

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

Nouri Najjar[†] Jevan Cherniwchan[‡]

Abstract

For much of the industrialized world, pollution from manufacturing has been falling despite increased output. We provide the first estimates of how 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. We find that regulation explains, at most, 61% of the clean-up of Canadian manufacturing, but the underlying channels differ across pollutants. We present a stylized model to illustrate how the channels through which regulations contribute to the clean-up depend on the cost of adopting cleaner production processes.

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 participants at various seminars and conferences for helpful comments and suggestions. We would also like to thank Michael Willox and Danny Leung at Statistics Canada for their efforts in making this project possible, and members of Environment and Climate Change Canada’s (ECCC) Legislative and Regulatory Affairs Division for helpful discussions on 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 SSHRC. 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.

[†]Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC, Canada, V6T 2E8. Email: npmnajjar@gmail.com.

[‡]Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta, Canada, T6G 2R6. Email: jevan@ualberta.ca. Tel: 780-492-2944.

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, emissions of most air pollutants from manufacturing 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, 2017). 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 industrialized world, manufacturing is becoming cleaner.

Environmental regulations have played a prominent role in the debate over what is causing this “clean-up” of manufacturing. Yet, while policy appears to have played a large role in the clean-up (Shapiro and Walker, 2015) and the clean-up appears to be due to reductions in the emission intensity of individual industries (Levinson, 2015; Brunel, 2017), at present, little is known about how regulations have affected industry emission intensity. This means the clean-up may not be replicable. If regulation reduced industry emission intensity by causing regulated facilities to exit or shrink instead of adopting cleaner production processes, politically constrained governments may not be willing to adopt the policy changes needed to clean-up their manufacturing sectors.

In this paper, we provide the first estimates of how within-industry responses to environmental regulation have contributed to the clean-up of manufacturing. To do so, we estimate the effects of a major revision to Canadian environmental policy - the Canada Wide Standards for Particulate Matter and Ozone (CWS) - on the production choices, pollution emissions, and entry and exit decisions of affected Canadian manufacturing plants over the period 2004-2010. By definition, these plant-level responses determine the emission intensity of individual industries. As such, we use our estimates to determine how environmental regulations have contributed to the clean-up of manufacturing by causing: (i) the adoption of cleaner production processes (a “process” effect), (ii) reallocations of output from regulated to unregulated plants

(a “reallocation” effect), and (iii) plant entry and exit (a “selection” effect).

Canada is an attractive setting to study how environmental regulations have contributed to the manufacturing clean-up for three reasons. First, the Canadian clean-up has been very similar to those documented in other countries. As we show below using the industry decomposition developed by Levinson (2009, 2015), total manufacturing emissions of most air pollutants in Canada have fallen substantially, primarily because of reductions in the emission intensity of individual industries. Second, the CWS - the major air quality regulation in place in Canada during our period of study - is similar in design to many of the air quality regulations that have been used in other jurisdictions that have experienced a clean-up. Third, by focusing on Canada, we are able to utilize a unique confidential dataset containing longitudinal information on the pollution emissions and productive activities of manufacturing plants. This data contains the detail needed to examine the process, reallocation, and selection effects induced by regulation.

To guide our empirical work, we develop a stylized model in which heterogeneous plants face regulatory constraints similar 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 in the design of the CWS. Under the CWS, regulated plants were required to either adopt technical changes to meet industry best practices, or to reduce activities generating the regulated pollutant. Third, it allows for differences across pollutants in the cost of adopting less-polluting production processes. This feature is important because fine scale particulate matter ($PM_{2.5}$) and nitrogen oxide¹ (NO_X), the two pollutants we study, have very different process costs, at least in Canada: NO_X process changes can be accomplished at a relatively low fixed cost, while $PM_{2.5}$ changes typically requires high fixed costs (Canadian Council of Ministers of

¹Nitrogen oxide is a main contributor to O_3 pollution, and a main target of the CWS.

the Environment, 1998; Environment Canada, 2002).

We then turn to estimate the effects of the CWS on affected manufacturing plants. To do so, we exploit variation in regulatory stringency created by the design of the CWS. The agreement was designed to ensure each region met a minimum level of air quality by establishing thresholds for the ambient concentration of $\text{PM}_{2.5}$ or 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”. As a result, plants in targeted industries and regions violating one of the CWS standards were subject to more stringent environmental regulation.² We identify the effects of regulation on these plants using a triple-difference research design that exploits the variation in regulatory stringency across time, region and industry. This allows us to control for factors such as localized recessions or industry demand shocks that would otherwise confound the effects of environmental regulation.

We estimate the effects of the CWS on the pollution emissions, emission intensity, output, and entry and exit decisions of affected Canadian manufacturing plants. We find robust evidence that the CWS reduced pollution emissions from affected manufacturing plants. For the average $\text{PM}_{2.5}$ emitting plant, the CWS is associated with a 15% reduction in $\text{PM}_{2.5}$ emissions. Furthermore, the CWS is associated with a 33% reduction in NO_x emissions from the average NO_x emitting plant. Our theory predicts that these reductions will be driven by different mechanisms. If the fixed costs of process changes are high, as in the case of $\text{PM}_{2.5}$, only relatively productive plants will adopt cleaner production processes following regulation. As a result, in this case, the CWS should have little to no effect on the emission intensity of the average plant. If the fixed costs of process changes are low, as with NO_x , then even less productive plants should adopt 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

²The annual permits required by plants to operate in each province were used to impose these regulations. We describe the CWS in more detail in Section 2.

have a significant effect on the emission intensity of the average affected $\text{PM}_{2.5}$ emitting plant, but is associated with a 29% reduction in the emission intensity of the average affected NO_X emitting plant. Our estimates of the effects of the CWS on plant output also fit with the predictions of our model; we find that the CWS was associated with a 11% reduction in output from the average affected $\text{PM}_{2.5}$ emitting plant, but had little to no effect on the output of the average NO_X emitter. As predicted, we also find that the CWS was associated with a significant reduction in the number of plants that emit $\text{PM}_{2.5}$, but had no significant 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. To make this contribution explicit, we develop an approach to translate the micro-level effects of the CWS to estimates of the aggregate channels of the clean-up. Our estimates for the responses of $\text{PM}_{2.5}$ emitters suggest that the effects of the CWS explain close to 21% of the reduction in the $\text{PM}_{2.5}$ intensity of manufacturing. 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 $\text{PM}_{2.5}$ clean-up was primarily driven by reallocation and selection effects, whereas the clean-up of NO_X was primarily due to process effects induced by regulation.

Our model suggests these differential responses to regulation are due to differences in fixed process costs across pollutants. While we have focused on this channel given the available evidence documenting the stark differences in the costs of abating $\text{PM}_{2.5}$ and NO_X , these costs are unobserved. Hence, to provide further evidence that our estimates are consistent with this mechanism, we also examine the heterogeneity across plants in response to regulation. Our model suggests the effects of regulation should only vary across plants of different productivity levels if the fixed costs of process changes 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 pollution from relatively low-

productivity regulated $PM_{2.5}$ plants fell primarily due to reductions in output, whereas pollution emissions from the moderately-productive $PM_{2.5}$ plants fell due to a reduction in emission intensity. In contrast, NO_X pollution intensity fell for both mid- and low-productivity plants. These results further suggest that our findings are driven by differences in fixed process costs.³

Altogether, our findings contribute to a burgeoning literature examining the sources of the clean-up of the manufacturing sector. This research stems from the work of Levinson (2009) who examined 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.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. Shapiro and Walker conclude that the effects of regulation explain most of the reduction in pollution emissions from US manufacturing. Our analysis complements this work directly by providing causal estimates of plant responses to environmental regulation, and showing how these regulations have contributed to the clean-up of manufacturing via process, reallocation and selection effects.

