\dot{\lambda}_{it}(n)$ and $\int_{n_{it}}^{n_{it-1}} s^z_{it-1}(n)dn$ can be constructed from $\hat{\beta}_x$ and $\hat{\beta}_n$, respectively. In the online appendix (see Section B.4), we show that

$$\dot{\lambda}_{it}(n) = \begin{cases} \frac{\hat{\beta}_x(1-s^{Treat}_{xit-1})+s^{Exit}_{xit-1}}{1-s^{Exit}_{xit-1}+\hat{\beta}_x s^{Treat}_{xit-1}}, & \text{if } n \text{ is treated} \\ \frac{s^{Exit}_{xit-1}-\hat{\beta}_x s^{Treat}_{xit-1}}{1-s^{Exit}_{xit-1}+\hat{\beta}_x s^{Treat}_{xit-1}}, & \text{otherwise,} \end{cases} \tag{7}$$

where s^{Treat}_{xit-1} and s^{Exit}_{xit-1} are the fraction of output in time $t-1$ from treated and exiting plants, respectively. Substituting Equation (6) and Equation (7) into Equation (5) gives estimates of the process, reallocation, and interaction effects. Letting s^{Treat}_{zit-1} be the share of industry i's pollution in time $t-1$ from treated plants, then the process effect is

$$\widehat{PE} = \hat{\beta}_e s^{Treat}_{zit-1}. \tag{8}$$

Similarly, the reallocation effect is given by

$$\widehat{RE} = \frac{s^{Exit}_{xit-1} + \hat{\beta}_x(s^{Treat}_{zit-1} - s^{Treat}_{xit-1})}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}}, \tag{9}$$

and the interaction effect is given by

$$\widehat{IE} = \hat{\beta}_e s^{Treat}_{zit-1} \left[\frac{\hat{\beta}_x(1 - s^{Treat}_{xit-1}) + s^{Exit}_{xit-1}}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}} \right] \tag{10}$$

Table 7: Counterfactual Estimates

	Process Effect (1)	Reallocation Effect (2)	Selection Effect (3)	Interaction Effect (4)	Total (5)
$PM_{2.5}$	0.034	0.109	0.073	-0.004	0.212
NO_X	0.409	0.140	0.085	-0.025	0.610

Notes: Table reports the share of the total change in manufacturing pollution intensity from 2004 to 2010 attributable to each CWS channel. The first row shows estimates for $PM_{2.5}$ and the second row for NO_X. Columns (1) through (4) show the estimates of each channel. Column (5) shows the total across all channels.

To construct an estimate of the selection effect, recall our estimate of $\hat{\beta}_n$ tells us the number of facilities that closed in an industry-CMA cell because of the CWS. Letting N^{Treat} be the number of regulated industry-CMA cells, then the selection effect is

$$\widehat{SE} = \hat{\beta}_n N^{Treat} \bar{s}^{Exit}_{zit-1}, \tag{11}$$

where $\bar{s}^{Exit}_{zit-1}$ is the average exiting plant's share of industry i's pollution in time $t-1$.

In Table 7 we present our estimates of each of the CWS channels relative to the observed change in manufacturing pollution intensity. The first row shows the fraction of the $PM_{2.5}$ clean-up due to the CWS and the second shows the fraction of the NO_X clean-up due to the CWS. Our estimates of the process effect, reallocation effect, selection effect and interaction effect for each pollutant are reported in columns (1)-(4), respectively. Column (5) reports the implied change in manufacturing pollution intensity that can be explained by the CWS.

The results of this exercise show that both the $PM_{2.5}$ and O_3 standards enacted under the CWS played a considerable role in the clean-up of Canadian manufacturing. The estimates in column (5) show that, from 2004 to 2010, the O_3 standard is responsible for 61% of the reduction in manufacturing NO_X intensity and the $PM_{2.5}$ standard is responsible for 21% of the reduction in manufacturing $PM_{2.5}$ intensity. However, the channels responsible varied considerably across pollutants. The process effect, for example, associated with NO_X regulation accounts for almost 41% of the clean-up. In contrast, the process effect accounts for just over 3% of the clean-up for $\text{PM}_{2.5}$. Instead, the $\text{PM}_{2.5}$ regulation primarily reduced aggregate emission intensity through a combination of reallocation and selection effects.

5.3.3 Explaining How Industries Clean-Up

The results presented above show that the channels through which the CWS caused the manufacturing sector to clean-up varied across pollutants. Our theoretical model in Section 4 provides a potential explanation for this: differences in the fixed costs of abatement across pollutants. Indeed, as we discussed in Section 3, engineering assessments of these pollutants argue abatement of NO_X can be accomplished at low-cost, while abatement of $\text{PM}_{2.5}$ pollution typically requires large fixed costs (Environmental Protection Agency, 1999a; Canadian Council of Ministers of the Environment, 1998; Environment Canada, 2002). We now turn to assess this mechanism further and examine other potential explantations for our findings to the extent possible with our data.

Differential Effects by Plant Productivity Level

We begin by testing our model's prediction that there should be large differences across plants in how they respond to regulation when abatement fixed costs are high, but that the responses should be relatively uniform when fixed costs are low. As we cannot observe the fixed costs of abatement directly, this is the most direct test of our hypothesized mechanism.

To test this prediction, we use an approach similar in spirit to that of Bustos (2011) and allow the effects of the CWS to differ across plants on the basis of their initial productivity. That is, we estimate the following regression

$$Y_{pict} = \sum_{q=1}^{3} \beta_{PM}^{Q_q}[T_{ict}^{PM} \times Q_q] + \sum_{q=1}^{3} \beta_{O3}^{Q_q}[T_{ict}^{O3} \times Q_q] + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \qquad (12)$$

where Q_q is an indicator that takes the value one if plants that are in productivity tercile q, T_{ict}^{j} takes a value of one for all plants in targeted industries for years in which their CMA violates standard j, $\beta_j^{Q_q}$ is the treatment effect of standard j on plants in productivity tercile

q, and the remaining variables are as defined for Equation (3).[51]

We use Equation (12) to examine the CWS' effects on plant pollution levels, emission intensity, and sales. Examining pollution levels allows us to assess whether the CWS affected emissions from plants of all productivity levels, whereas examining emission intensity and sales allows us to quantify the channels by which regulation affected each plant.

These results are shown in Table 8. Panel A reports our estimates for $\text{PM}_{2.5}$ emitters; Panel B for NO_X emitters. In each panel, we report estimates from three separate regressions. The first column in each panel shows our estimates from Equation (12) on plant emissions, the second column shows the effects on plant emissions per dollar of sales, and the third on plant sales. Natural logarithms are taken of all dependent variables. The first three rows in each panel report the effects of the $\text{PM}_{2.5}$ regulation (the $\beta_{PM}^{Q_q}$ coefficients in Equation (12)). The first row shows the effect on plants in the lowest productivity tercile, the second row the effect on plants in the middle tercile, and the third row the effects on plants in the highest tercile. Similarly, the final three rows report the effects of the O_3 regulation ($\beta_{O3}^{Q_q}$ in Equation (12)). The fourth row shows the effect on plants in the lowest productivity tercile, the fifth row the effect on plants in the middle tercile, and the sixth row the effects on plants in the highest tercile. As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

The results for the $\text{PM}_{2.5}$ standard show stark differences across $\text{PM}_{2.5}$ plants of different productivity levels. $\text{PM}_{2.5}$ regulation caused a drop in emissions among the bottom two-thirds of the productivity distribution, with a reduction in emissions of 16.3% for low productivity plants and 27.9% for middle productivity $\text{PM}_{2.5}$ emitters. In contrast, $\text{PM}_{2.5}$

---

[51] We construct Q_q by sorting plants in each sample into terciles based on their initial productivity level. We proxy a plant's initial productivity using value added per worker in the first year a plant enters each sample. To account for potential differences in average productivity levels across industries and time, we regress plants' initial productivity levels on entry-year and industry fixed effects, and use the residuals from this regression as our measure of plant productivity. Finally, we divide the distribution of initial productivity residuals into thirds, and place plants into three bins according to their place in the productivity distribution. These bins are used to construct the indicators Q_q. Note that because we construct these bins separately for $\text{PM}_{2.5}$ and NO_X emitters, the composition of plants in each tercile may vary across each pollutant sample.

Table 8: The Effects of the CWS by Plant Productivity Level

	Panel A: $PM_{2.5}$			Panel B: NO_X		
	(1) $PM_{2.5}$	(2) $PM_{2.5}$/ Sales	(3) Sales	(4) NO_X	(5) NO_X/ Sales	(6) Sales
$PM_{2.5}$ Std.						
x Q1	-0.163**	0.038	-0.201***	0.079	0.084	-0.005
	(0.083)	(0.102)	(0.073)	(0.091)	(0.118)	(0.084)
x Q2	-0.279**	-0.251*	-0.028	0.155	0.188	-0.032
	(0.134)	(0.143)	(0.056)	(0.120)	(0.126)	(0.096)
x Q3	-0.023	-0.016	-0.007	0.079	0.109	-0.030
	(0.100)	(0.101)	(0.057)	(0.134)	(0.134)	(0.056)
O_3 Std.						
x Q1	-0.281	-0.353	0.072	-0.457**	-0.412**	-0.045
	(0.210)	(0.222)	(0.074)	(0.207)	(0.207)	(0.205)
x Q2	0.076	0.065	0.011	-0.340**	-0.277*	-0.063
	(0.195)	(0.227)	(0.130)	(0.173)	(0.160)	(0.056)
x Q3	-0.093	-0.150	0.057	-0.183	-0.182	-0.001
	(0.237)	(0.232)	(0.071)	(0.177)	(0.180)	(0.167)
R^2	0.176	0.162	0.226	0.312	0.282	0.266
N	6501	6501	6501	3012	3012	3012

Notes: Table reports estimates of the effects of the CWS where the estimated treatment effects are allowed to vary by plant initial productivity level. Panel A shows the effects on $PM_{2.5}$ emitters and Panel B on NO_X emitters. For each panel, the first column reports estimates from a regression of the CWS regulations on the natural log of plant emissions, the second column shows estimates on the natural logarithm of the emissions-sales ratio, and the third shows estimates on the natural logarithm of plant sales. In all cases, the first row reports the effects of $PM_{2.5}$ regulations for plants in the bottom tercile of their industry's productivity distribution. The second row shows the effects of $PM_{2.5}$ regulations for plants in the middle tercile of their industry's productivity distribution. The third row shows the effects of $PM_{2.5}$ regulations for plants in the top tercile of their industry's productivity distribution. Rows four through six show similar estimates for the O_3 regulations. All regressions include plant, industry-year, and CMA-year fixed effects, and are weighted by the inverse of the NPRI-ASM match probability to control for potential sample bias. Standard errors clustered by CMA-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

regulation had no significant effect on the most productive $PM_{2.5}$ plants, suggesting they were unaffected by regulation.

The results in columns (2) and (3) indicate how the affected $PM_{2.5}$ polluters reduced their emissions varied considerably across the productivity distribution. The drop in emissions among the middle-productivity plants was almost entirely driven by a drop in plant emission intensity, with pollution intensity falling by 25.1%. The drop in emissions from low-productivity plants was driven by a reduction in output, with no significant change in pollution intensity and a 20.1% drop in output. These findings suggest changes in plant pollution intensity driven by regulation played a role in the particulate matter clean-up, however, only among relatively productive plants.

In contrast to the effects of $PM_{2.5}$ regulation, the O_3 standard had relatively uniform effects across NO_X emitters. NO_X emissions fell considerably across the entire productivity distribution, with estimated reductions of between 18-46%, though not significant for the most productive plants. The NO_X clean-up in response to the CWS was primarily driven by changes in plant production techniques, as plant-level changes in emission intensities explain 80-100% of the reduction in emissions.

The results in Table 8 are consistent with our theory, and the hypothesis that the channels of the CWS clean-ups varied across pollutants because of differences in abatement costs.[52] As the abatement of $PM_{2.5}$ requires paying a relatively high fixed cost, only relatively productive plants should choose to do so. These highly productive plants, in turn, experience a reduction in pollution intensity with a relatively small change in output and production inputs. The less productive plants, on the other hand, experience an increase in production costs, leading to a reduction in input use, output, and productivity. In contrast, as NO_X can be abated at a relatively low cost, there are smaller differences across plants of different productivity levels. For both pollutants, the most productive plants in an industry use state of the art technology, and are thus unaffected by the CWS.

Other Margins of Plant Adjustment

Lastly, we examine the effects of the CWS on several additional margins of plant adjustment, including changes in primary inputs, intermediate inputs, and productivity. Doing so allows us to examine a number of alternative explanations as to why the $PM_{2.5}$ and O_3 standards caused the manufacturing sector to clean-up through different channels.

Thus far, the hypothesis we have focused on is that $PM_{2.5}$ and NO_X have different abatement costs, which affects a plant's willingness to adopt cleaner production processes. An

---

[52] This conclusion holds even if we consider alternative specifications in which we split the productivity distribution into quartiles or quintiles, or use a quadratic interaction of plant productivity with the treatment indicators.

alternative hypothesis is that the opportunities for input substitution may vary across pollutants. For example, there could be readily available alternatives to the inputs that create NO_X pollution, but not for the inputs that create $PM_{2.5}$ pollution. If this were the case, then regulation would reduce NO_X intensity but not $PM_{2.5}$ intensity.

Examining the effect of the CWS on input use allows us to asses the above hypothesis. If this hypothesis were true, then the CWS should have caused an increase in spending on inputs for NO_X emitters.[53] In addition, examining the effect of the CWS on input use for plants of different productivity levels allows us to indirectly test our main hypothesis. While our model does not contain intermediate inputs, their use should be positively correlated with output. As our model predicts a reduction in output only for the least productive $PM_{2.5}$ emitters, this should also be accompanied by a reduction in spending on intermediate inputs for these less-productive plants.

The literature on the Porter Hypothesis provides an additional alternative hypothesis. This literature posits environmental regulation could cause an increase in innovative activities and productivity among regulated firms.[54] If the average plant became less productive in response to $PM_{2.5}$ regulation, but more productive in response to NO_X regulation, then this could generate the findings reported in Section 5.3.2. Examining the effect of the CWS on plant productivity allows us to test this hypothesis.

We examine these alternative hypotheses using data on the total number of plant employees[55], spending on both production materials and fuel and energy, value added per worker, and the probability a plant is involved in research and development.

Estimates of the effects of the CWS on productivity and input use for the average manufacturing plant are shown in Table 9. Panel A shows estimates of Equation (3) for $PM_{2.5}$ emitters and Panel B shows estimates for NO_X emitters. In each panel, we report estimates

---

[53] Here we have assumed plants would use the cheapest input in the absence of regulation.

[54] For a recent review of this literature, see Ambec et al. (2013)

[55] Although we do not observe plant capital stock information, given our relatively short period of study we expect capital adjustment to play a minor role in this context. While capital adjustment could play an important role over larger time horizons, the existing literature seems to find limited evidence of capital stock adjustments in response to environmental regulation. See, e.g., Greenstone (2002) and Levinson (1996).

from five separate regressions corresponding to the different mechanisms of interest. Natural logarithms are taken of the dependent variables in columns one to four. The first column shows the CWS' effects on employment, the second spending on materials, the third spending on energy, and the fourth labour productivity. The final column estimates the CWS effect on an indicator for whether the plant is involved in research and development using a linear probability model. In each specification, the first row reports the effect of the $PM_{2.5}$ regulation and the second row reports the effect of the O_3 regulation. As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

We also examine if the effects of the CWS on productivity and input use differ across the initial plant productivity distribution. These estimates are reported in Table 10.[56] Panel A shows the results for $PM_{2.5}$ emitters and Panel B for NO_X emitters. Each column in each panel corresponds to a different dependent variable, each measured in natural logarithms. Each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

As the estimates reported in Table 9 and Table 10 show, the main channels by which the average $PM_{2.5}$ emitting plant responded to $PM_{2.5}$ regulation appears to be through changes in intermediate input use and labor productivity. $PM_{2.5}$ regulation decreased spending on production materials by 11.9%, caused a drop in energy spending (although not significant at conventional levels), and reduced labor productivity (also not significant at conventional levels). $PM_{2.5}$ regulation also caused a significant reduction in labor productivity among $\mathrm{NO_X}$ emitters. There is no evidence of a change in employment or R&D propensity in response to the $PM_{2.5}$ standard.