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 NO_X emissions from manufacturing plants. In addition, the U.S. Clean Air Act appears to have 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

³We also 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 alternative explanations for why we observe different responses to the CWS across pollutants. As we show in the online appendix, we find little evidence to support these explanations.

⁴Others have argued trade may have caused changes to how plants produce their goods (by, for example, outsourcing some production or adopting new technologies), leading to a reduction in plant emission intensity (see Martin (2012) or Cherniwchan (2017)).

in response to regulation are due to changes in the level of output produced, or changes in the emission intensity of production.⁵

Lastly, our work also relates to a large literature examining the effects of air quality regulation on various aspects of manufacturing plant operations. Our work is most closely related to the papers that have studied the effects of regulation on either plant entry and exit (e.g. Henderson (1996), Becker and Henderson (2000)) or plant output (e.g. Greenstone (2002)), providing preliminary evidence on the importance of selection and reallocation effects. We build on these studies by also documenting the process effects induced by regulation, and showing that the effects of environmental regulation may vary across plants of different productivity levels.⁶

The remainder of this paper proceeds as follows. In Section 1, we document the clean-up of the Canadian manufacturing sector. Section 2 provides a brief overview of the CWS. In Section 3 we outline our theory. Section 4 presents our data, outlines our research design and empirical specification, and presents our empirical results. Finally, Section 5 concludes.

1 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)) and the European Union (e.g. Brunel (2017)), it has yet to be documented 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

⁵Our work is also related to that of Martin et al. (2014) who ask how the energy intensity of UK manufacturing plants were affected by a carbon tax.

⁶Other related work considers regulation's effect on input use and productivity (e.g. Berman and Bui (2001), Greenstone et al. (2012), and Walker (2013)), which are dimensions potentially related to the process effect.

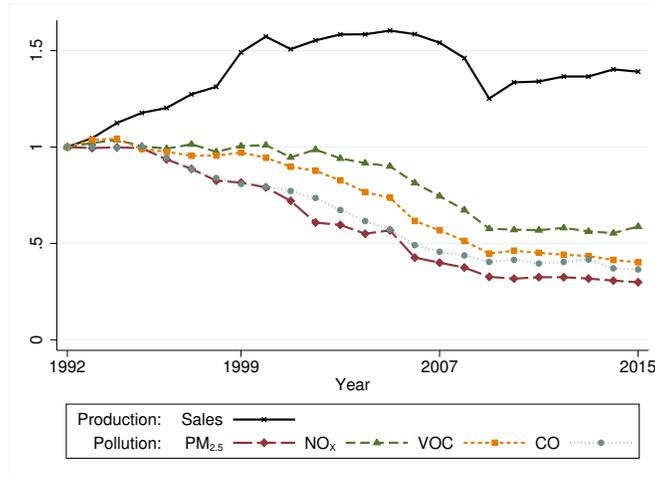


Figure 1: Output and Pollution from Canadian Manufacturing: 1992-2015

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.

the Canadian manufacturing sector, as well as changes in aggregate manufacturing output. As it shows, the emission intensity of Canada's manufacturing sector has fallen since 1992. 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 clean-up of Canadian manufacturing was similar to those 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 (2017) shows the emission intensity of European manufacturing fell by 3.4-5.5% 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. This approach allows us to determine if the observed reductions in aggregate emission inten-

sity are driven by a “composition effect” created by a reallocation of economic activity from dirty emission-intensive sectors to clean sectors with relatively low emission 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 industries⁷, indexed by i . Manufacturing emission 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

$$dE = \sum_i E_i d\theta_i + \sum_i \theta_i dE_i. \quad (1)$$

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 emission intensity attributable to both the composition and technique effects for 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 change in emissions and emission intensity that occurred for manufacturing 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 emission 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-

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

⁸The 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.

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

Pollutant	Δ Emissions (1)	Δ Emission Intensity (2)	Technique Effect (3)	Composition Effect (4)	Technique 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 emission 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)).

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 (2017), the clean-ups of US and European manufacturing are also primarily due to the technique effect.⁹

1.1 How do Industries Clean-Up?

The evidence presented above indicates that reductions in industry emission intensities are the predominant source of the clean-up of manufacturing. This suggests that plant-level changes are a key determinant of the observed reductions in manufacturing emission intensity.

To see this, it is useful to extend the logic of the decomposition 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 $x_i(n)$, $z_i(n)$, and $e_i(n)$ denote output, pollution, and pollution intensity from plant n . Lastly, let

⁹In addition, Shapiro and Walker (2015) perform a product-level decomposition, and find the clean-up in the US is primarily due to within-product reductions in pollution intensity.

$\lambda_i(n) = x_i(n)/X_i$ be plant n 's share of production in industry i and 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 plant emission intensities: $E_i = \int_0^{n_i} e_i(n)\lambda_i(n)dn$. The change in emission intensity of any industry i is then

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

We call the first term on the right-hand side of equation (2) the “process effect”. This captures the change in industry emission intensity due to changes in plant emission intensity resulting from the adoption of new production processes. As such, this term captures the direct effects of a shock; all else equal, industry emission intensity will fall if a shock such as environmental regulation induces plants to lower their emission intensities. The remaining two terms capture indirect changes in industry emission intensity. 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, this may cause a reduction in the relative output of affected plants. This would cause 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). In what follows, we present direct estimates of the regulatory process, reallocation, and selection effects in the clean-up of Canadian manufacturing.

2 Air Quality Regulation in Canada

In order to understand how environmental regulations contributed 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 2000-2012.¹⁰ 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.¹¹ 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).¹² Much like the National Ambient Air Quality Standards at the centre of the U.S. Clean Air Act Amendments (CAAAAs), 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

¹⁰It was 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.

¹¹For details of the CWS, see Canadian Council of Ministers of the Environment (2000a).

¹²The 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, while a CA must have a core population of at least 10,000. We use the term CMA for both.

¹³The standard for particulate matter required each CMA’s 24-hour $PM_{2.5}$ concentration lie below $30\mu g/m^3$. Achievement of the $PM_{2.5}$ standard was based on the 98th 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 4th 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 ppb (Environmental Protection Agency, 2016).

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.

Broadly, the CWS was a tiered regulatory approach in which the federal and provincial 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 regulated plants in targeted industries to meet these regional standards.

To regulate plants, the provinces used their annual operation permit systems. In each province, these systems required plants to prove compliance with certain environmental regulations in order to operate in any year (see, e.g. Environment Canada (2002)). To address the CWS, in most instances facilities could effectively follow one of two paths to meet the permitting requirements: they could either adopt technical changes recommended in their industry’s MERS, or reduce activities contributing to the problematic pollutant. 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 its region’s air quality.

The options available to plants for meeting the regulations imposed under the CWS also differed across pollutants due to technical constraints. For example, NO_x , the main ozone precursor targeted by the O_3 regulations, is primarily caused by the combustion of fossil fuels. Facilities can reduce NO_x emissions at a relatively low cost by adopting efficient combustion processes¹⁵ or by adopting relatively low-cost low- NO_x emissions burners (see, e.g. Envi-

¹⁴The 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).

¹⁵This may entail either changing the temperature or the fuel-oxygen ratio at which combustion occurs.

ronment Canada (2002), 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). In contrast, 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) and Environmental Protection Agency (1998, 2002, 1999b)).¹⁶

3 Theory

As we described above, the CWS 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 based on the work of Melitz (2003) and Bustos (2011) to help guide our empirical analysis. Below, we outline the features of the model and highlight its key empirical predictions. For the sake of brevity, we relegate details of the model’s solution and relevant derivations to the online appendix.

3.1 Setup

We consider a closed 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 labor be the numeraire.

The demand side of the economy is represented by a consumer who derives utility from the consumption of goods and disutility from aggregate pollution Z

¹⁶As 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 using a low-NO_x burner (Environmental Protection Agency, 2006).

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 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 of good. To enter the market, firms must pay a fixed cost f_e . Entrants then draw a productivity level φ from a common Pareto distribution $G(\varphi) = 1 - \varphi^{-k}$.

Upon observing their productivity draw, firms decide to exit or remain in the market. If they remain, firms are able to produce output x using one of two increasing returns to scale production technologies: business-as-usual b and state-of-the-art s . These technologies differ along two dimensions. First, they differ in their marginal labor costs: marginal labor costs are given by $1/\varphi$ for b and $1/[\alpha\varphi]$ for s , where $\alpha > 1$. Second, the two technologies feature different emission intensities: b 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 s features an emission intensity of $e_s = \kappa/[\gamma\varphi]$, where $\gamma > 1$, so total pollution from production is $z_s(\varphi) = [\kappa x]/[\gamma\varphi]$. Adopting the state-of-the-art technology is costly; upgrading to s requires that firms pay an additional fixed cost f_s .

If firms adopt b , they also have the option to retrofit (r) their technology so it has the same emission intensity as s . As such, the emission intensity of a retrofitted plant (e_r) is also $\kappa/[\gamma\varphi]$, and the total level of pollution generated by production is $z_r(\varphi) = [\kappa x]/[\gamma\varphi]$. However, retrofitting does not affect labor productivity, making it less costly than adopting s , meaning $f_r < f_s$.

After choosing their technology, firms produce output to maximize profits. In each period, firms face an exogenous probability of exit, δ .

3.2 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 the “no-regulation” equilibrium (*no*) in which pollution is not regulated. In this case labor is the

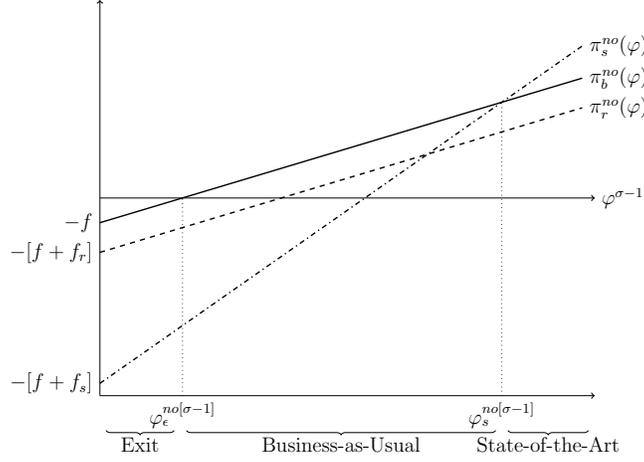


Figure 2: Technology Choices without Environmental Regulation

only variable cost of production. Below, we describe firm behaviour in this case; the industry equilibrium is presented in the online appendix.

Given the structure of consumer preferences, firms set prices at a constant mark-up over marginal costs. Hence, in the absence of regulation, firms that employ technologies b and r charge the same price: $p_b^{no}(\varphi) = p_r^{no}(\varphi) = 1/[\rho\varphi]$. If, instead, a firm employs technology s , it charges $p_s^{no}(\varphi) = 1/[\rho\alpha\varphi]$.

Firms choose between the three available technologies to maximize profits. If firms employ technology b , profits are given by $\pi_b^{no}(\varphi) = [1/\sigma]I [P\rho]^{\sigma-1} \varphi^{\sigma-1} - f$. Profits from employing technology r are $\pi_r^{no}(\varphi) = [1/\sigma]I [P\rho]^{\sigma-1} \varphi^{\sigma-1} - [f + f_r]$. Finally, profits from choosing technology s are given by $\pi_s^{no}(\varphi) = [1/\sigma]I [P\rho]^{\sigma-1} \varphi^{\sigma-1} \alpha^{\sigma-1} - [f + f_s]$.

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.¹⁷ As the figure shows, for productivity levels below φ_ϵ^{no} it is unprofitable for a firm to operate using any technology. Hence if a firm has a φ less than φ_ϵ^{no} , 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 $\varphi \in \{\varphi_\epsilon^{no}, \varphi_s^{no}\}$, then it will produce using technology

¹⁷To linearize this figure, we show profits as a function of $\varphi^{\sigma-1}$, not φ .

b. However, if a firm has a productivity level $\varphi > \varphi_s^{no}$, then the reduction in variable cost created by adopting *s* is great enough to justify the fixed cost of adoption, meaning that these firms adopt the technology *s*.

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

3.3 The Effects of Environmental Regulation

We now consider the effects of adopting environmental regulations similar to those imposed under the CWS. To build intuition, below we describe firm behaviour with regulation, holding industry prices fixed at the no-regulation level. The full effects of regulation allowing for industry prices to adjust are presented in the online appendix.

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 or the retrofitted baseline technologies), it is not subject to regulation because it is operating with the lowest emission intensity currently achievable for its productivity level. As a result, the marginal costs and profits from using these technologies are unaffected by regulation. 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.¹⁸ Given that firm prices feature a constant markup, this increase in marginal costs raises the price of output for firms producing with technology *b*. That is, $p_b^{no}(\varphi) < p_b^{cws}(\varphi) = [1 + \kappa\tau]/[\rho\varphi]$, so profits are now $\pi_b^{cws}(\varphi) = [1/\sigma]I [P\rho]^{\sigma-1} \varphi^{\sigma-1} [\frac{1}{1+\kappa\tau}]^{\sigma-1} - f$. This means, holding industry prices (*P*) fixed, the profit from using *b* falls for any level of productivity φ .

¹⁸Alternatively, 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. To ensure consistency with the design of the CWS, we assume that the revenue from the tax is not returned to consumers and is spent outside the model.

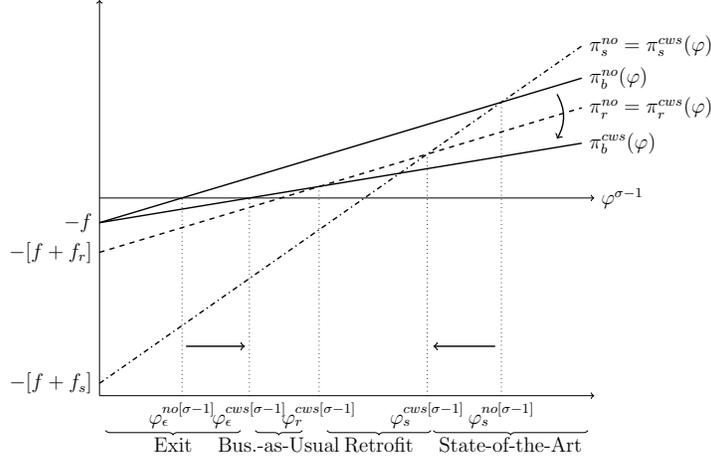


Figure 3: Technology Choices with CWS-Type Environmental Regulation

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 fixed. As the figure shows, a reduction in the profitability of technology b increases the productivity level for which it is unprofitable to enter the market from φ_ϵ^{no} to φ_ϵ^{cws} . As such, regulation creates a selection effect; firms with $\varphi \in \{\varphi_\epsilon^{no}, \varphi_\epsilon^{cws}\}$ exit in response to regulation. Moreover, the increase in technology b 's variable production cost makes technology upgrading a profitable alternative for some firms. As depicted, it is profit maximizing for firms with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^{no}\}$ to retrofit their technology in response to regulation. For these firms, the benefit of avoided emission tax payments outweighs the increase in fixed production costs. Similarly, firms with productivity $\varphi \in \{\varphi_c^{cws}, \varphi_s^{no}\}$ 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. These effects are instead displayed in Figure 4 for firms that survive regulation (those with $\varphi > \varphi_\epsilon^{cws}$). This figure depicts the effects of environmental regulation on firm revenues (Panel (a)) and emission intensity (Panel (b)) holding industry prices fixed. Both panels show that the most productive firms, with