The estimates of the effects of the $PM_{2.5}$ standard by productivity level are also consistent with our main hypothesis. These results show that the reductions in materials, energy inputs, and labor productivity in response to the $PM_{2.5}$ standard were driven by the least productive

---

[56] The effects on R&D are omitted, but are available upon request.

Table 9: Other Margins of Plant Adjustment

Panel A: $PM_{2.5}$

	Prim. Inputs	Inter. Inputs		Productivity	
	(1) Employment	(2) Materials	(3) Energy	(4) VA/Worker	(5) Pr(R&D)
PM 2.5 Standard	-0.040	-0.119*	-0.086	-0.098	0.033
	(0.064)	(0.064)	(0.056)	(0.073)	(0.040)
O3 Standard	0.071	-0.008	0.224**	0.039	-0.086
	(0.068)	(0.071)	(0.108)	(0.060)	(0.060)
R^2	0.188	0.218	0.151	0.185	0.155
N	6501	6499	6478	6501	6501

Panel B: NO_X

	Prim. Inputs	Inter. Inputs		Productivity	
	(1) Employment	(2) Materials	(3) Energy	(4) VA/Worker	(5) Pr(R&D)
PM 2.5 Standard	0.003	0.039	-0.094	-0.231***	0.061
	(0.069)	(0.077)	(0.093)	(0.085)	(0.060)
O3 Standard	-0.064	-0.069	0.085	-0.062	-0.143
	(0.157)	(0.154)	(0.264)	(0.117)	(0.119)
R^2	0.285	0.276	0.218	0.242	0.248
N	3012	3012	3009	3012	3012

Notes: Table reports estimates of the effects of the CWS on additional margins of adjustment for plants that emit either $PM_{2.5}$ or NO_X. For each group of emitters, each column shows the results of a different regression. The first column reports estimates from a regression of the CWS regulations on the natural log of the number of workers employed at the plant. The second and third columns report estimates of the CWS' effects on the natural log of spending on production materials and fuel and energy, respectively. The fourth column reports estimates of the CWS' effects on the natural log of value added per worker. The final column reports estimates of the CWS' effects on an indicator for whether the plant spends money on research and development, using a linear probability model. In all cases, the first row reports the effects of $PM_{2.5}$ regulations, and the second row reports the effects of the O_3 regulations. All regressions include plant, industry-year, and CMA-year fixed effects, and are weighted by the inverse of the NPRI-ASM match probability to control for potential sample bias. Standard errors clustered by CMA-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

plants. In response to $PM_{2.5}$ regulation, the least productive plants reduced spending on material inputs by 19.4% and energy inputs by 12.5%, and value added per worker fell by 24.7%. $PM_{2.5}$ regulation had no significant effect on these mechanisms at relatively more productive plants. Interestingly, $PM_{2.5}$ regulation had no significant effect on employment for the least productive plants, but reduced employment among the middle-productivity plants. Though output did not fall for the middle-productivity plants, regulation appears to have

Table 10: CWS Mechanisms by Plant Productivity Level

	Panel A: $PM_{2.5}$				Panel B: NO_X			
	(1) Emp.	(2) Materials	(3) Energy	(4) VA/ Worker	(5) Emp.	(6) Materials	(7) Energy	(8) VA/ Worker
$PM_{2.5}$ Std.								
x Q1	0.003	-0.194**	-0.125*	-0.247**	0.141	0.165*	-0.080	-0.418***
	(0.113)	(0.094)	(0.074)	(0.119)	(0.097)	(0.099)	(0.122)	(0.124)
x Q2	-0.093*	-0.044	-0.051	0.058	-0.031	-0.007	-0.189	-0.188**
	(0.055)	(0.070)	(0.073)	(0.061)	(0.079)	(0.094)	(0.146)	(0.092)
x Q3	-0.065	-0.049	-0.041	0.027	-0.116	-0.049	-0.044	-0.076
	(0.072)	(0.095)	(0.098)	(0.089)	(0.100)	(0.122)	(0.102)	(0.111)
O_3 Std.								
x Q1	0.131	0.004	0.311**	-0.058	-0.079	-0.004	-0.181	-0.177
	(0.086)	(0.092)	(0.128)	(0.078)	(0.200)	(0.194)	(0.304)	(0.156)
x Q2	0.024	-0.031	0.252	0.014	-0.057	-0.108	0.190	0.011
	(0.128)	(0.149)	(0.243)	(0.099)	(0.163)	(0.166)	(0.257)	(0.157)
x Q3	0.047	-0.016	0.109	0.136	-0.010	-0.053	0.237	-0.085
	(0.076)	(0.082)	(0.142)	(0.092)	(0.170)	(0.159)	(0.279)	(0.120)
R^2	0.189	0.219	0.152	0.188	0.288	0.277	0.220	0.245
N	6501	6499	6478	6501	3012	3012	3009	3012

Notes: Table reports estimates of the effects of the CWS where the estimated treatment effects are allowed to vary by plant initial productivity level. Panel A shows the effects on $PM_{2.5}$ emitters and Panel B on NO_X emitters. For each group of emitters, each column shows the results of a different regression. The first column reports estimates from a regression of the CWS regulations on the natural log of the number of workers employed at the plant. The second and third columns report estimates of the CWS' effects on the natural log of spending on production materials and energy, respectively. The final column reports estimates of the CWS' effects on the natural logarithm of value added per worker. In all cases, the first row reports the effects of $PM_{2.5}$ regulations for plants in the bottom tercile of their industry's productivity distribution. The second row shows the effects of $PM_{2.5}$ regulations for plants in the middle tercile of their industry's productivity distribution. The third row shows the effects of $PM_{2.5}$ regulations for plants in the top tercile of their industry's productivity distribution. Rows four through six show similar estimates for the O_3 regulations. All regressions include plant, industry-year, and CMA-year fixed effects, and are weighted by the inverse of the NPRI-ASM match probability to control for potential sample bias. Standard errors clustered by CMA-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

made them less labor-intensive, in addition to causing them to adopt cleaner production processes. A potential explanation for this is that the $PM_{2.5}$ process changes may have required new capital investments, thereby changing the plants' capital-labor ratio. Finally, the drop in productivity among NO_X emitters in response to the $PM_{2.5}$ standard appears to be driven by relatively less-productive plants.

The estimates reported in Table 9 and Table 10 also suggest O_3 regulation did not have a significant effect on input use, employment, labor productivity, or R&D propensity at the average affected plant. The exception to this is an increase in energy spending among $PM_{2.5}$ emitters. Allowing the effects of the CWS to vary across plant productivity levels, we still find no significant effect on NO_X emitter employment, input spending, or labor productivity. These results are inconsistent with the two additional hypotheses described above, as neither productivity nor input spending rise in response to regulation, which further suggests our results are driven by the fixed costs of abatement.

6 Conclusion

In this paper, we examine the channels through which environmental regulations have contributed to the "clean-up" of the Canadian manufacturing sector. We start by showing the Canadian manufacturing sector has cleaned-up considerably in recent decades, both in terms of aggregate pollution emissions, and pollution emissions per dollar of output (emission intensity). This clean-up was primarily driven by reductions in industry emission intensity, similar to the clean-ups observed in the U.S. and Europe. We then present a simple model to show how environmental regulation can cause a reduction in an industry's emission intensity through three channels: the reallocation in output across plants, plant entry and exit, or the adoption of cleaner production processes at surviving plants. Finally, we examine how Canadian manufacturing plants responded to a major revision to environmental policy, the Canada-Wide Standards for Particulate Matter and Ozone, and use the resulting empirical estimates to quantify the channels through which environmental regulations have contributed to the manufacturing clean-up. Given the similarity between the clean-ups and regulatory structures in Canada, the US, and Europe, we believe our results provide insights relevant for all three regions.

Our estimates imply that this policy explains approximately 60% of the drop in nitrogen oxide emission intensity of the Canadian manufacturing sector, and approximately 20% of the drop in particulate matter emission intensity. However, how this policy caused manufacturing to clean up varied considerably across pollutants. Over two-thirds of the nitrogen oxide clean-up caused by this policy was due to the adoption of cleaner production processes by surviving plants. In contrast, over 80% of the particulate matter clean-up caused by this policy was due to plant exit and the reallocation of output from regulated to unregulated plants.

These results suggests that transitioning to a less-pollution intensive economy may require large changes in an industry's composition. However, the degree to which an industry's composition will need to change likely depends on the costs of adopting cleaner production processes. When these costs are low, as we argue is the case for nitrogen oxide process improvements, process improvements may yield considerable reductions in industry pollution intensity, even in the absence of plant exit or reallocation across plants.

This work also highlights the importance of linked pollution and production data in assessing the effects of environmental regulation. The mechanisms by which plants respond to regulation appears to vary considerably across emitters of different pollutants, and across plants that emit a common pollutant. Accounting for this heterogeneity is likely important in both the design and assessment of environmental policy, and doing so requires rich information on firm economic and environmental performance.

References

Ambec, S., M. a. Cohen, S. Elgie, and P. Lanoie (2013). The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Review of Environmental Economics and Policy 7*(1), 2–22.

Angle, R. (2014). Ambient air quality objectives. In E. Taylor and A. McMillan (Eds.), *Air Quality Management Canadian Perspectives on a Global Issue*, Chapter 14, pp. 289–301. Dordrecht Heidelberg New York London: Springer.

Anoulies, L. (2017). Heterogeneous firms and the environment: A cap-and-trade program. *Journal of Environmental Economics and Management 84*, 84–101.

Batrakova, S. and R. B. Davies (2012). Is there an environmental benefit to being an exporter? Evidence from firm-level data. *Review of World Economics 148*(3), 449–474.

Becker, R. and V. Henderson (2000). Effects of air quality regulation on polluting industries. *Journal of Political Economy 108*(2), 379–421.

Berman, E. and L. Bui (2001). Environmental regulation and productivity: Evidence from oil refinieries. *Review of Economics and Statistics 83*(3), 498–510.

Boyd, D. R. (2003). *Unnatural Law: Rethinking Canadian Environmental Law and Policy.* Vancouver, BC, Canada: UBC Press.

Brankov, E., R. F. Henry, K. L. Civerolo, W. Hao, S. Rao, P. Misra, R. Bloxam, and N. Reid (2003). Assesing the effects of transboundary ozone pollution between ontario, canada and new york, usa. *Environmental Pollution 123*, 403–411.

Brunel, C. (2016). Pollution offshoring and emissions reductions in e.u. and u.s. manufacturing. *Environmental and Resource Economics*.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms. *American Economic Review 101*(1), 304–340.

Canadian Council of Ministers of the Environment (1998). National Emission Guideline for Commercial/Industrial Boilers and Heaters. `http://www.ccme.ca/files/Resources/air/emissions/pn_1286_e.pdf`. [Online; accessed 25-September-2017].

Canadian Council of Ministers of the Environment (2000a). Canada-Wide Standards for Particulate Matter and Ozone. `http://www.ccme.ca/files/Resources/air/pm_ozone/pmozone_standard_e.pdf`.

Canadian Council of Ministers of the Environment (2000b). Joint Initial Actions to Reduce Pollutant Emissions That Contribute to Particulate Matter and Ground-level Ozone. `http://www.ccme.ca/files/Resources/air/pm_ozone/pmozone_joint_actions_e.pdf/`.

Canadian Council of Ministers of the Environment (2002). Guidance Document on Achievement Determination: Canada-wide Standards for Particulate Matter and Ozone. `http://www.ccme.ca/files/Resources/air/pm_ozone/pn_1330.pdf`.

Cherniwchan, J. (2017). Trade liberalization and the environment: Evidence from NAFTA and U.S. Manufacturing. *Journal of International Economics 105*, 130–149.

Cherniwchan, J., B. R. Copeland, and M. S. Taylor (2017). Trade and the environment: New methods, measurements, and results. *Annual Review of Economics 9*, 59–85.

Copeland, B. R. and M. S. Taylor (1994). North-south trade and the environment. *Quarterly Journal of Economics 109*(3), 755–787.

Environment and Climate Change Canada (2013). National Air Pollution Surveillance Program (NAPS). `http://www.ec.gc.ca/rnspa-naps/`. [Online; accessed 15-August-2015].

Environment and Climate Change Canada (2016a). National Pollutant Release Inventory (NPRI) Toolbox. `https://www.ec.gc.ca/inrp-npri/default.asp?lang=En&n=65A75CDF-1`. [Online; accessed 28-June-2016].

Environment and Climate Change Canada (2016b). Reported Facility Greenhouse Gas Data. `http://www.ec.gc.ca/ges-ghg/default.asp?lang=En&n=8044859A-1`. [Online; accessed 28-June-2016].

Environment and Climate Change Canada (2016c). Summary of National Pollutant Release Inventory Reporting requirements. `https://www.ec.gc.ca/inrp-npri/default.asp?lang=en&n=629573FE-1`. [Online; accessed 28-June-2016].

Environment Canada (2002). Multi-Pollutant Emission Reduction Analysis Foundation (MERAF) for the Iron and Steel Sector. Unpublished report.

Environment Canada (2004). Towards More Innovative Air Quality Management: Proposal for a Pulp and Paper Air Quality Forum. `http://publications.gc.ca/collections/Collection/En84-4-2004E.pdf`. [Online; accessed 25-September-2017].

Environment Canada (2005). Great Lakes Basin Airshed Management Framework Pilot Project. `https://www.epa.gov/sites/production/files/2015-07/documents/great_lakes_basin_airshed_management_framework.pdf`. [Online; accessed 25-September-2017].

Environmental Protection Agency (1998). Particulate Matter Controls. Baghouses and Filters. EPA/452/B-02-001. `https://www3.epa.gov/ttncatc1/cica/files/cs6ch1.pdf`. [Online; accessed 29-August-2017].

Environmental Protection Agency (1999a). Nitrogen Oxides (NOx), Why and How they are Controlled. EPA-456/F-99-006R. `https://www3.epa.gov/ttncatc1/cica/files/fnoxdoc.pdf`. [Online; accessed 29-August-2017].

Environmental Protection Agency (1999b). Particulate Matter Controls. Electrostatic Precipitators. EPA/452/B-02-001. `https://www3.epa.gov/ttncatc1/cica/files/cs6ch3.pdf`. [Online; accessed 29-August-2017].

Environmental Protection Agency (2002). Particulate Matter Controls. Wet Scrubbers for Particulate Matter. EPA/452/B-02-001. `https://www3.epa.gov/ttncatc1/cica/files/cs6ch2.pdf`. [Online; accessed 29-August-2017].

Environmental Protection Agency (2006). Chapte 3. Regulatory Impact Analysis for the Particulate Matter NAAQS. `https://www3.epa.gov/ttn/ecas/regdata/RIAs/Chapter%203--Controls.pdf`. [Online; accessed 3-November-2017].

Environmental Protection Agency (2016). NAAQS Table. `https://www.epa.gov/criteria-air-pollutants/naaqs-table/`. [Online; accessed 10-March-2016].