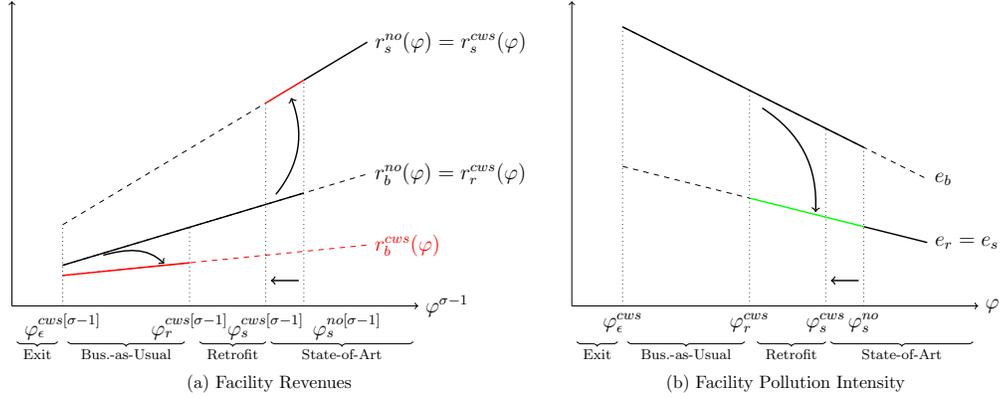


Figure 4: Revenues and Pollution Intensity for Surviving Firms with CWS-Type Environmental Regulation

productivity $\varphi > \varphi_s^{no}$, are unaffected by regulation because they use the most advanced technology regardless of regulatory regime. In contrast, regulation causes the least productive firms, with productivity $\varphi < \varphi_r^{cws}$, to produce less because they use technology b under either regime and regulation causes the variable cost of production with this technology to increase. However, as shown in Panel (a), this production is reallocated to firms with productivity $\varphi \in \{\varphi_s^{cws}, \varphi_r^{cws}\}$ because variable costs for these firms fall. Lastly, regulation induces a process effect by causing the emission intensity of firms in the middle of the productivity distribution, those with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^{no}\}$, to fall. This occurs because these firms either retrofit from technology b or adopt technology s .

While the discussion above is illustrative, in the online appendix we show that if $f_r > 0$ then in equilibrium:

1. Regulation causes some firms to exit.
2. On average, revenue and emission intensity fall in response to regulation.
3. The effects of regulation vary across the productivity distribution. Revenues fall for the least productive surviving firms using the business-as-usual technology, while emission intensity falls for firms in the middle of

the productivity distribution that retrofit or upgrade to the state-of-the-art technology.

In addition to the above results, in the online appendix, we show that the fixed cost of retrofitting (f_r) plays an important role in determining the channels through which regulation causes an industry to clean-up. Specifically, decreasing f_r increases the measure of firms that adopt a clean production process in response to regulation and reduces the measure of firms that exit in response to regulation.¹⁹ 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.

4 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 environmental regulations should primarily reduce industry emission intensity via reallocation and selection effects when the fixed costs of process changes are high. In contrast, when these fixed costs are low, the industry clean-up should be driven by process effects. In this section, we explore these predictions empirically by estimating the effects of the CWS on plant emissions, emission 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 Canadian manufacturing.

4.1 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.²⁰

¹⁹This result requires restricting the size of f_s .

²⁰It 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

We begin by comparing the average outcomes of plants in regulated CMAs while regulated 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 causal effect of the CWS. Yet, such absence is unlikely; there is strong reason to believe that this before-and-after comparison could also capture the effects of regional, industry, or aggregate economic shocks.

To control for 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 captures the effects of any unobserved time-varying provincial or CMA-level heterogeneity, such as changes in regional economic conditions or concurrent changes in provincial policy that would otherwise confound the effects of the CWS.

Similarly, to control for industry shocks, 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 comparison captures the effects of industry-specific shocks, such as increased foreign competition created by international trade, or revisions to federal policies that target certain sectors, that would otherwise confound identification.

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

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.

ing aggregate trend in pollution and production. 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 effects of regulation using the following equation

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

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 .²¹ T_{ict}^j is treatment indicator that takes the value of one for plants that are in industries targeted by the CWS for years in which their CMA exceeded standard j .

Equation (3) also includes plant (ρ_p), CMA-year (ξ_{ct}), industry-year²² (λ_{it}) fixed effects, and an error term (ε_{pict}). The plant fixed effects account for any time-invariant unobserved plant, industry or CMA heterogeneity. 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.^{23,24}

²¹We employ this transformation to address skewness in the distribution of each variable.

²²The CWS defined targeted industries at the 3- or 4-digit North American Industry Classification System (NAICS) level. We create an industry indicator to account for this. 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.

²³It 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.

²⁴We are able to separately estimate the effect of both standards because there are cities

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, $PM_{2.5}$ 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.²⁵ 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 transboundary pollution, Canada’s contribution to which involved ensuring the CWS was met (International Joint Commission, 2002). This means it is unlikely a single plant is able to directly manipulate its treatment status.

4.2 Data and Measurement

Our analysis relies on a unique confidential micro-dataset that contains information on the pollution emissions and productive activities 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). The NPRI contains information on the emissions of various pollutants from Canadian manufacturing plants.²⁶ The ASM provides information on output, production costs, employment, and other plant characteristics for most manufacturing plants in Canada. The matched dataset contains longitudinal, plant-level information on $PM_{2.5}$ and NO_X emissions, production, and other plant characteristics over the period 2004-2010. Further details on the dataset are given in the online appendix.

Descriptive statistics for the key variables that we employ are reported in

that exceed one, both, or none of the standards. Of all treated CMA-years in our sample, approximately 80% violated only one standard, while the remaining 20% violated both.

²⁵See, for example, Brankov et al. (2003) or Johnson et al. (2007).

²⁶By law, any facility that emits one of the covered pollutants above a minimum threshold must report to the NPRI.

Table 2: Summary Statistics

	PM _{2.5} (1)	NO _x (2)	Full ASM (3)
Emissions (tonnes)	25.83 (103.43)	262.14 (646.14)	
Sales (\$1 mill.)	194.62 (890.55)	342.15 (1,305.95)	11.12 (123.56)
Value Added (\$1 mill.)	62.46 (241.82)	102.11 (346.27)	4.29 (34.34)
Employment	280.11 (634.85)	382.03 (868.68)	35.69 (125.27)
VA/Worker (\$1,000)	200.18 (243.63)	265.41 (297.06)	84.78 (166.11)
<i>N</i>	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.

Table 2. Each column in Table 2 presents statistics for a different sample of plants. The first column corresponds to the set of plants that emit PM_{2.5}, while the second column corresponds to plants that emit NO_x.²⁷ 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 reports summary statistics for the entire 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 pol-

²⁷The 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 the online appendix for further details).

lution 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 process changes 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. Both samples of polluters sell more goods, employ more workers, and have higher value added per worker than the average manufacturing plant.²⁸ This is, in part, due to the reporting requirements of 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, 2016). While this means we systematically exclude small facilities, our analysis covers plants that account for the majority of manufacturing pollution in Canada.²⁹

4.2.1 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 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.³⁰

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 did not 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

²⁸This is still true when we consider medians instead of means.

²⁹In 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.

³⁰For more details on the construction of the hourly pollution concentration measures, see the online appendix.