Forslid, R., T. Okubo, and K. H. Ulltveit-Moe (2014). Why are Firms that Export Cleaner? International Trade and CO2 Emissions. *CESifo Working Paper 4817*.

Foster, L., J. Haltiwanger, and C. J. Krizan (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In D. E. Hulten, Charles and M. Harper (Eds.), *New Developments in Productivity Analysis*. University of Chicago Press.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review 98*(1), 394–425.

Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What do emissions market deliver and to whom? Evidence from Southern California NOx trading program. *American Economic Review 102*, 965–993.

Gibson, M. (2016). Regulation-induced pollution substitution. Working Paper 2016-04, Williams College.

Government of Canada (2003). Clean Air in Canada: Progress Report on Particulate Matter and Ozone. `http://publications.gc.ca/pub?id=9.664805&sl=0`. [Online; accessed 25-September-2017].

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 & 1977 Clean Air Act amendments and the census of manufacturers. *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: Papers and Proceedings 93*(2).

Greenstone, M., J. A. List, and C. Syverson (2012). The effects of environmental regulation on the competitiveness of U.S. manufacturing. Working Paper 18392, National Bureau of Economic Research.

Grether, J.-M., N. A. Mathys, and J. de Melo (2009). Scale, technique and composition effects in manufacturing so_2 emissions. *Environmental and Resource Economics 43*(2), 257–274.

Grossman, G. and A. Krueger (1991). Environmental impacts of a north american free trade agreement. Working Paper 3914, National Bureau of Economic Research.

Hanna, R. (2010). U.S. environmental regulation and FDI: Evidence from a panel of US-based multinational firms. *American Economic Journal: Applied Economics 2*(3), 158–89.

Henderson, J. V. (1996). Effects of air quality regulation. *The American Economic Review 86*(4), 789–813.

Holladay, J. S. and L. D. LaPlue (2017). Emissions Demand Within and Across Sectors: Identifying the Effective Channels. Mimeo.

International Joint Commission (2002). Agreement Between the Government of Canada and the Government of the United States of America on Air Quality. `http://www.ijc.org/rel/agree/air.html`.

Johnson, D., D. Mignacca, D. Herod, D. Jutzi, and H. Miller (2007). Characterization and idenification of trends in average ambient ozone and fine particulate matter levels through trajectory cluster analysis in eastern canada. *Journal of the Air & Waste Management Association 58*(8), 907–918.

Konishi, Y. and N. Tarui (2015). Emissions trading, firm heterogeneity, and intra-industry reallocations in the long ru. *Journal of the Association of Environmental and Resource Economists 2*(1), 1–42.

Lanoie, P. and A. Rochon-Fabien (2012). Promoting Pollution Prevention in Small Businesses: Costs and Benefits of the "Enviroclub" Initiative. *Canadian Public Policy 38*(2), 217–232.

Levinson, A. (1996). Environmental regulations and industry location: International and domestic evidence. In J. Bhagwati and R. Hudec (Eds.), *Fair Trade and Harmonization: Prerequisite for Free Trade?* MIT Press.

Levinson, A. (2009). Technology, international trade, and pollution from u.s. manufacturing. *American Economic Review 99*(5), 2177–2192.

Levinson, A. (2015). A direct estimate of the technique effect: Changes in the pollution intensity of u.s. manufacturing, 1990-2008. *Journal of the Association of Environmental and Resource Economists 2*(1), 43–56.

Martin, L. A. (2012). Energy efficiency gains from trade: Greenhouse gas emissions and india's manufacturing sector. Mimeo.

Martin, R., L. B. de Preux, and U. J. Wagner (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics 117*, 1–14.

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

Popp, D. (2003). Pollution control innovations and the clean air act of 1990. *Journal of Policy Analysis and Management 22*(4), 641–660.

Shadbegian, R. J. and W. B. Gray (2006). Assessing multi-dimensional performance: Environmental and economic outcomes. *Journal of Productivity Analysis 26*(3), 213–234.

Shapiro, J. S. and R. Walker (2015). Why is Pollution from U.S. Manufacturing Declining? The Roles of Trade, Regulation, Productivity, and Preferences.

Statistics Canada (2015). Census metropolitan area (CMA) and census agglomeration (CA). `https://www12.statcan.gc.ca/census-recensement/2011/ref/dict/geo009-eng.cfm/`. [Online; accessed 10-March-2016].

Taylor, E. and A. McMillan (2014). *Air Quality Management Canadian Perspectives on a Global Issue*. Dordrecht Heidelberg New York London: Springerr.

Walker, R. (2013). The transitional costs of sectoral reallocation: Evidence from the Clean Air Act and the workforce. *Quarterly Journal of Economics 128*(4), 1787–1835.

Yujian, L. and L. Bo (2007). A normalized Levenshtein distance metric. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*(6), 1091–1095.

Online Appendices

The online appendices are available for download from this link.

Appendix A Theory

A.1 Equilibrium

In this section we provide the full solution to our theoretical model. We start by showing the productivity level below which firms exit, followed by the productivity levels at which firms choose the retrofitted and state-of-the-art technologies.

In equilibrium, the least productive facility that chooses to produce must earn zero profits. This zero profit condition can be used to find the productivity level below which facilities exit the market.[57] Defining the exit cut-off productivity level under regime l as φ_{ϵ}^{l}, setting $\pi_{b}^{l}(\varphi_{\epsilon}^{l}) = 0$ and rearranging gives the production cut-off productivity levels as $\varphi_{\epsilon}^{n} = [\frac{\sigma f}{I}]^{\frac{1}{\sigma-1}} \frac{1}{\rho P}$ and $\varphi_{\epsilon}^{cws} = [\frac{\sigma f}{I}]^{\frac{1}{\sigma-1}} \frac{1+\tau\kappa}{\rho P}$ for the no regulation and CWS regimes, respectively.

We next turn to a facility's choice of technology, starting with retrofitting then moving onto state-of-the-art technology. Under the CWS regime, for a given productivity level revenues are highest for the state-of-the-art technology, followed by the retrofitted technology, and lowest for the business-as-usual technology. As a result, the least productive facilities use the standard technology, and there is a productivity level at which they are indifferent between the standard and abatement technologies. Letting φ_{r}^{cws} be the retrofitting cut-off productivity level, setting $\pi_{b}^{cws}(\varphi_{r}^{cws}) = \pi_{r}^{cws}(\varphi_{r}^{cws})$ and solving gives $\varphi_{r}^{cws} = [\frac{\sigma f_r}{I\Delta_2}]^{\frac{1}{\sigma-1}} \frac{1}{\rho P}$, where $\Delta_2 = 1 - \frac{1}{[1+\tau\kappa]^{\sigma-1}} > 0$.

As for the choice of state-of-the-art technology, under the no regulation regime, producing facilities decide between using the business-as-usual technology with a low fixed cost, or the state-of-the-art technology with a high fixed cost. Let φ_{s}^{n} be the state-of-the-art technology cut-off productivity level, then setting $\pi_{b}^{n}(\varphi_{h}^{n}) = \pi_{s}^{n}(\varphi_{h}^{n})$ gives $\varphi_{s}^{n} = [\frac{\sigma f_s}{\Delta_1 I}]^{\frac{1}{\sigma-1}} \frac{1}{\rho P}$, where $\Delta_1 = \alpha^{\sigma-1} - 1 > 1$. Under the CWS regime, as revenue under retrofitting exceeds that under business-as-usual technology, facilities decide between using the retrofitted technology and state-of-the-art technology. Defining the state-of-the-art technology cut-off under the CWS regime as φ_{s}^{cws}, setting $\pi_{r}^{cws}(\varphi_{s}^{cws}) = \pi_{s}^{cws}(\varphi_{s}^{cws})$ gives $\varphi_{s}^{cws} = [\frac{\sigma(f_s - f_r)}{\Delta_1 I}]^{\frac{1}{\sigma-1}} \frac{1}{\rho P}$.

Finally, we can express the retrofitting and state-of-the-art technology cut-offs as a function of the exit cut-off. Doing so gives the retrofitting cut-off as $\varphi_{r}^{cws} = \frac{\varphi_{\epsilon}^{cws}}{1+\tau\kappa}[\frac{f_r}{\Delta_2 f}]^{\frac{1}{\sigma-1}}$, the state-of-the-art technology cut-off under no regulation as $\varphi_{s}^{n} = \varphi_{\epsilon}^{n}[\frac{f_s}{\Delta_1 f}]^{\frac{1}{\sigma-1}}$, and the state-of-the-art technology cut-off under the CWS as $\varphi_{s}^{cws} = \frac{\varphi_{\epsilon}^{cws}}{1+\tau\kappa}[\frac{f_s - f_r}{\Delta_1 f}]^{\frac{1}{\sigma-1}}$.[58]

Free Entry Condition

To solve for the model's equilibrium requires using the free entry condition. We first show that revenues for any facility (and thus profits) can be written as a monotonic function of the exit cut-off. To see this, recall revenues for a facility with productivity φ using technology t in regime l are given by $R_{t}^{l}(\varphi) = \frac{I(\rho P)^{\sigma-1}}{c_{t}^{l}(\varphi)^{\sigma-1}}$. Using zero profits for the exiting facility gives $R_{b}^{l}(\varphi_{\epsilon}^{r}) = \sigma f$, which gives revenues as a function of the production cut-off as $R_{t}^{l}(\varphi) = [\frac{c_{b}^{l}(\varphi_{\epsilon}^{r})}{c_{t}^{l}(\varphi)}]^{\sigma-1}\sigma f$. Thus, the main outcomes of interest in the model (exit, technology adoption, and revenues) can be expressed as a function of the exit cut-off, which can be

---

[57] Note that in both regimes the least productive surviving facility must use business-as-usual technology.

[58] To ensure sensible orderings of the technology cut-offs, we impose the assumptions: $f_r > [[1+\tau\kappa]^{\sigma-1} - 1]f$, $f_s > \Delta_1 f$, and $f_s > \frac{\Delta_1 + \Delta_2}{\Delta_2} f_r$.

solved using the free entry condition.

As a result of free entry, in expectation facilities earn zero discounted profits. This means the fixed entry cost paid to draw a productivity parameter, f_ϵ, must equal the present value of expected profits. Under regulatory regime l, free entry implies

$$f_\epsilon = \frac{1 - G(\varphi_u^r)}{\delta}\bar{\pi}^l, \tag{13}$$

where $\bar{\pi}^l$ are a facility's expected profits conditional on surviving in regime r. Note that under the no regulation regime $\bar{\pi}^n = \bar{\pi}_b^n + \bar{\pi}_s^n = \int_{\varphi_\epsilon^n}^{\varphi_s^n} \pi_b^n(\varphi)\frac{g(\varphi)}{1-G(\varphi_\epsilon^n)}d\varphi + \int_{\varphi_s^n} \pi_s^n(\varphi)\frac{g(\varphi)}{1-G(\varphi_\epsilon^n)}d\varphi$. Under the CWS regime, $\bar{\pi}^{cws} = \bar{\pi}_b^n + \bar{\pi}_r^n + \bar{\pi}_s^n = \int_{\varphi_\epsilon^{cws}}^{\varphi_r^{cws}} \pi_b^{cws}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\epsilon^{cws})}d\varphi + \int_{\varphi_r^{cws}}^{\varphi_s^{cws}} \pi_r^{cws}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\epsilon^{cws})}d\varphi + \int_{\varphi_s^{cws}} \pi_s^{cws}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\epsilon^{cws})}d\varphi$.

To solve Equation (13) requires solving for expected profits. Note that the expected profits from technology t under regime l can be written as $\bar{\pi}_t^l = \frac{I(\rho P)^{\sigma-1}}{\sigma}\int_{\hat{\varphi}_t^l}^{\tilde{\varphi}_t^l} c_t^l(\varphi)^{1-\sigma}\frac{g(\varphi)}{1-G(\varphi_\epsilon^l)}d\varphi - f_t\frac{G(\tilde{\varphi}_t^l)-G(\hat{\varphi}_t^l)}{1-G(\varphi_\epsilon^l)}$, where $\hat{\varphi}_t^l$ and $\tilde{\varphi}_t^l$ are the least and most productive facilities using technology t under regime l, respectively. Substituting in the production cut-off productivity levels (φ_ϵ^l) derived from the zero profit condition can remove the aggregate variables I and P. In addition, the abatement and high technology cut-offs can be written as a function of φ_ϵ^l, which gives $\bar{\pi}_t^l$ as a function of one endogenous variable, φ_ϵ^l. After some algebra, it can be shown that expected profits under the no regulation regime are

$$\bar{\pi}^n = \frac{\sigma-1}{k-\sigma+1}f\left[1 + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right], \tag{14}$$

and under the CWS regime are

$$\bar{\pi}^{cws} = \left[\frac{\sigma-1}{k-\sigma+1}\right]f\left[1 + [1+\tau\kappa]^k\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]\right]. \tag{15}$$

To ensure expected profits are positive in both regimes we impose the restriction $k > \sigma-1$.

Substituting $\bar{\pi}^l$ into Equation (13) and using $1 - G(\varphi) = \varphi^{-k}$ solves for the exit cut-off φ_ϵ^l. The production cut-off under the no regulation regime is

$$\varphi_\epsilon^n = \left[\frac{\sigma-1}{k-\sigma+1}\right]^{\frac{1}{k}}\left[\frac{f}{\delta f_\epsilon}\right]^{\frac{1}{k}}\left[1 + \Delta_1^{\frac{k}{\sigma-1}}(\frac{f}{f_s})^{\frac{k-\sigma+1}{\sigma-1}}\right]^{\frac{1}{k}}, \tag{16}$$

and under the CWS regime is

$$\varphi_{\epsilon}^{cws}=\left[\frac{\sigma-1}{k-\sigma+1}\right]^{\frac{1}{k}}\left[\frac{f}{\delta f_{\epsilon}}\right]^{\frac{1}{k}}\left[1+[1+\tau\kappa]^{k}\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]\right]^{\frac{1}{k}}. \tag{17}$$

Given φ_{ϵ}^{l} all remaining endogenous variables (the technology cut-off productivity levels, plant revenues, and industry prices) can be solved.

A.2 The Effect of Environmental Regulation

To examine the effects of regulation we compare equilibrium outcomes under the CWS regime to those under the no regulation regime. Recall that the outcomes of interest under a given regime can be fully characterized by the exit cut-off φ_{ϵ}^{l}. To begin, we show how the exit cut-off changes under this transition.

Exit Cut-Off

The exit cut-off increases under the transition ($\varphi_{\epsilon}^{cws}>\varphi_{\epsilon}^{n}$). To show this, first note that

$$\frac{\varphi_{\epsilon}^{cws}}{\varphi_{\epsilon}^{n}}=\left[\frac{1+[1+\tau\kappa]^{k}\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}{1+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}}\right]^{k}.$$

A sufficient condition to ensure $\frac{\varphi_{\epsilon}^{cws}}{\varphi_{\epsilon}^{n}}>1$ is $[1+\tau\kappa]^{k}\,\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}>\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}$. This sufficient condition can be expressed as $[1+\tau\kappa]^{k}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}>1$, which is satisfied given the assumptions of the model.

Technology Upgrading

Next, we show the effect of regulation on a facility's choice of technology. The effect on technology retrofitting is straightforward. In the no regulation regime, no facility retrofits. In the CWS regime, a facility with $\varphi_r^{cws}\geq\varphi<\varphi_s^{cws}$ retrofits. If the fixed costs satisfy two conditions ($f_s>\frac{\Delta_1+\Delta_2}{\Delta_2}f_r$ and $f_r>\left[[1+\tau\kappa]^{\sigma-1}-1\right]f$), then regulation causes a positive measure of facilities to abate.

Regulation has an ambiguous effect on high technology adoption. To see this, note that the ratio of high technology cut-offs under the CWS and no regulation regimes are given by

$$\frac{\varphi_h^{cws}}{\varphi_h^{n}}=\frac{1+[1+\tau\kappa]^{k}\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}+\Delta_1^{\frac{k}{\sigma-1}}\left(\frac{f}{f_s-f_r}\right)^{\frac{k-\sigma+1}{\sigma-1}}\right]}{[1+\tau\kappa]^{\frac{1}{k}}\left[\frac{f_s-f_r}{f_s}\right]^{\frac{1}{k[\sigma-1]}}\left[1+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}.$$

Thus, $\frac{\varphi_h^{cws}}{\varphi_h^{n}}>1$ if the fixed cost of production, f, is large enough to satisfy

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[1 - \frac{1}{[1+\tau\kappa]^{\frac{1}{k}}}\left[\frac{f_s}{f_s - f_r}\right]^{\frac{1}{k[\sigma-1]}}\right]\left[\frac{\left[[1+\tau\kappa]^{\sigma-1} - 1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\frac{1}{f_r^{\frac{k-\sigma+1}{\sigma-1}}}\left[\frac{f_s}{f_s - f_r}\right]^{\frac{1}{k[\sigma-1]}} + \left[[1+\tau\kappa]^{\frac{k^2-1}{k[\sigma-1]}}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k[k-\sigma+1]}{k[\sigma-1]}} - 1\right]\frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma-1}{\sigma-1}}}\right].$$

If f is relatively small, then regulation increases the measure of facilities using the high technology. If f is large enough, then regulation reduces the measure of facilities using the high technology.

Revenues

To examine the effect of regulation on facility revenues, recall that revenues can be expressed as a function of the exit cut-off: $R_t^l(\varphi) = \left[\frac{c_b^l(\varphi_\epsilon^l)}{c_t^l(\varphi)}\right]^{\sigma-1}\sigma f$. As a result, the change in the exit cut-off caused by regulation is sufficient, once properly weighted, to capture the effect of regulation on a facility's revenues. To see this, note that for a facility using technology t, taking the ratio of equilibrium revenues under the CWS to those under the no regulation regime gives $\frac{R_t^{cws}(\varphi)}{R_t^n(\varphi)} = \left[\frac{c_b^{cws}(\varphi_\epsilon^{cws})}{c_b^n(\varphi_\epsilon^n)}\right]^{\sigma-1}$. Given the marginal cost functions are monotonic transformations of φ, this reduces to the ratio of exit cut-offs.

The one complication to this is that facilities may change technology from one regime to another. As a result, deriving the change in revenues due to regulation for a facility with a given productivity level requires comparing their revenues given their optimal technology choices. We show the full range of possible technological transitions, and resulting revenue changes below.

For facilities that use business-as-usual technology in both the no regulation and CWS regimes, the effect of regulation is given by $\frac{R_b^{cws}(\varphi)}{R_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^n}{\varphi_\epsilon^{cws}}\right]^{\sigma-1} < 1$, where the result follows by $\frac{\varphi_\epsilon^n}{\varphi_\epsilon^{cws}} > 1$ (shown above). Regulation reduces revenues for the facilities that use business-as-usual technology.

For facilities that retrofit business-as-usual technology to technology r, the effect of regulation is given by $\frac{R_r^{cws}(\varphi)}{R_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^n[1+\tau\kappa]}{\varphi_\epsilon^{cws}}\right]^{\sigma-1}$, which is ambiguous but always greater than the ratio for the business-as-usual technology. Note that

$$\frac{\varphi_\epsilon^n\,[1+\tau\kappa]}{\varphi_\epsilon^{cws}} = \left[\frac{[1+\tau\kappa]^{\frac{1}{k}}\left[1 + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}{1 + [1+\tau\kappa]^k\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}\right]^k.$$

With some algebra, one can show that regulation reduces revenues for these facilities if and only if the fixed cost of production is large enough. That is, f must satisfy

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{[1+\tau\kappa]^{\frac{1}{k}}-1}{[1+\tau\kappa]^{\frac{1}{k}}}\right]\Bigg[\left[\frac{\left[[1+\tau\kappa]^{\sigma-1}-1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right]\left[\frac{1}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}$$
$$+\left[[1+\tau\kappa]^{\frac{k^2-1}{k}}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}-1\right]\frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}}\Bigg]^{-1}.$$

For facilities that use state-of-the-art technology in both regimes, $\frac{R_s^{cws}(\varphi)}{R_s^n(\varphi)} = \left[\frac{\varphi_\epsilon^n[1+\tau\kappa]}{\varphi_\epsilon^{cws}}\right]^{\sigma-1}$, which is the same as the retrofitting facilities above. Thus, regulation reduces revenues for these facilities if and only if the fixed cost of production is large enough

If $\frac{\varphi_s^{cws}}{\varphi_s^n} > 1$, then regulation causes some facilities to switch from standard technology to high technology. Revenues must rise for these facilities. To see this, note that the effect of regulation on revenues is given by $\frac{R_s^{cws}(\varphi)}{R_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^n[1+\tau\kappa]\alpha}{\varphi_\epsilon^{cws}}\right]^{\sigma-1}$. This is greater than one if and only if

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{\alpha^{\frac{1}{k}}\,[1+\tau\kappa]^{\frac{1}{k}}-1}{\alpha^{\frac{1}{k}}\,[1+\tau\kappa]^{\frac{1}{k}}}\right]\Bigg[\left[\frac{\left[[1+\tau\kappa]^{\sigma-1}-1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right]$$
$$\left[\frac{1}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\frac{1}{\alpha^{\frac{1}{k}}}+\left[[1+\tau\kappa]^{\frac{k^2-1}{k}}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}-\alpha^{\frac{1}{k}}\right]\frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}}\frac{1}{\alpha^{\frac{1}{k}}}\Bigg]^{-1},$$

which must be satisfied if $\frac{\varphi_s^{cws}}{\varphi_s^n} > 1$ (the only condition under which this scenario is plausible).

If $\frac{\varphi_s^{cws}}{\varphi_s^n} < 1$, then regulation causes some facilities to downgrade from state-of-the-art to the retrofitted technology. The change in revenue for these facilities is $\frac{R_r^{cws}(\varphi)}{R_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^n[1+\tau\kappa]}{\varphi_\epsilon^{cws}}\alpha\right]^{\sigma-1}$, which is less than one if and only if

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{[1+\tau\kappa]^{\frac{1}{k}}-\alpha^{\frac{1}{k}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right]\Bigg[\left[\frac{\left[[1+\tau\kappa]^{\sigma-1}-1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right]\left[\frac{1}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\alpha^{\frac{1}{k}}$$
$$+\left[[1+\tau\kappa]^{\frac{k^2-1}{k}}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}-\frac{1}{\alpha^{\frac{1}{k}}}\right]\frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}}\alpha^{\frac{1}{k}}\Bigg]^{-1}$$

Notice that this cut-off value for f is lower than that required to ensure revenues for retrofitters falls. There is a range of values for f for which facilities that retrofit business-as-usual technology gain revenue while those that switch from state-of-the-art technology to the retrofitted technology lose revenue. Note also that imposing $\alpha > [1+\tau\kappa]$ is sufficient to guarantee $\frac{R_r^{cws}(\varphi)}{R_s^n(\varphi)} < 1$.

Pollution Intensity

The effect of regulation on facility pollution intensity is determined by the adoption of new technology. For facilities that use the business-as-usual technology in both regimes, pollution intensity is unaffected. Facilities that change from business-as-usual to either the retrofitted or state-of-the-art technology experience a decline in pollution intensity, given by $\frac{e_r^{cws}(\varphi)}{e_b^n(\varphi)} = \frac{e_s^{cws}(\varphi)}{e_b^n(\varphi)} = \frac{1}{\gamma} < 1$. Lastly, there is no change in pollution intensity for facilities that downgrade from state-of-the-art to retrofitted technology, or use state-of-the-art technology in both regimes. That is, there is no change in pollution intensity for the facilities using state-of-the-art technology under the no regulation regime.

Average Revenues

Average revenues fall under the CWS regime. To see this, note that average revenues under regime n can be written as $\bar{R}^n = \int_{\varphi_\epsilon^n}^{\varphi_s^n} R_b^n(\varphi)g(\varphi)d\varphi + \int_{\varphi_s^n} R_s^n(\varphi)g(\varphi)d\varphi$, where R_b^n is the revenue for a business-as-usual facility and R_s^n is the revenue for a state-of-the-art facility. Substituting the expression for revenues gives $\bar{R}^n = \frac{k}{[\varphi_\epsilon^n]^{\sigma-1}}[\int_{\varphi_\epsilon^n}^{\varphi_s^n} \varphi^{\sigma-2-k}d\varphi + \alpha^{\sigma-1}\int_{\varphi_s^n} \varphi^{\sigma-2-k}d\varphi] = \frac{k}{k-\sigma+1}\frac{\Lambda^n}{[\varphi_\epsilon^n]^{\sigma-1}}$. Similarly, average revenues under regime cws can be written as $\bar{R}^{cws} = \int_{\varphi_\epsilon^{cws}}^{\varphi_r^{cws}} R_b^{cws}(\varphi)g(\varphi)d\varphi + \int_{\varphi_r^{cws}}^{\varphi_s^{cws}} R_r^{cws}(\varphi)g(\varphi)d\varphi + \int_{\varphi_s^{cws}} R_s^{cws}(\varphi)g(\varphi)d\varphi$. Substituting the expression for revenues and rearranging gives $\bar{R}^{cws} = \frac{k}{k-\sigma+1}\frac{\Lambda^{cws}}{[\varphi_\epsilon^{cws}]^{\sigma-1}}$. Taking the ratio of revenues under the cws regime to those under regime n gives $\frac{\bar{R}^{cws}}{\bar{R}^n} = \frac{\Lambda^{cws}}{\Lambda^n}[\frac{\varphi_\epsilon^n}{\varphi_\epsilon^{cws}}]^k$. But $\frac{\varphi_\epsilon^n}{\varphi_\epsilon^{cws}} = \frac{\Lambda^n}{\Lambda^{cws}}$, which means $\frac{\bar{R}^{cws}}{\bar{R}^n} = [\frac{\Lambda^n}{\Lambda^{cws}}]^{k-1}$. This expression is less than one, as $k > \sigma - 1 > 1$.

Role of Retrofitting Fixed Costs

Lastly, we examine the role of retrofitting fixed costs, f_r. Lowering f_r lowers the retrofitting cut-off productivity level. As a result, under the CWS regime, lower f_r increases the measure of facilities that switch from business-as-usual to retrofitted technology. To see this, differentiate the retrofitting cut-off with respect to f_r to get

$$\begin{aligned}
\frac{\partial\varphi_r^{cws}}{\partial f_r} &= \left[\frac{1}{1+\tau\kappa}\right]\left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}\left[\frac{\partial\varphi_\epsilon^{cws}}{\partial f_r} + \left[\frac{1}{\sigma-1}\right]\frac{\varphi_\epsilon^{cws}}{f_r}\right] \\
&= \left[\frac{1}{1+\tau\kappa}\right]\left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}\left[\frac{1}{k[\varphi_\epsilon^{cws}]^k}\frac{\partial[\varphi_\epsilon^{cws}]^k}{\partial f_r} + \left[\frac{1}{\sigma-1}\right]\frac{\varphi_\epsilon^{cws}}{f_r}\right].
\end{aligned}$$

Thus, $\frac{\partial\varphi_r^{cws}}{\partial f_r} > 0$ if and only if

$$\frac{\partial\left[\varphi_\epsilon^{cws}\right]^k}{\partial f_r} > -\left[\frac{k}{\sigma-1}\right]\frac{\left[\varphi_\epsilon^{cws}\right]^k}{f_r}, \tag{18}$$

where $\frac{\partial[\varphi_\epsilon^{cws}]^k}{\partial f_r} = [\frac{f}{\delta f}][1+\tau\kappa]^k f^{\frac{k-\sigma+1}{\sigma-1}}\left[[\frac{\Delta_1}{f_s-f_r}]^{\frac{k}{\sigma-1}} - [\frac{\Delta_2}{f_r}]^{\frac{k}{\sigma-1}}\right]$. With some algebra, one can show that Equation (18) reduces to

$$\left[\frac{k}{k-\sigma-1}\right]\frac{1}{f_r}\left[\frac{1}{f}\right]^{\frac{k-\sigma+1}{\sigma-1}}\left[\frac{1}{1+\tau\kappa}\right]^{\sigma-1}+\left[1+\left[\frac{k}{k-\sigma-1}\right]\frac{1}{f_r}\frac{1}{f_s-f_r}\right]\left[\frac{\Delta_1}{f_s-f_r}\right]^{\frac{k}{\sigma-1}}> \\ -\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{\Delta_2}{f_r}\right]^{\frac{k}{\sigma-1}},$$

which is always satisfied.

In addition, lowering f_r lowers the exit cut-off under the CWS regime if f_s isn't too large. Differentiating φ_ϵ^{cws} with respect to f_r gives

$$\frac{\partial\varphi_\epsilon^{cws}}{\partial f_r}=\left[\frac{\sigma-1}{k-\sigma+1}\right]^{\frac{1}{k}}\left[\frac{f}{\delta f_\epsilon}\right]^{\frac{1}{k}}\left[\Lambda^{cws}\right]^{\frac{1-k}{k}}\left[1+\tau\kappa\right]^{k}\left[\frac{k-\sigma+1}{\sigma-1}\right]f^{\frac{k-\sigma+1}{\sigma-1}} \\ \left[\Delta_2^{\frac{k-\sigma+1}{\sigma-1}}\left[\frac{1}{f_s-f_r}\right]^{\frac{k-2[\sigma-1]}{\sigma-1}}-\Delta_1^{\frac{k-\sigma+1}{\sigma-1}}\left[\frac{1}{f_r}\right]^{\frac{k-2(\sigma-1)}{\sigma-1}}\right],$$

which is greater than zero if and only if $f_s < \left[1+[\frac{\Delta_1}{\Delta_2}]^{\frac{k-[\sigma-1]}{k-2[\sigma-1]}}\right]f_r$. Note that if $k > 2[\sigma-1]$ this means the model requires both a maximum and minimum constraint on f_s to produce the above result and maintain $\varphi_r^{cws} < \varphi_s^{cws}$. If $k < 2[\sigma-1]$, then imposing $f_s > [\frac{\Delta_1+\Delta_2}{\Delta_2}]f_r$ ensures both results.

Appendix B Empirics

B.1 Data Appendix

Micro Data

Our micro-data was created by merging two existing datasets: the National Pollutant Release Inventory (NPRI) and the Annual Survey of Manufactures (ASM). We describe each here, and provide details on how these two sources were matched.

The NPRI is Canada's main source for pollution information, and the only source of air pollution micro-data in the country. It records plant-level pollution activities for over 300 pollutants, including criteria air contaminants, toxins, and heavy metals. All plants in Canada that emit at least one covered pollutant (above that pollutant's minimum emissions threshold) and employ at least 10 individuals are required by law to report to the NPRI (Environment and Climate Change Canada, 2016c). In addition, all plants that use stationary combustion equipment must report to the NPRI, regardless of their number of employees. Failure to report, or the submission of incorrect data, may result in a penalty of between \$25,000 and \$12,000,000.[59] The federal ministry of environment performs inspections to confirm the completeness of submitted data. From 2000 to 2010, there were 2,198 NPRI inspections completed, resulting in 1,270 written warnings.[60].

---

[59] For details, see sections 272 and 273 of the Canadian Environmental Protection Act.