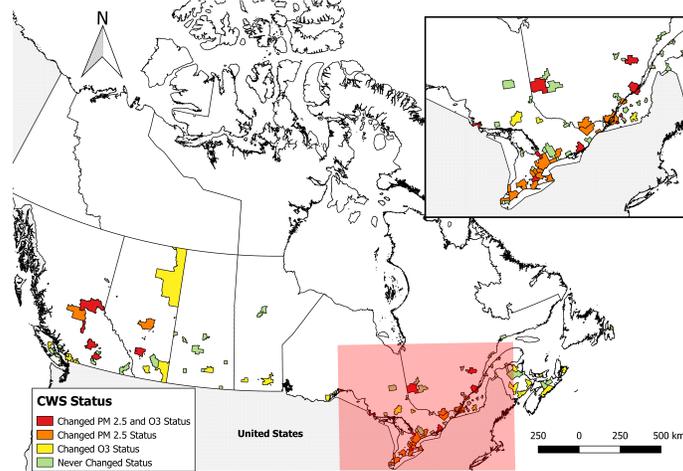


Figure 5: Regulatory Status Changes under the CWS

Notes: Figure depicts $PM_{2.5}$ and O_3 standard status changes for each CMAA from 2000 to 2010. Red CMAAs changed status under both the $PM_{2.5}$ and O_3 standards. Orange CMAAs only changed status for the $PM_{2.5}$ standard. Yellow CMAAs only changed status for the O_3 standard. Green CMAAs did not 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.

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.

4.3 Empirical Results

4.3.1 Plant Responses to the CWS

We begin our main empirical analysis by estimating the effects of the CWS on plant pollution emissions, emission intensity and output. These estimates are reported in Table 3 for two samples of plants. The first sample (in Panel A) are plants that emit $PM_{2.5}$, which is the main contributor to $PM_{2.5}$ pollution. The second sample (in Panel B) are plants that emit NO_x , which is the main contributor to O_3 pollution. The first column in each panel reports estimates of the effects of the CWS on plant pollution emissions. The second and third columns of each panel report estimates of the effects of the CWS on plant emission intensity, measured as the emissions-sales ratio (column two), or the emissions-value added ratio (column three). Finally, columns four and five

Table 3: The Effects of the CWS on Manufacturing Plants

	Panel A: PM _{2.5}				
	(1)	(2)	(3)	(4)	(5)
PM _{2.5} Std.	-0.151 (0.076)	-0.043 (0.096)	-0.013 (0.110)	-0.108 (0.050)	-0.138 (0.065)
O ₃ Std.	-0.113 (0.164)	-0.169 (0.169)	-0.224 (0.189)	0.056 (0.060)	0.111 (0.070)
R^2	0.175	0.161	0.156	0.224	0.221
N	6501	6501	6501	6501	6501
	Panel B: NO _x				
PM _{2.5} Std.	0.106 (0.069)	0.127 (0.080)	0.333 (0.098)	-0.022 (0.059)	-0.227 (0.083)
O ₃ Std.	-0.325 (0.179)	-0.286 (0.153)	-0.200 (0.157)	-0.039 (0.161)	-0.125 (0.188)
R^2	0.311	0.281	0.260	0.265	0.294
N	3012	3012	3012	3012	3012

Notes: Table reports estimates of the effects of the CWS on manufacturing plants. Each panel reports results for a different sample of emitters. Each column of each panel displays estimates from a different regression. In column one, the dependent variable is the natural log of pollution emissions. In columns two and three the dependent variable is the natural log of plant emission intensity, measured as the emission-sales ratio and the emissions-value added ratio respectively. The dependent variable in column four is the natural log of plant sales, while the dependent variable in column five is the natural log of plant value added. The first row of each panel reports the effects of the PM_{2.5} standard, and the second row reports the effects of the O₃ 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.

report estimates of the CWS's effects on plant output, measured in terms of sales (column four) or value-added (column five). In all cases, the dependent variable is log transformed. 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.³¹ In all cases, standard errors clustered at the CMA-industry level are reported in parentheses.

The estimates reported in column (1) of Table 3 indicate that the CWS led to significant reductions in the emissions of both PM_{2.5} and NO_x from

³¹In 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 the online appendix.

affected plants. For example, the estimate reported in the first row of Panel A indicates that the CWS particulate matter regulations are associated with a 15.1% reduction in $PM_{2.5}$ emissions from affected plants. Similarly, the estimate reported in the second row of Panel B indicates that the CWS ozone regulations are associated with 32.5% decrease in NO_X emissions from affected plants. The estimates reported in Panels A and B also suggest that there were no significant cross-effects of either standard. That is, O_3 regulation did not significantly affect particulate matter emissions and PM regulation did not significantly affect NO_X emissions.

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 2014.³² This suggests that the CWS had similar effects on pollution levels as the environmental policies enacted elsewhere.³³

The remaining columns of Table 3 highlight the mechanisms driving these emissions reductions. The estimates reported in Columns (2) and (3) of Panel A suggest that the CWS had little-to-no-effect on the emission intensity of the average $PM_{2.5}$ emitting plant, while the estimates reported in Columns (4) and (5) suggest that $PM_{2.5}$ regulation is associated with significant reductions in output from these plants. For example, our preferred estimate, reported in Column (4) indicates that the CWS particulate matter regulation is associated

³²Greenstone (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.

³³It is also worth noting that the estimates reported in column (1) 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 the online appendix.

with a 10.8% decrease in sales from that average plant that emitted $\text{PM}_{2.5}$.³⁴ In contrast, the estimates reported in Columns (2) to (5) of Panel B suggest that the CWS led to significant reductions in the emission intensity of the average NO_x emitting plants, but had no clear effect on output. For example, our preferred estimate of the effects of the CWS on emission intensity, reported in column (2) of Table 3, indicates that the CWS ozone regulations are associated with a 28.6% decrease in the level of NO_x emitted per unit of output.^{35,36} In contrast, our preferred estimate of the effects of CWS on output, reported in column (4), show the CWS ozone regulations are associated with a 3.9% decrease in output, but this estimate is imprecisely measured.³⁷

Lastly, we examine plant entry and exit decisions by estimating a variant of our main specification (Equation (3)). Specifically, we estimate

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

where N_{ict} is the number of active plants in industry i in CMA c and year t , 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

³⁴It is worth noting that 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 our attention on sales as a measure of output given that the previous literature has primarily employed shipments as its measure of output. 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.

³⁵Though 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.

³⁶From the estimate reported in Column (3) of Panel B, it also appears $\text{PM}_{2.5}$ regulation caused a significant increase in NO_x intensity measured in value added terms. This estimate is not robust; it is driven by a very small number of plants that are regulated by the $\text{PM}_{2.5}$ standard and emit NO_x , but not $\text{PM}_{2.5}$. For these plants, $\text{PM}_{2.5}$ regulation caused a large increase in NO_x emissions and decrease in value added. Dropping these plants yields reduces the estimated effect the $\text{PM}_{2.5}$ standard on NO_x emissions to 0.052 with a standard error of 0.073.

³⁷Note that $\text{PM}_{2.5}$ regulation also appears to have 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, and is not robust.

Table 4: The Effects of the CWS on Plant Exit

	Panel A: Emit PM _{2.5}		Panel B: Emit NO _x	
	(1)	(2)	(1)	(2)
PM _{2.5} Std.	-1.134 (0.626)	-0.347 (0.169)	-0.188 (0.293)	-0.031 (0.119)
O ₃ Std.	0.726 (0.547)	0.142 (0.147)	-0.457 (0.489)	-0.135 (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 O₃ regulations. All regressions include industry-year and CMA-year fixed effects, and an indicator for industry-CMAs that are ever treated by the CWS. Standard errors clustered by CMA are reported in parentheses.

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_{O_3}) measure the net exit or entry of plants in 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. These estimates are reported in columns (1) and (2) of Table 4, respectively for two groups of plants: those that emit PM_{2.5} (Panel A) and those that emit NO_x (Panel B). In all cases, standard errors clustered by CMA are reported in parentheses.