[60] These figures are from the authors' calculations computed using data from the Canadian Environmental Protection Act annual reports. These reports are available here: `http://www.ec.gc.ca/lcpe-cepa/default.asp?lang=En&n=477203E8-1`

For each pollutant, plants are required to report their releases by medium (to air, water, and land), quantities sent for disposal and recycling, methods used to compute releases, and abatement activities[61]. Detailed guidelines on how to compute emissions for each pollutant are provided for each sector and production activity (for a detailed list by sector, see: Environment and Climate Change Canada (2016a)). Each plant is also required to report a number of characteristics, including plant name, business number, industry, and location.

The ASM was used as Statistics Canada's manufacturing census until 2012, and provides longitudinal information for the majority of manufacturing plants in Canada.[62] Before 2004, every manufacturing plant in the country was sampled annually. The sampling strategy changed in 2004 so that a new random sample of the smallest plants was taken in each year, rather than collecting information for every plant annually. All large plants were sampled annually. For the plants that weren't sampled yearly, where possible, administrative tax files were used to fill-in missing sales and expenditure data. We restrict our analysis to 2004 onwards to avoid any issues with the methodological change.

The ASM collects information on sales, production costs (including energy expenditures by fuel type), employment, the distribution of sales by province and country, and plant characteristics (including plant name, business number, industry, and location). Sales, value added, and cost variables are expressed in 2007 Canadian dollars using industry price deflators from Statistic Canada's Industry Multifactor Productivity Program.

To match the two datasets, Statistics Canada developed a cross-walk file between them following a multi-stage linking strategy. The majority of plants were linked using business number, year, and location information. A second round of linking was done using two-variable combinations of the above three variables (business number and location, etc). A final round of linking was done using plant names. Approximately 80% of manufacturing plants in the NPRI were successfully linked to the ASM.

There are two potential issues that arise from the imperfect link between the NPRI and the ASM. The first issue is to do with the representativeness of the matched sample. If the probability of a successful match is non-random, then the matched sample will not be representative of the universe of polluters. This means descriptive statistics from the matched sample will not be reflective of polluters in general. Rather they will be informative about the subset of polluters that were successfully matched.

The second issue is more problematic, as it could lead to biased estimates of the CWS' effects. This issue arises if the match probability is correlated with the CWS' treatment effect. Note that if the effect of the CWS is homogenous, then the match probability cannot be correlated with treatment, and the estimated effect of the CWS from the matched data will be an unbiased estimate of the true effect of the CWS. That is, this issue only arises when the effect of treatment varies across plants.

In the case of the CWS, there is substantial heterogeneity in the treatment effects. As we show in the main body of the paper, the treatment effects vary by plant productivity. Moreover, plant productivity is correlated with plant size, and the probability of a successful match also appears to be correlated with plant size. As a result, the match probability is potentially correlated with the treatment effect. This sample bias induced by the imperfect

---

[61] Reporting of abatement activities was discontinued in 2010.

[62] The ASM was discontinued in 2012 and was replaced with a repeated cross-section survey.

match should be addressed so as to obtain unbiased estimates of the CWS' effects. We correct for this bias using a simple weighting strategy.

To see how weighting corrects for this sample bias, consider the estimation of a treatment effect, β, that varies across two groups, g_1 and g_2. Let the treatment effect in g be given by β^g. The average treatment effect is a weighted average of the two groups' treatment effects

$$\beta = Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2}, \tag{19}$$

where $Pr(g)$ is the probability an observation is in group g.

The treatment effect in the matched sample is given by

$$\begin{aligned} \beta^{match} &= Pr(g_1|match)\beta^{g_1} + Pr(g_2|match)\beta^{g_2} \\ &= \frac{Pr(match|g_1)Pr(g_1)}{Pr(match)}\beta^{g_1} + \frac{Pr(match|g_2)Pr(g_2)}{Pr(match)}\beta^{g_2}, \end{aligned} \tag{20}$$

where the second equality follows by Bayes' theorem, $Pr(match)$ is the probability of a successful match, and $Pr(match|g)$ is the probability an observation in group g is successfully matched.

If the probability of a successful match is random, then $Pr(match|g_1) = Pr(match|g_2) = Pr(match)$, and $\beta^{match} = Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2} = \beta$. That is, there is no bias and the imperfect match does not matter. If the probability of a successful match is non-random, then $Pr(match|g_1) \neq Pr(match|g_2)$, and $\beta^{match} \neq \beta$.

Now, suppose the match probabilities ($Pr(match|g)$) were known for each group, and were used to construct weights defined as the inverse of the probability an observation was successfully matched. In this case, the weight for group g would be $\omega_g = \frac{Pr(match)}{Pr(match|g)}$. Clearly, performing a simple weighted regression on the matched data using these weights would produce an unbiased estimate of the true treatment effect. The weighted treatment effect from the matched data would be

$$\begin{aligned} \beta^{match,weighted} &= \omega_{g1}Pr(g_1|match)\beta^{g_1} + \omega_{g2}Pr(g_2|match)\beta^{g_2} \\ &= Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2}, \end{aligned} \tag{21}$$

which is the true treatment effect, β.

The real issue is that these match probabilities are generally not known. In our case, however, we can recover a reasonable approximation of these probabilities because our concern is that the match probabilities and treatment effects vary by plant size, and we observe a reasonable measure of size (pollution) for both the universe of polluters and the matched sample.

We operationalize this weighting procedure by splitting the distribution of pollution into ten evenly spaced bins in both the full NPRI and the matched NPRI-ASM. We then compute the match probability in each bin as the number of plants in that bin in the matched sample divided by the total number of plants in that bin in the full NPRI. The weights are taken as the inverse of this ratio for each bin. We compute these weights for each of the four pollutant

Table A1: Mean Emissions in Matched Dataset

	Universe of Polluters	Matched Sample	
		Weighted	Unweighted
$PM_{2.5}$ Emissions	23.0	+12%	+26%
NO_X Emissions	276.4	-5%	+1%

Notes: Table reports the mean emissions in tonnes from the universe of polluters in the NPRI and the matched NPRI-ASM samples. Column 1 shows the mean emissions from the full NPRI. Column 2 shows the difference in mean emissions in the matched data with weighting. Column 3 shows the difference in mean emissions in the matched data without weighting.

samples.

To show the effect of our weighting procedure, Table A1 compares the average plant emissions of each of the CWS pollutants from the full NPRI, the unweighted matched sample, and the weighted matched sample. The first column shows the mean emissions for the universe of polluters, and the second the percentage differences between the mean emissions in the matched sample using our weighting procedure and the universe of polluters. The third column shows the percentage differences between the mean emissions in the matched sample without weighting and the universe of polluters.

The match problem appears most severe for particulate matter emissions, with unweighted average emissions approximately 25% higher in the NPRI-ASM matched data than in the universe of polluters. Weighting reduces this over-estimate considerable, to 12% for $PM_{2.5}$. The match problem is relatively small for NO_X emissions, and weighting has a relatively small effect on the average emissions of these pollutants.

Air Quality Data

The NAPS is a network of 286 air quality monitoring stations located across Canada, and is Canada's main source for air quality data. Each monitoring station is operated by a provincial authority, and the federal environment ministry oversees the network. Hourly monitor-level data is available from 1974 onward for ozone, most Criteria Air Contaminants (including fine and large scale particulate matter), and some heavy metals (for data, see: Environment and Climate Change Canada (2013)).

We construct regional air quality measures using the following methods. For $PM_{2.5}$, we construct the 98^{th} percentile of each CMA's 24-hour concentration in a given year.[63] For O_3, we construct the 4^{th} highest 8-hour concentration reported in a CMA in a given year.[64] For any CMA that contains more than one monitor, we follow the rule defined by the CWS and compute the average pollution concentration across all monitors for the $PM_{2.5}$ measurements and the maximum concentration for the O_3 measurements (Canadian Council of Ministers

---

[63] The 24-hour concentration is the 24-hour average taken from midnight to midnight for each day. This calculation collapses the hourly data to the daily frequency.

[64] For each monitor, running eight-hour averages are computed for each hour, and reported as the value associated with the last hour used in the calculation. That is, for January 1st, 2000, there is no reported value from midnight to 7am, the 8am value is the average from midnight to 8am, the 9am value is the average from 1am to 9am, etc.

of the Environment, 2002, p. 12).

B.2 Robustness

There are two potential concerns with the identification strategy we've pursued. The first is that the assignment mechanism may not be exogenous because treatment depends on the air quality of the city in which a plant resides, and the emissions of relatively large plants may directly affect their CMA's air quality. The second is that treatment of one plant may affect the outcomes of another because they are owned by the same company. In this section we present the results from a number of robustness checks aimed at addressing these concerns. In the interest of space we only provide the estimation results for the average effects of the CWS on emissions, but for each robustness check we also describe the CWS' effects on output and by plant-productivity level.[65]

Main Robustness

We first present the results of our two core robustness checks. In brief, we show that the above results are not simply driven by a non-linear relationship between a CMA's air quality and the production choices of the plants therein, and that there is no evidence of a difference in outcomes between our treatment and control groups before treatment occurs. These results show, in effect, that there is a break in trend between the treatment and control groups at the time at which treatment starts, and at the air quality threshold values used in the standards.

We first test that the above findings are not the result of a non-linear relationship between a CMA's air quality and plant production choices.[66] We do this by estimating a flexible triple-difference regression in which we allow the potential effect of treatment to vary by the air quality of the CMA in which the plant is located. If, as we've claimed, being above a CWS threshold results in greater regulatory stringency, then flexibly estimating our triple-difference regression should produce estimates that are insignificant below the policy's threshold, but significant (and negative) above the threshold. In effect, this allows us to test, rather than assert, that the CWS air quality thresholds matter.

To accomplish this we assign each plant-year observation into bins according to their CMA's air quality in that year, and then estimate a version of Equation (3) in which the target industry indicators are interacted with these air quality bins. This amounts to estimating a number of difference-in-difference regressions that, for a given year, compare outcomes for plants in targeted industries to those in non-targeted industries within CMAs with a given range of air quality, and then comparing this to the same difference in an omitted group of CMAs. Every year in the sample is pooled, and the coefficient on each bin is identified from regions changing air quality bins over time.

---

[65] These regression results are available upon request.

[66] Such a relationship could arise if plants select into regions based on unobserved regional characteristics that are correlated with air quality. For example, if the most productive polluters select into clean regions to avoid future regulation, then comparing outcomes in dirty regions to clean regions may simply reflect differential trends between high-productivity and low-productivity plants.

This specification is given by,

$$
\begin{aligned}
Y_{pict} = & \sum_b \beta^b_{PM}[K_i \times I(\underline{A^{PM}_b} \leq a^{PM}_{ct} < \overline{A^{PM}_b})] \\
& + \sum_b \beta^b_{O3}[K_i \times I(\underline{A^{O3}_b} \leq a^{O3}_{ct} < \overline{A^{O3}_b})] + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict},
\end{aligned} \tag{22}
$$

where b indexes air quality bin numbers, K_i selects all industries targeted by the CWS, a^j_{ct} is the air quality measured in CMA c for pollutant j in year t, $\underline{A^j_b}$ is the air quality lower bound for bin b for pollutant j, $\overline{A^j_b}$ is the air quality upper bound for bin b for pollutant j, and $I(\underline{A^j_b} \leq a^j_c < \overline{A^j_b})$ is an indicator for all CMA-years with air quality that corresponds to bin b for pollutant j.[67] The coefficient β^b_j gives the effects of standard j in air quality bin b.

In estimating Equation (22), we omit the "cleanest" air quality bin for each of the standards. For the $\mathrm{PM}_{2.5}$ standard, we break the air quality distribution into seven equal-sized bins from 18 to 36 $\mu g/m^3$. For the O_3 standard, we break the air quality distribution into six equal-sized bins from 57 to 77 ppb [68]

In Figure 6 we plot the results of the estimation of Equation (22) for emissions using the full sample of polluters from the NPRI. Figure 6 shows the coefficients and confidence intervals for $\mathrm{PM}_{2.5}$ and $\mathrm{NO_X}$ emissions, respectively. Only the coefficients for the $\mathrm{PM}_{2.5}$ standard are shown for $\mathrm{PM}_{2.5}$ emissions, and the O_3 standard for $\mathrm{NO_X}$ emissions. Each figure also displays the fraction of observations in each bin treated over the sample, to show that there are plants eventually treated over the entire distribution of air quality. The dependent variable in each regression is the natural log of plant emissions and standard errors are clustered at the CMA-industry level.

The results show a break that occurs just below the $\mathrm{PM}_{2.5}$ standard's threshold for $\mathrm{PM}_{2.5}$ emissions and at the precise level of the O_3 standard's threshold for $\mathrm{NO_X}$ emissions. This means there is no significant difference in trends between treated and control plants until a CMA's air quality reaches that of the standard's threshold. The observed effect of the CWS appears to be coming from a break in trend for the plants in CMA-years above the standard's thresholds. As these thresholds were not used for any other policy, this suggests the results in Table 3 reflect the effects of increased regulation driven by violation of the CWS thresholds, rather than some other relationship between a CMA's air quality and the emissions of manufacturing plants therein.

As our next robustness check, we adopt a common approach in program evaluation and perform an event-study analysis in which the effect of treatment is allowed to vary over time. This type of robustness check is useful for two reasons. First, it allows us to test whether there is a significant difference in outcomes between our treatment and control groups before

---

[67] For example, suppose $\mathrm{PM}_{2.5}$ air quality ranged from 20 to 40 μg/m^3, and we split this into two equal-sized bins. The upper and lower bounds for bin one would be $\overline{A^{PM}_1} = 30$ and $\underline{A^{PM}_1} = 20$, respectively. The upper and lower bounds for bin two would be $\overline{A^{PM}_1} = 40$ and $\underline{A^{PM}_1} = 30$, respectively. Bin one would select all plants in CMAs with air quality below 30 $\mu g/m^3$, and bin two would select all plants in CMAs with air quality above 30 μg/m^3.

[68] For the $\mathrm{PM}_{2.5}$ regulation we include all CMA-years with air quality above 36 $\mu g/m^3$ in the top bin. For the O_3 regulation we include all CMA-years with air quality above 77 ppb in the top bin.

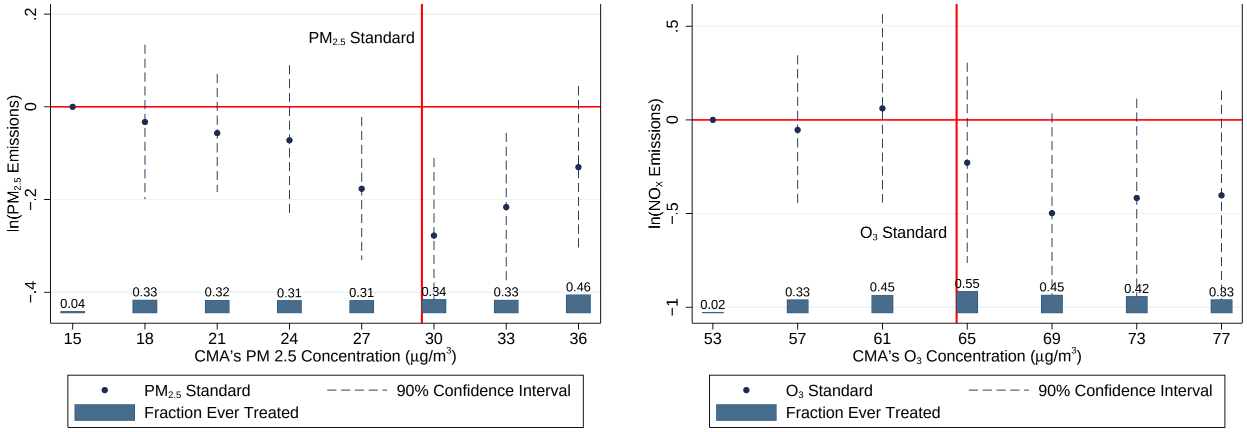

(a) $PM_{2.5}$ Emissions (b) O_3 Emissions

Figure 6: Mean Pollution Concentrations by Year

Notes: Figure shows the results of flexible DDD estimation of the $PM_{2.5}$ standard's effect on $PM_{2.5}$ emissions and the O_3 standard's effects on NO_X emissions allowing the treatment effects to vary by CMA air quality. Diamonds show the triple-difference estimation coefficients by CMA air quality bins, with a 90% confidence interval. Coefficients in each bin are relative to the excluded group (air quality below 18 $\mu g/m^3$ for $PM_{2.5}$ and below 57 ppb for O_3). Standard errors are clustered by industry-CMA. The histogram shows the fraction of observations in each bin treated by the respective standard at some point over the sample.

treatment occurs. If we've constructed a valid control group, there should be no significant pre-treatment differences. Secondly, it allows us to determine if the effects of treatment persist into the future.