The estimates reported in Table 4 suggest that the CWS had a significant effect on the entry and exit decisions of some manufacturing plants. For example, the estimates reported in column (1) of Panel A show PM_{2.5} regulation reduced the number of PM_{2.5} emitters in the average affected industry-CMA by 1.134 plants.³⁸ In contrast, O₃ regulation had no significant effect on the net exit or entry of NO_x emitters.

³⁸Note that there is no significant difference in the incident rate ratio between the ordinary least squares and Poisson regressions.

Altogether, the results reported in Table 3 and Table 4 suggest that CWS had different effects on plants that emit different pollutants. On average, emissions of $\text{PM}_{2.5}$ and NO_x from affected plants fell, but these reductions were primarily due to changes in emission intensity for NO_x emitters, and output for $\text{PM}_{2.5}$ emitters. Moreover, we find that the CWS led to the exit of plants that emit $\text{PM}_{2.5}$. These findings are consistent with our theory, which predicts the effects of a CWS-style regulation will depend on the fixed costs of adopting cleaner processes. If fixed costs are high, as in the case of $\text{PM}_{2.5}$, then only relatively productive plants will adopt cleaner production processes in response regulation; most plants will experience a reduction in output or exit because they are relatively unproductive. As a result, in this case, a CWS-style regulation should have little to no effect on the emission intensity of the average plant, but cause average output to fall and plants to exit, as we find. If, instead, fixed costs are low, as with NO_x , then even less productive plants should adopt cleaner production processes. As we observe, in this case, the emission intensity of the average plant should fall in response to the CWS, and regulation should have little effect on output or entry and exit.

4.3.2 Aggregate Implications

The results presented in Table 3 and Table 4 also suggest that the CWS contributed to the clean-up of Canadian manufacturing through different channels for different pollutants. To quantify the relative importance of these channels, 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 effect. We then compare these estimates to the observed change in manufacturing pollution intensity.³⁹

To develop an empirical analogue to Equation (2), we follow an approach

³⁹For 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.

used in much of the labor literature and consider total changes in emission intensity over time.⁴⁰ To start, let t index time, 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. Given this setup, the percentage change in an industry's emission intensity, $\dot{E}_{it} = \frac{E_{it}-E_{it-1}}{E_{it-1}}$, can be expressed as⁴¹

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

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 1.1. 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 empirical estimates 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 effects of the CWS on plant emission intensity and output (from Table 3), and selection (from Table 4), respectively. Moreover, recall that, given our identification assumptions, $\hat{\beta}_e$ captures the average change in emission intensity due to the CWS, meaning we can write

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

In addition, an estimate of $\dot{\lambda}_{it}(n)$ and $\int_{n_{it}}^{n_{it-1}} s_{zit-1}^z(n)dn$ can be constructed

⁴⁰For a review of this approach, see Foster et al. (2001).

⁴¹A detailed derivation is given in the online appendix.

from $\hat{\beta}_x$ and $\hat{\beta}_n$, respectively. In the online appendix, we show

$$\lambda_{it}(n) = \begin{cases} \frac{\hat{\beta}_x(1-s_{xit-1}^{Treat})+s_{xit-1}^{Exit}}{1-s_{xit-1}^{Exit}+\hat{\beta}_x s_{xit-1}^{Treat}}, & \text{if } n \text{ is treated} \\ \frac{s_{xit-1}^{Exit}-\hat{\beta}_x s_{xit-1}^{Treat}}{1-s_{xit-1}^{Exit}+\hat{\beta}_x s_{xit-1}^{Treat}}, & \text{otherwise,} \end{cases} \quad (7)$$

where s_{xit-1}^{Treat} and s_{xit-1}^{Exit} 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_{zit-1}^{Treat} be the share of industry i 's pollution at time $t-1$ from treated plants, the process effect is given by

$$\widehat{PE} = \hat{\beta}_e s_{zit-1}^{Treat}. \quad (8)$$

Similarly, the reallocation effect is given by

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

and the interaction effect is given by

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

To construct an estimate of the selection effect, recall our estimate of $\hat{\beta}_n$ captures the average 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, the selection effect is given by

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

where \bar{s}_{zit-1}^{Exit} is the average exiting plant's share of industry i 's pollution.

Our estimates of each channel's contribution to the observed change in manufacturing emission intensity is reported in Table 5. The first row reports the fraction of the PM_{2.5} clean-up due to the CWS, while the second reports

Table 5: 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.

the fraction of the NO_X clean-up due to the CWS. Columns (1)-(4) report our estimates of the process effect, reallocation effect, selection effect and interaction effect for each pollutant. 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 O₃ 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 O₃ 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 PM_{2.5}. Instead, the PM_{2.5} regulation primarily reduced aggregate emission intensity through a combination of reallocation and selection effects.

4.3.3 Mechanism: Additional Evidence

The findings presented above suggest that the CWS contributed to the clean up of Canadian manufacturing via different channels for different pollutants. We have attributed this finding to differences in the fixed costs of process changes across pollutants. However, we do not observe these costs directly. As such we now probe the underlying mechanism to the extent possible with our data.

We begin by examining how the effects of the CWS vary across the plant productivity distribution. Our model suggests that regulation should: (i) have little effect on the most productive plants, (ii) reduce emissions by reducing output from the least productive plants and the emission intensity of plants in the middle of the productivity distribution when fixed costs are high, and (iii) reduce emissions by lowering the emission intensity of both low and middle productivity plants when fixed process costs are low.

To test these predictions, 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

$$Y_{pict} = \sum_{q=1}^3 \beta_{PM}^{Q_q} [T_{ict}^{PM} \times Q_q] + \sum_{q=1}^3 \beta_{O_3}^{Q_q} [T_{ict}^{O_3} \times Q_q] + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \quad (12)$$

where Q_q is an indicator that takes the value one if plants that are in productivity tercile q , $\beta_j^{Q_q}$ measures the effect of standard j on plants in productivity tercile q , and the remaining variables are as defined for Equation (3).⁴²

We use Equation (12) to examine the CWS' effects on plant pollution levels, emission intensity, and sales. The resulting estimates are reported in Table 6, for each sample of emitters. Each panel corresponds to a different sample: Panel A reports our estimates for PM_{2.5} emitters, while Panel B for NO_x emitters. The first column of each panel reports estimates of the CWS on plant emissions, the second reports estimates of the effects of the CWS on plant emissions per dollar of sales, and the third reports estimates of the effects of the CWS plant sales. As before, natural logarithms are taken of all

⁴²We 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 PM_{2.5} and NO_x emitters, the composition of plants in each tercile may vary across each pollutant sample.

dependent variables. The first three rows in each panel report the effects of the PM_{2.5} regulation (the β_{PM}^{Qq} coefficients in Equation (12)): the first reports the estimated effect of the CWS PM_{2.5} regulation on plants in the lowest productivity tercile, while the second and third rows report this effect for plants in the middle and highest terciles, respectively. Similarly, the final three rows report the effects of the O₃ regulation ($\beta_{O_3}^{Qq}$ in Equation (12)), again by productivity tercile. 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 estimates reported in Panel A of Equation (12) indicate the PM_{2.5} standard had substantially different effects on PM_{2.5} emitting plants of different productivity levels. PM_{2.5} regulation caused a drop in emissions among the bottom two-thirds of the productivity distribution, with an average reduction in emissions of 16.3% for low productivity plants and 27.9% for middle productivity PM_{2.5} emitters. In contrast, PM_{2.5} regulation led to a 2.3% reduction for the most productive PM_{2.5} plants, but this effect is imprecisely estimated.

The estimates reported in columns (2) and (3) indicate the mechanisms through which affected surviving PM_{2.5} polluters reduced their emissions varied considerably across plants of different productivities. 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.