This is particularly demanding in this setting because the majority of treated CMAs begin the sample period under treatment, particularly for the O_3 standard (see Section B.2 for more details). As a result, we must rely on a relatively small group of treated plants for the event-study analysis and are only able to perform this robustness check for the $PM_{2.5}$ standard.

We implement the event-study approach by determining the first year a plant exceeds the $PM_{2.5}$ standard's threshold, then comparing treated plants to untreated plants in each of the years before a plant is treated and each of the years after a plant is treated (for which they are are still treated). This regression is estimated by fitting the following generalized triple-difference estimator to the data

$$Y_{pict} = \sum_{k=-3} \beta^k_{PM} T^{PM}_{ick} + \beta_{O3} T^{O3}_{ict} + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict}, \tag{23}$$

where T^{PM}_{ick} is an indicator for the years before $(k < 0)$ or after $(k \geq 0)$ a plant is treated for standard j, and T^{O3}_{ict} captures the average effect of the O_3 standard. We exclude the year prior to treatment for the $PM_{2.5}$ standard $(k = -1)$, so the coefficients of interest (β^k_{PM}) report the semi-elasticity of treatment k years before or after treatment relative to the year before treatment. In other words, $\beta^k_{PM}$is the triple-difference coefficients relative to the year

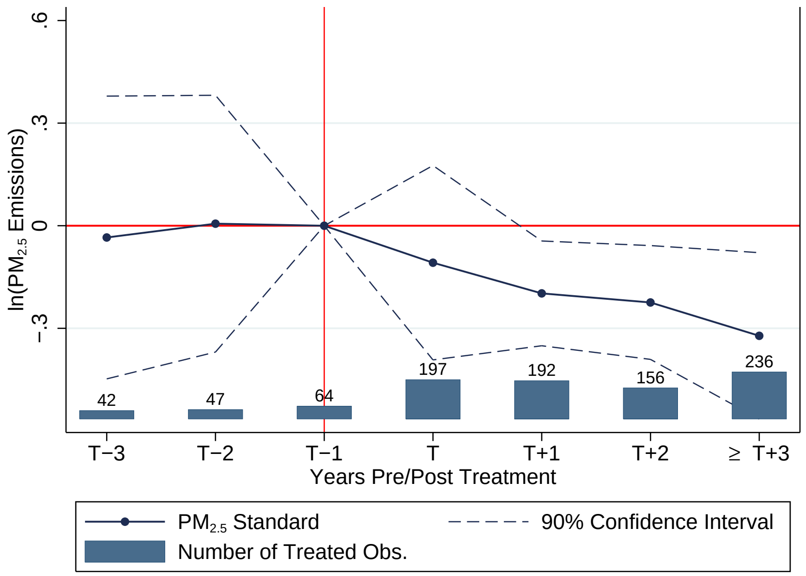

Figure 7: The Effect of $PM_{2.5}$ Regulation on $PM_{2.5}$ by Years Pre/Post Regulation

Notes: Figure shows the results of a flexible DDD estimation of the $PM_{2.5}$ standard for $PM_{2.5}$ emissions allowing the treatment effect to vary by years pre/post regulation. Diamonds show the triple-difference estimation coefficients by years before and after treatment, with a 90% confidence interval in light blue. Treated plants with no pre-treatment data are omitted. All coefficients are relative to the year before treatment (T-1), indicated by a vertical red line. Standard errors are clustered by industry-CMA. The histogram shows the number of observations in each bin treated by the respective standard at some point over the sample.

before a plant is first treated by the standard.[69]

We estimate Equation (23) from three periods before a plant is treated onward. Separate coefficients are estimated up to three periods post treatment, and all periods greater than three years after treatment are pooled. We drop all observations that occur prior to three periods before a plant is treated. All plants in CMAs that began the sample period under treatment are dropped from the regression.

The results of the effects of the $PM_{2.5}$ standard on $PM_{2.5}$ emitters are shown in Figure 7. The dependent variable is the natural logarithm of $PM_{2.5}$ emissions and standard errors are clustered by CMA-industry.

Figure 7 shows strong evidence that there was no significant difference in pre-treatment trends for our treatment and control groups for the PM standard, with the pre-treatment coefficients hovering tightly around zero. In addition, there was a clear break in $PM_{2.5}$ emissions starting in the year of treatment and persisting (and even potentially growing) following treatment.

Treatment Assignment Mechanism

Of particular importance for our identification strategy is that plants may preemptively make production changes so as to avoid regulation. We first show that a large number

---

[69] Note that in our basic specification, Equation (3), the triple-difference coefficient compares the average over all years during which a plant is treated to the average over all years during which a plant is not treated.

Table A2: Regulation Cohorts

	Panel A: % Reg. in 1st Year		Panel B: % Reg. by 2005	
	(1) $PM_{2.5}$	(2) NO_X	(4) $PM_{2.5}$	(5) NO_X
$PM_{2.5}$ Standard	50%	52%	84%	80%
O_3 Standard	56%	68%	63%	87%

Notes: Table reports the regulation cohorts for each standard and group of emitters. Panel A shows the percentage of treated plants treated in the first year of the sample. Panel B shows the percentage of treated plants treated by 2005. The first column within each panel shows the results for $PM_{2.5}$ emitting plants, the second column for NO_X plants. Each cell shows the fraction of plants that are ever regulated by each standard by the year in question. The first row reports results for the $PM_{2.5}$ standard and the second for the O_3 standard.

of treated plants are in CMAs that begin the regulatory period in exceedance of a CWS threshold and eventually "clean-up" and fall below the threshold. Moreover, restricting our treated group to only contain these early regulated plants leaves the core results unaffected. This is important, because as long as the plants in these CMAs were unaware of the policy before it was announced then treatment should be as good as random for this group.

Table A2 shows the fraction of treated plants that are treated early in the policy. Panel A shows the plants treated in the first year of the sample, and Panel B shows the plants treated by the middle of the CWS phase-in. For each standard and pollutant, over half of the treated plants start the sample treated. That fraction increases to between 80% and 90% by 2005 for all standard-pollutant pairs with the exception of the $PM_{2.5}$ emitters treated by the O_3 standard, for which two-thirds are treated by 2005.

Restricting treatment to plants that start the sample treated (dropping all plants treated later from the sample) leaves the results qualitatively unchanged, and actually increases the magnitude of the main effects (though not significantly). The results for the average effect of the CWS on emissions of each pollutant are shown in Table A3. For this group, the $PM_{2.5}$ standard reduced emissions of $PM_{2.5}$ by 17% and PM_{10} by 22%, and the O_3 standard reduced emissions of NO_X by 56%. The average effect of the CWS on scale, and the effects on emissions and scale by plant productivity levels have the same sign and are similar magnitude to the main results.

The above results suggest that preemptive changes by regulated plants to avoid regulation aren't likely to be a problem. Nevertheless, an identification problem could still arise if our effects are primarily driven by large emitters for whom changes in emissions directly affect CMA air quality. This could be problematic for two reasons. Firstly, it would mean influential plants could have potentially manipulated the length of time they were treated, meaning treatment is not exogenous. Secondly, our results could be spurious if large emitters are on a different trend relative to small emitters owing to some other factors beyond regulation, and treatment is positively correlated with large emitter status.

Fortunately, we can test for both of the above concerns. To address the first we employ a blunt-force robustness approach, dropping plants that emit a large fraction of their CMA's emissions. Dropping large plants lowers the potential for bias by removing plants who

Table A3: CWS Effect on Emissions for Initial Treatment Cohort

	(1) $PM_{2.5}$	(2) NO_X
$PM_{2.5}$ Standard	-0.169*	0.0132
	(0.087)	(0.072)
O_3 Standard	-0.059	-0.560*
	(0.082)	(0.330)
R^2	0.268	0.336
N	6538	2881

Notes: Table reports estimates of the effects of the CWS on plant pollution emissions for the cohort of plants treated at the beginning of the sample. All plants treated after the beginning of the sample are dropped. Each panel reports results for a different sample of emitters. In each regression, the dependent variable is the natural log of pollution emissions. The first row reports the effects of the $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

are potential drivers of their city's air quality problem. As there is no obvious size cut-off above which a plant becomes "influential", we start by dropping plants that account for more than 20% of their CMA's emissions and continue tightening until we reach a 1% threshold.[70] We report the results for emissions in Table A4. The effect of the $PM_{2.5}$ standard is remarkably robust. For $PM_{2.5}$ emitters, the effect is negative and statistically significant in each specification, and there is no significant difference between each of the results in Table A4 and the effect in the full sample. The effects of the O_3 standard are also consistent with the main results in the paper, although they are less robust than the PM standard. The O_3 standard is only significant in the first specification for the NO_X emitters; however, the results are qualitatively unchanged and there is no significant difference between the first three specifications and the effects in the full sample. The O_3 regulation's effect on NO_X emissions, however, disappears if we drop plants that emit more than 1% of their CMA's emissions.

The additional core results of the paper are also robust to dropping large emitters. For $PM_{2.5}$ emitters, the average effects on output and by plant productivity-levels for emissions and output are qualitatively unchanged in each of the size thresholds employed in Table A4. The same is true of the O_3 standard's effects for NO_X emitters, with the exception of the most stringent size threshold. As in Table A4, dropping NO_X emitters that account for more than 1% of their city's emissions causes the effect of the O_3 standard to disappear. The O_3 standard's effects appear to be largely driven by plants that emit between 1% and 5% of their city's emissions.

To address the second concern described above, we estimate a version of our main spec-

---

[70] For reference, the average plant fraction of city emissions is: 7% for $PM_{2.5}$ and 10% for NO_X.

Table A4: CWS Effect on Emissions Dropping Large Emitters

	Drop 20%	Drop 10%	Drop 5%	Drop 1%
	Panel A: $PM_{2.5}$			
	(1)	(2)	(3)	(4)
$PM_{2.5}$ Standard	-0.164**	-0.205***	-0.203***	-0.133*
	(0.0651)	(0.0698)	(0.0750)	(0.0749)
R^2	0.220	0.217	0.215	0.246
N	6342	5905	5399	4052
	Panel B: NO_X			
	(1)	(2)	(3)	(4)
O3 Standard	-0.273**	-0.205	-0.219	-0.0696
	(0.115)	(0.129)	(0.134)	(0.133)
R^2	0.334	0.345	0.357	0.468
N	2433	2192	1978	1341

Notes: Table reports estimates of the effects of the CWS on plant pollution emissions dropping large emitters. Each panel reports results for a different sample of emitters. In each regression, the dependent variable is the natural log of pollution emissions. Column one drops all plant-years that account for more than 20% of their CMA's emissions. Column two drops all plant-years that account for more than 10% of their CMA's emissions. Column three drops all plant-years that account for more than 5% of their CMA's emissions. Column four drops all plant-years that account for more than 1% of their CMA's emissions. The first row reports the effects of the $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. The effect of the $PM_{2.5}$ standard is shown for PM emitters, and the O_3 standard is shown for O_3 NO_X emitters. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

ification that allows for separate CMA-year fixed effects for relatively dirty and relatively clean emitters. We accomplish this by determining the fraction of their CMA's annual emissions each plant accounts for, then placing each plant into one of three bins reflecting small, medium, and large emitters. Small emitters produce less than 1% of their CMA's emissions (for the respective pollutant). Medium emitters produce between 1-20% of their CMA's emissions. Large emitters produce more than 20% of their CMA's emissions. We then include a full set of emitter size-by-CMA-by-year fixed effects in our regressions. We are able to do this because, while targeted industries are those that are relatively dirty, how dirty they are relative to other industries varies across the country. In some regions, plants in non-targeted industries are larger emitters than plants in targeted industries, which gives us variation in treatment that is not perfectly correlated with how dirty a plant is relative to other plants in their region.

The results are presented in Table A5. Flexibly controlling for emitter size-by-CMA fixed effects produces similar results to our baseline specification, albeit with a minor attenuation in the main treatment effects. $PM_{2.5}$ regulation significantly reduced PM emissions from affected plants, and O_3 regulation significantly reduced NO_X emissions from affected plants. Consequently, we conclude our results are unlikely to be reflective of differential trends across

Table A5: CWS Effect on Emissions with Large Emitter Trends

	(1) $PM_{2.5}$	(2) NO_X
$PM_{2.5}$ Standard	-0.128**	0.0573
	(0.0590)	(0.0841)
O_3 Standard	-0.0644	-0.277**
	(0.0776)	(0.134)
R^2	0.563	0.652
N	6296	2243

Notes: Table reports estimates of the effects of the CWS on plant pollution emissions controlling for separate trends within each CMA for small, medium, and large emitters. Small emitters are those that account for less than 1% of their CMA's pollution for a given pollutant. Medium emitters emit between 1-20%, and large emitters are those that emit above 20%. Each panel reports results for a different sample of emitters. In each regression, the dependent variable is the natural log of pollution emissions. The first row reports the effects of the $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. All regressions include plant, industry-year and emitter size-by-CMA-by-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

large and small emitters.

Multi-Plant Firms

Approximately 50% of the plants in our sample are directly owned by a firm that owns at least two plants in the manufacturing sector. These multi-plant firms create a potential identification problem because the treatment of one plant may alter the potential outcomes of another plant owned by the same firm. We address this here by identifying the plants owned by these multi-plant firms, and then testing whether the treatment effects differ for plants owned by mutli- and single-plant firms.[71,72]

We use the parent company name information reported in the NPRI to identify multi-plant firms. This information is entered as a text string, which is messy and imprecise. To improve our matching, we use a string-similarity algorithm called the Levenshtein Edit Distance. The Levenshtein Distance measure, in essence, tracks the number of changes required to convert one string to another. Two strings requiring few changes would have a relatively small distance.[73] We classify firms in two ways. In our first approach we classify firms as multi-plant if they own more than one plant that emit the same pollutant (either $PM_{2.5}$ or NO_X). In our second approach we classify firms as multi-plant if they own more than one plant in our dataset (that is, that emit any of 300 pollutants tracked in the NPRI).

---

[71] An alternative approach is to simply drop all multi-plant firms. Doing this produces similar results.

[72] Our data only allows us to identify the immediate parent of a plant, rather than the ultimate corporate parent. As such, our definition of a multi-plant firm is a firm that is the immediate parent of more than one plants, rather than the parent of another firm that owns another plant.

[73] For details on the Levenshtein Distance measure, see Yujian and Bo (2007).

In both approaches we present results using both coarse matching, which produces more matches but is open to more false positives, and fine matching, which is more conservative but more likely to miss correct matches.

We estimate a version of our main specification in which we include a time-varying indicator that selects all plant-years owned by a multi-plant firm, and an interaction between the multi-plant indicator and our treatment indicators. For $PM_{2.5}$ emitters we estimate the $PM_{2.5}$ standard's effect on plant emissions, and for NO_X emitters we estimate the O_3 standard's effect on plant emissions. The results are shown inTable A6. In all specifications there is no significant difference in the treatment effects between the plants owned by single-plant firms and those owned by multi-plant firms. The potential failure of SUTVA through the common-ownership channel does not appear to be an issue in this context.

Additional Robustness

Recall that an identification problem exists if there is an unobservable characteristic that varies by CMA-industry-year and is correlated with treatment under the CWS. We believe there are two potential identification problems of concern that our robustness checks may not have addressed. The first is to do with differential trade shocks. If plants in targeted industries are more likely to export, and the CMAs that eventually exceed the CWS are more connected with Canada's major trading partners (i.e. the US), then exchange rate fluctuations would have a larger effect on the treated plants than the untreated plants. This is a potential issue because over the CWS phase-in period the Canadian dollar appreciated significantly with respect to the US dollar (in 2000, one Canadian dollar was worth 67 cents US, but by 2010 one Canadian dollar was worth 97 cents US). This appreciation in the Canadian dollar made Canadian goods more expensive, which could have depressed relatively export-intensive manufacturing plants. Note, however, that this is only an identification problem if the treated plants (those in the targeted industries in dirty regions) are more trade-exposed than the untreated plants. As we have plant-level data on exports to the US, we can test whether this is true. Testing for differences in trade exposure between our treated and untreated groups, we find plants that are eventually treated by the CWS are less export-intensive than those that are untreated.[74] As a result, differential trade shocks should be less costly to the treated plants, and would bias our results upwards, if at all.

The second remaining potential identification problem is to do with local industrial policy. If local authorities enact policy to protect regulated plants, then this will create industry-by-CMA-by-year variation that is correlated with CWS assignment. However, the goal of these policies would presumably be to support regulated plants, thereby biasing our results upward. As this type of local industrial policy would lead to attenuation bias, it is not a major concern.

B.3 Additional Results

Air Quality Improvements and the CWS

While the variation depicted in Figure 5 is suggestive of substantial changes to environmental regulation, these changes may not be binding if regulations are not enforced. Given that we do not observe regulatory enforcement decisions, we provide descriptive evidence that ambient air quality changes in different CMAs are consistent with changes in environmental

---

[74] These results are available on request

Table A6: CWS Effect on Emissions - Multi-Plant Firms

	Panel A: $PM_{2.5}$			
	Same Pollutant		Any Pollutant	
	Coarse Matching (1)	Fine Matching (2)	Coarse Matching (3)	Fine Matching (4)
$PM_{2.5}$ Std.	-0.200**	-0.197**	-0.236**	-0.230**
	(0.0937)	(0.0949)	(0.0996)	(0.0998)
$PM_{2.5}$ Std. x Multi-Plant	0.0645	0.0594	0.116	0.107
	(0.0997)	(0.103)	(0.0960)	(0.0971)
Multi-Plant	0.0690	0.0634	-0.00161	0.00190
	(0.0767)	(0.0767)	(0.0825)	(0.0813)
R^2	0.938	0.938	0.938	0.938
N	7058	7058	7058	7058
	Panel B: NO_X			
	Same Pollutant		Any Pollutant	
	Coarse Matching (5)	Fine Matching (6)	Coarse Matching (7)	Fine Matching (8)
O_3 Std.	-0.434**	-0.421**	-0.386**	-0.376**
	(0.181)	(0.182)	(0.187)	(0.187)
O_3 Std. x Multi-Plant	0.132	0.112	0.0206	-0.00447
	(0.0875)	(0.0875)	(0.101)	(0.0990)
Multi-Plant	0.0823	0.0840	0.129*	0.135*
	(0.0666)	(0.0657)	(0.0744)	(0.0725)
R^2	0.978	0.978	0.978	0.978
N	2779	2779	2779	2779

Notes: Table reports estimates of the effects of the CWS on plant pollution emissions allowing treatment to vary by the number of plants owned by the plant's parent firm.Each panel reports results for a different sample of emitters. In each regression, the dependent variable is the natural log of pollution emissions. Column one drops all plant-years that account for more than 20% of their CMA's emissions. Column two drops all plant-years that account for more than 10% of their CMA's emissions. Column three drops all plant-years that account for more than 5% of their CMA's emissions. Column four drops all plant-years that account for more than 1% of their CMA's emissions. The first row reports the effects of the $PM_{2.5}$ standard, and the second row reports the effects of the O_3 standard. The effect of the $PM_{2.5}$ standard is shown for PM emitters, and the O_3 standard is shown for O_3 NO_X emitters. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

regulation under the CWS.

To do this, for each standard we sort each of the CMAs into one of two groups: "clean" CMAs that never violated the relevant standard, and "dirty" CMAs that violated the standard at least once over the phase-in period. Doing this allows us to assess whether the changes in air quality across Canada matched with the design of the CWS. Over the period 2000 to 2011, amongst CMAs that never exceeded the $PM_{2.5}$ standard there was no significant change in mean $PM_{2.5}$ concentrations, nor were there significant changes in mean O_3 concentrations amongst CMAs that never exceeded the O_3 standard. Mean $PM_{2.5}$ concentrations in the clean CMAs was approximately 15 $\mu g/m^3$ in each year; mean O_3 concentrations were between 55 and 58 ppb in each year. In contrast, mean $PM_{2.5}$ in the dirty CMAs fell from approximately 30 $\mu g/m^3$ in the beginning of the decade to approximately 22 $\mu g/m^3$ at the end. Similarly, mean O_3 in the dirty CMAs fell from approximately 80 ppb in the beginning of the decade to approximately 68 ppb at the end of the phase-in.[75]

As a further check, we also examine the distributions of CMA air quality for the first half (2000-2005) and second half (2006-2011) of the decade. We again separate CMAs into two groups: (i) clean CMAs where ambient pollution concentrations were never above the CWS, and (ii) dirty CMAs that exceeded the CWS at least once. We estimate these distributions using kernel density estimation with a Gaussian kernel.

These distributions are depicted in Figure 8. As the figure shows, there was almost no change in either of the $PM_{2.5}$ or O_3 distributions for clean CMAs over the entire phase-in period. The same, however, cannot be said for the dirty CMAs. The $PM_{2.5}$ distribution shifted drastically from the beginning to the end of the phase-in, with almost all of the CMA-year observations lying below the CWS threshold in the second half of the phase-in. The right tail of the O_3 distribution shifted leftward, and the mass of CMA-years near the CWS threshold increased substantially. By the end of the phase-in period, most CMAs in Canada had met the $PM_{2.5}$ standard, and the dirty O_3 cities were moving towards compliance. This provides further evidence of changes in air quality consistent with the CWS.

A potential concern with these figures is that they could merely reflect different trends across regions owing to other factors beyond the CWS. A primary concern is that the CMAs exceeding one of these thresholds may be more heavily populated or industrialized than those below the threshold. To show this is not the driver of the documented change in air quality, in Figure 9 we show the change in mean $PM_{2.5}$ (panel (a)) and O_3 (panel (b)) concentrations for CMAs with a population of at least 300,000 people. As above, the panel on the left displays pollution concentrations of CMAs that never exceed the relevant standard, and the right shows concentrations of the dirty CMAs. The figure shows a pronounced drop in both $PM_{2.5}$ and O_3 air quality for the heavily populated regions that exceed the respective thresholds, and no change in air quality for the clean CMAs.

While there are potentially other factors that may explain the changes in air quality shown above, a proper treatment of the effect of the CWS on air quality is beyond the scope of this paper.

The CWS' Effect on Other Pollutants

In this section we present our estimates of the effect of the CWS on plant-level emissions of pollutants not directly regulated by the CWS. These results are useful for two reasons.

---

[75] Both changes are statistically significant at the 95% confidence level.

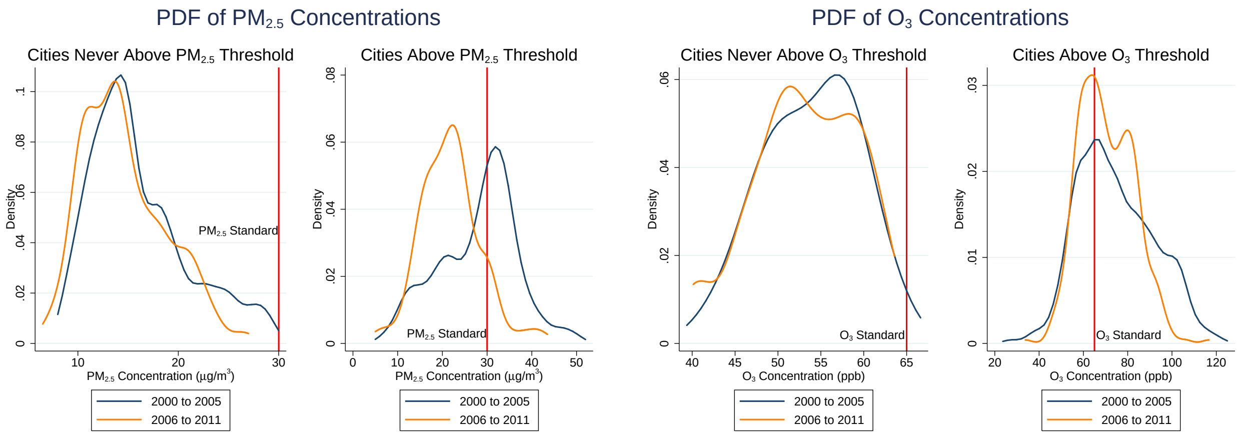

(a) $PM_{2.5}$ Concentrations (b) O_3 Concentrations

Figure 8: Distribution of Pollution Concentrations by Year

Notes: Figure depicts kernel density estimates of the distributions of $PM_{2.5}$ (panel (a)) and O_3 (panel (b)) concentrations across CMAs in the first half (2000-2005) and second half (2006-2011) of the CWS phase in period. For each pollutant, the panel on the left displays pollution concentrations of CMAs that never exceed the relevant standard. The right panel displays the pollution concentrations of CMAs that exceed the relevant standard at least once. The vertical red lines represents the relevant threshold for the each standard.

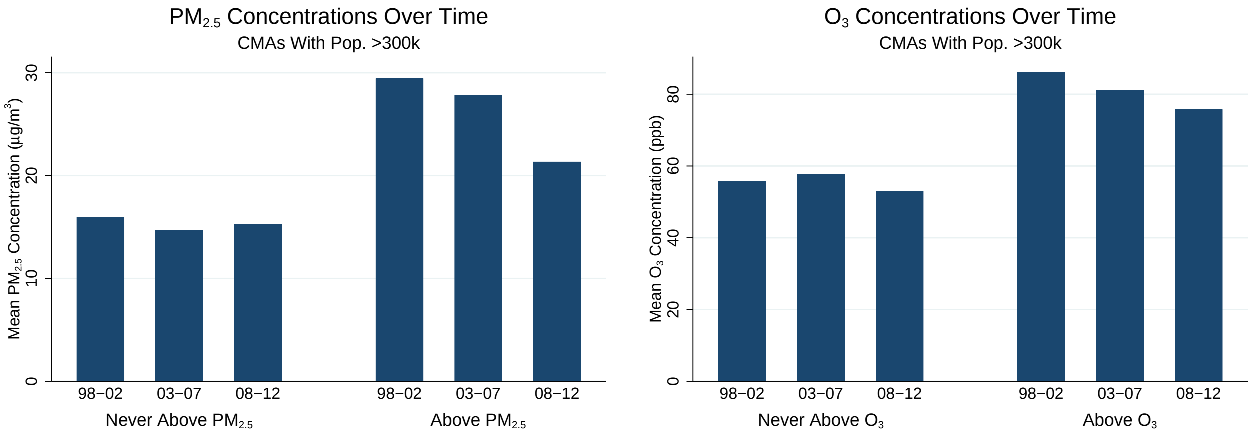

(a) $PM_{2.5}$ Concentrations (b) O_3 Concentrations

Figure 9: Mean Pollution Concentrations by Year for Large CMAs

Notes: Figure depicts mean $PM_{2.5}$ (panel (a)) and O_3 (panel (b)) concentrations by year for CMAs with a population of at least 300,000 people. For each pollutant, the panel on the left displays pollution concentrations of CMAs that never violated the relevant standard. The right panel displays the pollution concentrations of CMAs that violated the relevant standard at least once.

Firstly, ambient $PM_{2.5}$ and O_3 pollution may be formed through chemical reactions in the atmosphere between other pollutants besides $PM_{2.5}$ and NO_X, in particular other criteria air contaminants (CACs). Secondly, this allows us to assess whether there were positive or negative spillovers in response to the CWS. A positive spillover would occur if plants substitute toward unregulated pollutants, whereas a negative spillover would occur if emissions were were correlated across pollutants. The former is typically referred to as regulation-induced substitution, and the latter as co-pollutant effects.

We consider emissions of other important air pollutants collected in the NPRI, including other CACs and heavy metals, as well as greenhouse gas (GHG) emissions. For CACs, we consider emissions of large-scale particulate matter (PM_{10}), volatile organic compounds (VOCs), sulphur dioxide (SO_2), and carbon monoxide (CO). For heavy metals, we consider lead and zinc.

GHG information is not available in the NPRI, however, Environment and Climate Change Canada collect GHG emission data for the largest plants in the country, and is publicly available through the Greenhouse Gas Reporting Program (GHGRP).[76] The GHGRP reports emissions for several GHGs (including carbon dioxide, methane, and nitrogen dioxide), total facility GHGs, and provides a crosswalk file to match plants in the GHGRP with plants in the NPRI. We use this crosswalk file to merge the facility GHG data to the NPRI. Virtually all manufacturing plants in the GHGRP over our sample were successfully matched to the NPRI. We report the effect of the CWS on total GHG emissions, carbon dioxide emissions, and nitrogen dioxide emissions.

The results of these regression are shown in Table A7. Each column reports our estimates of the effect of the CWS on a different pollutant. The dependent variable in each of these regressions is the natural log of plant pollution emissions for the relevant pollutant, and standard errors clustered at the CMA-industry level are reported in parentheses.

The results in Table A7 show the $PM_{2.5}$ standard caused a significant drop in PM_{10} emissions. This, to some extent, is a mechanical result: by definition, reported PM_{10} emissions *include* emissions of $PM_{2.5}$. Nonetheless, these results provide added confidence to our main results, and indicate there was no significant substitution from fine to large scale particulate matter emissions in response to the CWS. $PM_{2.5}$ regulation had only minor effects on emissions of the other air pollutants and greenhouse gases. $PM_{2.5}$ regulation caused a small (insignificant) increase in SO_2, CO, and lead, a small (insignificant) drop in VOCs, and had virtually no effect on GHGs. The only pollutant showing a sizeable response to PM regulation is zinc, which fell by 23%, although this is not significant at conventional levels.

The O_3 standard, however, had no effect on heavy metals or PM_{10}, but caused a large reduction in emissions of other CACs and greenhouse gases. O_3 regulation caused a 37% reduction in total GHG emissions, a 21% reduction in VOCs (which is a potential ozone precursor), a 44% reduction in CO emissions (which is also a potential ozone precursor), and a 51% reduction in SO_2 emissions. The drop in GHGs was driven by reductions in both CO_2 (32%) and N_2O (66%), the latter being both a GHG and an ozone precursor.

---

[76] See Environment and Climate Change Canada (2016b) for data.

Table A7: The Effects of the CWS on Plant Emissions of Unregulated Pollutants

	Panel A: CACs				Panel B: Metals		Panel C: GHGs		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	PM_{10}	VOCs	SO_2	CO	Lead	Zinc	GHGs	CO_2	$N_2 0$
$PM_{2.5}$ Std.	-0.17***	-0.07	0.05	0.07	0.05	-0.23	-0.02	0.02	0.04
	(0.06)	(0.08)	(0.21)	(0.09)	(0.32)	(0.18)	(0.19)	(0.19)	(0.15)
O_3 Std.	0.01	-0.22*	-0.51**	-0.44*	-0.09	0.01	-0.37**	-0.32**	-0.66***
	(0.09)	(0.12)	(0.21)	(0.24)	(0.56)	(0.79)	(0.17)	(0.15)	(0.19)
R^2	0.22	0.22	0.45	0.38	0.37	0.37	0.61	0.62	0.68
N	8,003	7,045	2243	3352	1411	1496	701	701	613

Notes: Table reports estimates of the effects of the CWS on plant emissions of pollutants not directly regulated by the CWS. Each column reports estimates from a regression of the CWS regulations on the natural log of the emissions a different pollutant. Panel A shows the effects on other criteria air contaminants (large scale particulate matter, volatile organic compounds, sulphur dioxide, and carbon monoxide). Panel B shows the effects on heavy metals (lead and zinc). Panel C shows the effects on greenhouse gases (total emissions, carbon dioxide, and nitrous oxide). In all cases, the first row reports the effects of $PM_{2.5}$ regulations, and the second row reports the effects of the O_3 regulations. All regressions include plant, industry-year, and CMA-year fixed effects. Standard errors clustered by CMA-industry are reported in parentheses. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels.

Table A8: Cross-Pollutant Elasticities

	(1) PM_{10}	(2) VOCs	(3) SO_2	(4) CO	(5) Total Metals	(6) GHGs
PM 2.5	0.821***	0.267***	0.211***	0.0568	0.339***	0.121***
	(0.013)	(0.032)	(0.0485)	(0.0356)	(0.0851)	(0.0342)
R^2	0.708	0.086	0.040	0.004	0.043	0.118
N	2,613	1,207	739	1109	584	199
	(1) PM_{10}	(2) VOCs	(3) SO_2	(4) CO	(5) Total Metals	(6) GHGs
NOx	0.561***	0.533***	0.766***	0.700***	0.403	0.333***
	(0.052)	(0.057)	(0.0658)	(0.0494)	(0.248)	(0.0591)
R^2	0.156	0.143	0.234	0.246	0.012	0.245
N	1,035	869	737	1008	341	209

Notes: Table reports cross-pollutant elasticities between regulated pollutants (either $PM_{2.5}$ or NO_X) and unregulated pollutants. The estimates are computed by regressing the natural log of plant emissions for each unregulated pollutant on the natural log of plant emissions for each regulated pollutant, including a plant fixed effect. Only early years are use (before 2006). The top panel shows the elasticities for $PM_{2.5}$ emissions; the bottom panel the elasticities for NO_X emissions.

There is a clear explanation for the observed effects on these other pollutants: co-pollutant effects. The correlation between changes in $PM_{2.5}$ emissions and PM_{10} emissions is very high, but very low for other pollutants. Whereas the correlation between NO_X emissions and emissions of other CACs and GHGs is relatively high. We show this by estimating simple co-pollutant elasticities between the regulated pollutants ($PM_{2.5}$ and NO_X) and unregulated pollutants. Our approach is to estimate the within-plant cross-pollutant elasticity for each regulated-unregulated pollutant pair, by estimating the following equation

$$ln(z_{u,i,t}) = \alpha_{u,r} ln(z_{r,i,t}) + \lambda_{u,i} + \epsilon_{u,i,t},$$

where u indexes an unregulated pollutant included in Table A7, r indexes either $PM_{2.5}$ or NO_X, $\lambda_{u,i}$ is a pollutant-plant fixed-effect, and $\alpha_{u,r}$ is our estimate of the cross-pollutant elasticity between unregulated pollutant u and regulated pollutant r.

We restrict our sample to years before 2006 to try to limit the potential interference of the CWS in changing these cross pollutant elasticities (recall, most of the CWS regulations were implemented between 2005 and 2007). The results from these regressions are shown in Table A8.

The findings show an intuitive result: the cross pollutant elasticity between $PM_{2.5}$ and PM_{10} is over 80%, whereas the cross pollutant elasticities between $PM_{2.5}$ and other pollutant emissions are relatively low (between 0.05 and 0.33). In contrast, the cross-pollutant elasticities are much higher for NO_X emissions (between 0.33 and 0.77). The one outlier is that we find a relatively high correlation between PM_{10} and NO_X, despite finding no significant effect of O_3 regulation on PM_{10} emissions.

B.4 CWS Counterfactuals

First, we present details on the plant-level decomposition. Recall that the change in an industry's pollution intensity is given by

$$\Delta E_{it} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn - \int_0^{n_{it}} e_{it-1}(n)\lambda_{it-1}(n)dn - \int_{n_{it}}^{n_{it-1}} e_{it-1}(n)\lambda_{it-1}(n)dn.$$

This can be written as

$$\begin{aligned} \Delta E_{it} = &\int_0^{n_{it}} (\lambda_{it}(n) - \lambda_{it-1}(n))e_{it}(n)dn - \int_0^{n_{it}} \lambda_{it}(n)e_{it-1}(n)dn \\ &+ \int_0^{n_{it}} (e_{it}(n) - e_{it-1}(n))\lambda_{it-1}(n)dn + \int_0^{n_{it}} \lambda_{it}(n)e_{it}(n)dn \\ -&\int_0^{n_{it}} (e_{it}(n) - e_{it-1}(n))\lambda_{it-1}(n)dn - \int_{n_{it}}^{n_{it-1}} e_{it-1}(n)\lambda_{it-1}(n)dn. \end{aligned}$$

With some algebra, this reduces to

$$\begin{aligned} \Delta E_{it} = &\int_0^{n_{it}} e_{it-1}(n)\Delta\lambda_{it}(n)dn + \int_0^{n_{it}} \lambda_{it-1}(n)\Delta e_{it}(n)dn \\ &+ \int_0^{n_{it}} \Delta\lambda_{it}(n)\Delta e_{it}(n)dn - \int_{n_{it}}^{n_{it-1}} e_{it-1}(n)\lambda_{it-1}(n)dn. \end{aligned}$$

Dividing by E_{it-1} gives the desired decomposition.

To express $\hat{\lambda}_{it}(n)$ as a function of our estimates, note that

$$\begin{aligned} \hat{\lambda}_{it}(n) &= \frac{\lambda_{ft}(n)}{\lambda_{ft-1}(n)} - 1 \\ &= \frac{x_{ft}(n)}{x_{ft-1}(n)}\frac{X_{it-1}}{X_{it}} - 1. \end{aligned}$$

By assumption, if n is untreated, then $x_{ft}(n) = x_{ft-1}(n)$, and if n is treated, then $x_{ft}(n) = (1+\beta_x)x_{ft-1}(n)$. Plugging this into $X_{it} = \int_0^{n_{it}} x_{it}(n)dn$ gives

$$\begin{aligned} X_{it} &= (1+\beta_x)\int_{treated} x_{it-1}(n)dn + \int_{untreated} x_{it-1}(n)dn \\ &= X_{it-1} - \int_{n_{it}}^{n_{it-1}} x_{it-1}(n)dn + \beta_x \int_{treated} x_{it-1}(n)dn. \end{aligned}$$

Rearranging gives $\frac{X_{it}}{X_{it-1}} = 1 - s^{Exit}_{xt-1} + \beta_x s^{Treat}_{xt-1}$. With some algebra it can be shown that $\hat{\lambda}_{it}(n)$ is as in the text.

Appendix C Policy Details

Nova Scotia

In 2004, Nova Scotia adopted emissions taxes for particulate matter and ozone-precursor pollutants (nitrogen oxides and volatile organic chemicals) (N.S. Reg. 31/2005). The emissions taxes were tiered such that small emitters were exempt, mid-size emitters paid a flat fee, and large emitters paid a flat fee plus a tax of $2.70/tonne for emissions above a given threshold. In 2005, Nova Scotia also strengthened its Air Quality Regulations, which were first passed in 1995 (N.S. Reg. 28/2005). There were three substantive changes; the provincial sulphur dioxide cap was reduced from 189,000 tonnes to 141,750 tonnes, the sulphur dioxide emissions cap for the electricity generation sector was strengthened, and nitrogen oxide and mercury emission caps were added for the electricity generation sector.

New Brunswick

New Brunswick amended their provincial air quality regulations (regulation 97-133 under the New Brunswick Clean Air Act) in 2005 to increase the emissions fees assessed for particulate matter (and sulphur dioxide) emitters. New Brunswick uses a staggered annual emissions fee schedule, with the highest annual fees being levied against the largest emitters. The 2005 amendment increased these fees by between 30%-900%, depending on the class of emitter. For example, the annual fees for the largest emitters rose from $42,000 to $60,000, for mid-range emitters from $15,000 to $28,00, and for the smallest emitters from zero to $500. For details, see part five of the regulation.

Ontario

In 2005, Ontario adopted site-specific air quality standards (regulation O Reg 419/05). These standards targeted many different pollutants, including ozone, ozone pre-cursors (including nitrogen oxides and various volatile organic compounds), and particulate matter[77]. The regulation contained more stringent standards for a number of industries, including several of the industries targeted by the CWS.[78] In addition to more stringent standards, plants in these industries must submit annual emissions reports to the Ontario environment ministry.

In 2006 Ontario introduced a limited NO_X and SO_2 trading program for the twenty largest emitters in four of the five CWS-targeted industries (regulation O Reg 194/05). While permit trading allowed flexibility in compliance with the policy, permits were allocated based on the pollution intensity of each facility, such that cleaner plants received relatively more permits.

Quebec

Quebec developed the Clean Air Regulations – which included local air quality standards and site-specific emissions standards – during the phase-in period (regulation QLR Q-2, r 4.1). Air quality standards were developed for a large number of pollutants, including ozone and particulate matter (the $PM_{2.5}$ standard was set at the level of the CWS and the ozone standard was set slightly more stringent than the CWS at 62.5 ppb). Emissions standards were developed for many different industries and industrial processes, including particle

---

[77] A standard was set for total particulate matter, but no standard was set for $PM_{2.5}$. The Ontario Ministry of Environment's rationale for omitting a $PM_{2.5}$ standard was to avoid duplicating the existing CWS (point 8 in `http://www.airqualityontario.com/downloads/AmbientAirQualityCriteria.pdf`).

[78] In particular, pulp and paper, electric power generation, iron and steel manufacturing, and base metal smelting.

emissions from a variety of sources (chapter II), VOCs from a variety of sources (chapter IV), pollutants from combustion plants (chapter VI), and pollutants from incinerators (chapter VII). Although the regulations were first published in 2005, it took six years before they were officially made law.

Prince Edward Island

Prince Edward Island amended their Air Quality Regulations in 2004 to add particulate matter emissions fees for fuel-burning equipment (for details, see schedule D of `http://www.gov.pe.ca/law/regulations/pdf/E&09-02.pdf` and `http://www.gov.pe.ca/photos/original/leg_table_regs.pdf`).[79]

Newfoundland

In 2004, Newfoundland amended its Air Pollution Control Regulations, which had been in place since 1996 (NLR 39/04). The original regulations contained air quality standards for $PM_{2.5}$ that were more stringent than the CWS and a one-hour ozone standard. The amendments left the $PM_{2.5}$ standard unchanged and added an eight-hour ozone standard of 43.5 ppb (more stringent than the CWS). The Newfoundland standards allow the province's minister of the environment to regulate individual facilities should regional air quality exceed one of the standards (see paragraph 3.(3) of the regulations). The amendments also added NO_x emission intensity standards for all new or modified fossil fuel fired boilers and heaters (paragraph 19). In 2014, the regulations were amended further to add an annual $PM_{2.5}$ standard equal to that under the Canadian Ambient Air Quality Standards, which replaced the CWS (for details, see: `http://www.assembly.nl.ca/legislation/sr/regulations/rc040039.htm#3_`).

Manitoba

Manitoba uses objectives and guidelines to manage air quality, rather than provincial regulations. In 2005, the province added the CWS' ozone and $PM_{2.5}$ standards to this list of objectives (for details, see: `https://www.gov.mb.ca/conservation/envprograms/airquality/pdf/criteria_table_update_july_2005.pdf`).

Saskatchewan

Saskatchewan's Clean Air Regulations have imposed ambient air quality standards in the province since 1989 (see: `http://www.qp.gov.sk.ca/documents/English/Regulations/Repealed/C12-1R1.pdf`). These standards remained in place until they were repealed in 2015 by the Environmental Management and Protection Regulations (regulation E-10.22 REG 2) (source: `http://www.qp.gov.sk.ca/documents/English/Regulations/Regulations/E10-22R2.pdf`). The new regulations imposed more stringent air quality standards, including adopting the Canadian Ambient Air Quality Standards for $PM_{2.5}$ and ozone (see Table 20 of the Saskatchewan Environmental Quality Standard, `https://envonline.gov.sk.ca/Pages/SEQS/Table20-SEQS-SAAQS.pdf`).

Alberta

Alberta primarily manages air quality using ambient air quality objectives and guidelines, that are enforced through the provincial permitting and licensing process. Industrial facilities must be designed and operate so as to ensure the provinces ambient air quality objectives are met; however, they are given relative freedom in deciding how to

---

[79] Two amendments were made: EC161/04 and EC423/04.

manage their pollution. More stringent permitting regulations were passed in 2003 (Alberta Regulation 276/2003), under the Environmental Protection and Enhancement Act. In 2007, the CWS' $PM_{2.5}$ standard was adopted as an objective (for details, see: `http://environment.gov.ab.ca/info/library/5726.pdf`). An ozone objective has been in place since 1975, and was reviewed in 2007 but left unchanged. Firms can be fined for violating the conditions of an operation permit, such as failure to comply with air pollutant-related constraints. For example, in 2012 a refinery was fined for failing to install proper air pollution control equipment (for details, see `https://www.alberta.ca/release.cfm?xID=32232CC295887-C17E-3ABE-EC7823B5948337D0`).

British Columbia

British Columbia manages air quality using a combination of air quality objectives, local airshed management plans, and industrial codes of practice. Air quality objectives are non-binding standards that set the air quality levels to which regulators should aim. Over the phase-in period, the province adopted the $PM_{2.5}$ and O_3 Canada-Wide Standards as air quality objectives (for details, see: `http://www.bcairquality.ca/regulatory/air-objectives-standards.html`). Towards the end of the phase-in, the province adopted additional, more stringent, $PM_{2.5}$ objectives (see: `http://www.bcairquality.ca/regulatory/pm25-objective.html`). Provincial regulators achieve these objectives using mandatory codes of practice or other regulations (for details, see paragraph 4.3.2 of `http://www.bcairquality.ca/reports/pdfs/pm25-implement-guide.pdf`). These provincial regulations can target specific industries, regions, or facilities[80]. In addition, local regulators develop local airshed management plans to meet the air quality objectives. Over the CWS phase-in period, airshed management plans were developed for thirteen regions in the province (see `http://www.bcairquality.ca/airsheds/bc-airsheds.html`).

Subsidies

Over the CWS period, some provinces provided subsidies to encourage plants to adopt the cleaner production techniques suggested in the industry MERS. These subsidies were relatively small, and were intended to offset the costs of developing an abatement plan, but not cover the capital and operating costs involved in abatement. Examples included the Enviroclubs initiative in Quebec (see (Lanoie and Rochon-Fabien, 2012) for details), and the Business Air Quality Program Pilot in Ontario (Environment Canada, 2005).

---

[80] Regulations and codes of practice exist for the pulp and paper, wood product manufacturing, asphalt, and agricultural sectors. For details, see `http://www.env.gov.bc.ca/epd/codes/index.htm`