The estimates reported in Panel B suggest that the O₃ standard also had different effects on NO_x emitting plants of different productivity levels. As the results reported in Column (4) show, NO_x emissions appear to have fallen considerably for low and middle productivity plants, with estimated reductions of 34 and 46%, respectively. In contrast, the regulation had much smaller effects on the most productive plants; emissions from these plants fell by 18%

Table 6: The Effects of the CWS by Plant Productivity Level

	Panel A: PM _{2.5}			Panel B: NO _x		
	(1)	(2)	(3)	(4)	(5)	(6)
PM _{2.5} Std. x Q1	-0.163 (0.083)	0.038 (0.102)	-0.201 (0.073)	0.079 (0.091)	0.084 (0.118)	-0.005 (0.084)
PM _{2.5} Std. x Q2	-0.279 (0.134)	-0.251 (0.143)	-0.028 (0.056)	0.155 (0.120)	0.188 (0.126)	-0.032 (0.096)
PM _{2.5} Std. x Q3	-0.023 (0.100)	-0.016 (0.101)	-0.007 (0.057)	0.079 (0.134)	0.109 (0.134)	-0.030 (0.056)
O ₃ Std. x Q1	-0.281 (0.210)	-0.353 (0.222)	0.072 (0.074)	-0.457 (0.207)	-0.412 (0.207)	-0.045 (0.205)
O ₃ Std. x Q2	0.076 (0.195)	0.065 (0.227)	0.011 (0.130)	-0.340 (0.173)	-0.277 (0.160)	-0.063 (0.056)
O ₃ Std. x Q3	-0.093 (0.237)	-0.150 (0.232)	0.057 (0.071)	-0.183 (0.177)	-0.182 (0.180)	-0.001 (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 O₃ 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.

on average, although this estimate is imprecise. The estimates reported in Columns (5) and (6) suggest the observed NO_x reductions at low and middle productivity plants was primarily driven by changes in plant emission intensity, as emission intensity changes explain over 80% of the reduction in emissions.

The results in Table 6 are consistent with our theory, and the hypothesis that the channels of the CWS clean-ups varied across pollutants because of differences in fixed process costs.⁴³ As we noted above, when fixed process

⁴³This 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.

costs are high, as with $\text{PM}_{2.5}$, only relatively productive plants that have not already adopted leading technologies should choose to do so. These relatively productive plants, in turn, experience a reduction in pollution intensity with a relatively small change in output and production inputs. 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. In both cases, the most productive plants use state of the art technology, and are thus unaffected by the CWS.

It is worth noting that there are other possible explanations for the findings in Section 4.3.1, such as differences in opportunities for input substitution across pollutants or differential responses to innovation induced by regulation (i.e. the Porter Hypothesis). In the online appendix we examine these alternative hypotheses in detail and find no evidence to suggest they are driving our results, providing further evidence that our findings are driven by differences in the fixed costs of process changes across pollutants.

5 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 environ-

mental 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 are informative for all three regions.

Our estimates imply that this policy explains approximately 60% of the reduction in nitrogen oxide emission intensity of the Canadian manufacturing sector, and approximately 20% of the reduction 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.

Our analysis suggests that these stark differences in responses to policy are due to differences in the fixed costs of adopting cleaner production processes across pollutants. Such differences may have broader implications for the clean-up of manufacturing; other hypothesized channels, such as technical change or international trade, may be affected by how fixed process costs influence plant decisions. However, we leave an examination of such effects to future work.

References

- Anouliès, L. (2017). Heterogeneous firms and the environment: A cap-and-trade program. *Journal of Environmental Economics and Management* 84, 84–101.
- 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 refineries. *Review of Economics and Statistics* 83(3), 498–510.
- Brankov, E., R. F. Henry, K. L. Civerolo, W. Hao, S. Rao, P. Misra, R. Bloxam, and N. Reid (2003). Assessing the effects of transboundary ozone pollution between ontario, canada and new york, usa. *Environmental Pollution* 123, 403–411.

- Brunel, C. (2017). Pollution offshoring and emissions reductions in e.u. and u.s. manufacturing. *Environmental and Resource Economics* 68(3), 621–641.
- 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.
- 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/.
- 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.
- Environment and Climate Change Canada (2016). Summary of National Pollutant Release Inventory Reporting requirements. <https://www.ec.gc.ca/inrp-npri/default.asp?lang=en&n=629573FE-1>.
- Environment Canada (2002). Multi-Pollutant Emission Reduction Analysis Foundation (MERAFA) for the Iron and Steel Sector. Unpublished report.
- Environmental Protection Agency (1998). Particulate Matter Controls. Baghouses and Filters. EPA/452/B-02-001. <https://www3.epa.gov/ttnecat1/cica/files/cs6ch1.pdf>.
- Environmental Protection Agency (1999a). Nitrogen Oxides (NOx), Why and How they are Controlled. EPA-456/F-99-006R. <https://www3.epa.gov/ttnecat1/cica/files/fnoxdoc.pdf>.

- Environmental Protection Agency (1999b). Particulate Matter Controls. Electrostatic Precipitators. EPA/452/B-02-001. <https://www3.epa.gov/ttn/catc1/cica/files/cs6ch3.pdf>.
- Environmental Protection Agency (2002). Particulate Matter Controls. Wet Scrubbers for Particulate Matter. EPA/452/B-02-001. <https://www3.epa.gov/ttn/catc1/cica/files/cs6ch2.pdf>.
- Environmental Protection Agency (2006). Chaptre 3. Regulatory Impact Analysis for the Particulate Matter NAAQS. <https://www3.epa.gov/ttn/ecas/regdata/RIAs/Chapter%203--Controls.pdf>.
- Environmental Protection Agency (2016). NAAQS Table. <https://www.epa.gov/criteria-air-pollutants/naaqs-table/>.
- Foster, L., J. Haltiwanger, and C. J. Krizan (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In C. Hulten, E. Dean, and M. Harper (Eds.), *New Developments in Productivity Analysis*. University of Chicago Press.
- 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(2), 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&s1=0>.
- 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), 442–448.
- 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 so₂ emissions. *Environmental and Resource Economics* 43(2), 257–274.
- Henderson, J. V. (1996). Effects of air quality regulation. *The American Economic Review* 86(4), 789–813.
- 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 identification 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 run. *Journal of the Association of Environmental and Resource Economists* 2(1), 1–42.
- 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.
- Shapiro, J. S. and R. Walker (2015). Why is Pollution from U.S. Manufacturing Declining? The Roles of Trade, Regulation, Productivity, and Preferences. Working Paper 20879, National Bureau of Economic Research.
- 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.

Appendices for: Environmental Regulations and the Clean-Up of Manufacturing: Plant-Level Evidence from Canada.

Nouri Najjar* Jevan Cherniwchan†

For Online Publication

This document contains additional information on the theory (Appendix A), empirics (Appendix B) and additional results (Appendix C) referenced in the main text.

Appendix A Theory

In this section we provide further details of the theoretical model presented in Section 3.

A.1 The No-Regulation Equilibrium

We begin with the no-regulation equilibrium (*no*) described in Section 3.2. We start by deriving the productivity cutoffs for firm exit and technology upgrading depicted in Figure 2. Given our assumptions, the marginal firm uses the business-as-usual technology b , meaning the no-regulation exit cutoff φ_ϵ^{no} can be determined by noting that $\pi_b^{no}(\varphi_\epsilon^{no}) = 0$. Substituting for $\pi_b^{no}(\varphi_\epsilon^{no})$ and rearranging yields

$$\varphi_\epsilon^{no} = \left[\frac{\sigma f}{I} \right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{no}}. \quad (1)$$

*Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC, Canada, V6T 2E8. Email: npmnajjar@gmail.com.

†Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta, Canada, T6G 2R6. Email: jevan@ualberta.ca. Tel: 780-492-2944.

APPENDIX

Firms upgrade to the state-of-the-art technology s or retrofit (r) the baseline technology when it is profit maximizing to do so. It is easy to verify that $\pi_b^{no}(\varphi) > \pi_r^{no}(\varphi)$ for any productivity level given $f_r > 0$, meaning firms never choose to retrofit in the no-regulation equilibrium. The productivity cutoff for technology-upgrading, φ_s^{no} is defined by $\pi_b^{no}(\varphi_s^{no}) = \pi_s^{no}(\varphi_s^{no})$. Substituting and rearranging yields

$$\varphi_s^{no} = \left[\frac{\sigma f_s}{\Delta_1 I} \right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{no}}, \quad (2)$$

where $\Delta_1 = \alpha^{\sigma-1} - 1 > 1$.

By combining equations (1) and (2), we can write the technology upgrading cutoff as a function of the exit cutoff

$$\varphi_s^{no} = \varphi_\epsilon^{no} \left[\frac{f_s}{\Delta_1 f} \right]^{\frac{1}{\sigma-1}}. \quad (3)$$

Given that the CWS relied on leading state of the art technologies as a benchmark for regulation, we assume $f_s > \Delta_1 f$ to ensure that the technology upgrading cutoff is always greater than the exit cutoff, meaning both technologies are used in the no-regulation equilibrium.

We next turn to solve for the industry equilibrium. To do so, we exploit the fact that because of free entry, in expectation, firms earn zero discounted profits. This means that in equilibrium, the fixed entry cost f_ϵ , must equal the present value of expected profits. That is

$$f_\epsilon = \frac{1 - G(\varphi_\epsilon^{no})}{\delta} \bar{\pi}^{no}, \quad (4)$$

where $\bar{\pi}^{no}$ are a firm's expected profits conditional on surviving, given by

$$\bar{\pi}^{no} = \int_{\varphi_\epsilon^{no}}^{\varphi_s^{no}} \pi_b^{no}(\varphi) \frac{g(\varphi)}{1 - G(\varphi_\epsilon^{no})} d\varphi + \int_{\varphi_s^{no}}^{\infty} \pi_s^{no}(\varphi) \frac{g(\varphi)}{1 - G(\varphi_s^{no})} d\varphi. \quad (5)$$

After substituting for $\pi_b^{no}(\varphi)$ and $\pi_s^{no}(\varphi)$ and exploiting equations (1) and (3),

APPENDIX

it is possible to show (with considerable algebra)

$$\bar{\pi}^{no} = \frac{[\sigma - 1]f}{k - \sigma + 1} \Psi_1, \quad (6)$$

where

$$\Psi_1 = \left[1 + \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s} \right]^{\frac{k-\sigma+1}{\sigma-1}} \right]. \quad (7)$$

Hence, the exit cutoff φ_ϵ^{no} can be obtained by substituting equation (6) into equation (4) and noting $1 - G(\varphi) = \varphi^{-k}$. This yields

$$\varphi_\epsilon^{no} = \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_1 \right]^{\frac{1}{k}}. \quad (8)$$

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

With an expression for φ_ϵ^{no} it is possible to solve for the technology upgrading cutoff and the aggregate price index. The technology upgrading cutoff φ_s^{no} can be determined by substituting equation (8) into equation (3)

$$\varphi_s^{no} = \left[\frac{f_s}{\Delta_1 f} \right]^{\frac{1}{\sigma-1}} \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_1 \right]^{\frac{1}{k}}. \quad (9)$$

The price index can be obtained by first noting that $I = L$, and then substituting equation (8) into equation (1) and rearranging to obtain

$$P^{no} = \left[\frac{\sigma f}{L} \right]^{\frac{1}{\sigma-1}} \left[\rho \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_1 \right]^{\frac{1}{k}} \right]^{-1}. \quad (10)$$

A.2 The CWS Equilibrium

Having solved the No-Regulation equilibrium, we now turn to solve the equilibrium featuring CWS-style regulations used in the analysis presented in Section 3.3. We again start by deriving productivity cutoffs; however, in this case, we now also consider the cut-off for retrofitting in addition to the cut-offs

APPENDIX

for firm exit, and technology upgrading.

Again, the marginal firm uses the business-as-usual technology b , meaning the CWS exit cutoff φ_ϵ^{no} can be determined by noting that $\pi_b^{cws}(\varphi_\epsilon^{cws}) = 0$. Substituting for $\pi_b^{cws}(\varphi_\epsilon^{cws})$ and rearranging yields

$$\varphi_\epsilon^{cws} = \left[\frac{\sigma f}{I} \right]^{\frac{1}{\sigma-1}} \left[\frac{1 + \tau\kappa}{\rho P^{cws}} \right]. \quad (11)$$

While firms are endowed with the business-as-usual technology, they will retrofit or upgrade to the state-of-the-art technology if it is profitable to do so. The productivity cutoff for retrofitting with regulation, φ_r^{cws} is defined by $\pi_b^{cws}(\varphi_b^{cws}) = \pi_r^{cws}(\varphi_r^{cws})$. Substituting and rearranging yields

$$\varphi_r^{cws} = \left[\frac{\sigma f_r}{I \Delta_2} \right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{cws}}, \quad (12)$$

where $\Delta_2 = 1 - \frac{1}{[1+\tau\kappa]^{\sigma-1}} > 0$. Similarly, the productivity cutoff for technology upgrading with regulation, φ_s^{cws} is defined by $\pi_r^{cws}(\varphi_r^{cws}) = \pi_s^{cws}(\varphi_s^{cws})$. Substituting and rearranging yields

$$\varphi_s^{cws} = \left[\frac{\sigma [f_s - f_r]}{\Delta_1 I} \right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{cws}}. \quad (13)$$

Both φ_r^{cws} and φ_s^{cws} can be expressed as functions of the exit cutoff productivity φ_ϵ^{cws} . An expression for the retrofitting cutoff, φ_r^{cws} , can be obtained by combining equation (12) with (11)

$$\varphi_r^{cws} = \frac{\varphi_\epsilon^{cws}}{1 + \tau\kappa} \left[\frac{f_r}{\Delta_2 f} \right]^{\frac{1}{\sigma-1}}. \quad (14)$$

Similarly, an expression for the upgrading cutoff, φ_s^{cws} , can be obtained by combining equation (13) with (11)

$$\varphi_s^{cws} = \frac{\varphi_\epsilon^{cws}}{1 + \tau\kappa} \left[\frac{f_s - f_r}{\Delta_1 f} \right]^{\frac{1}{\sigma-1}}. \quad (15)$$

APPENDIX

To ensure that the model is consistent with the CWS, we impose the additional assumptions: $f_r > [(1 + \tau\kappa)^{\sigma-1} - 1]f$, and $f_s > \frac{\Delta_1 + \Delta_2}{\Delta_2} f_r$. This ensures the technology-upgrading cutoff is always greater than the retrofitting cutoff, meaning all three technologies (b , s , and r) are used in the CWS equilibrium.

To solve for the industry equilibrium, we again exploit the fact that in expectation, firms earn zero discounted profits due to free entry. Hence, in the CWS equilibrium, the fixed entry cost f_ϵ must equal the present value of expected profits

$$f_\epsilon = \frac{1 - G(\varphi_\epsilon^{cws})}{\delta} \bar{\pi}^{cws}, \quad (16)$$

where $\bar{\pi}^{cws}$ are a firm's expected profits conditional on surviving, given by

$$\begin{aligned} \bar{\pi}^{cws} &= \int_{\varphi_\epsilon^{cws}}^{\varphi_r^{cws}} \pi_b^{cws}(\varphi) \frac{g(\varphi)}{1 - G(\varphi_\epsilon^{cws})} d\varphi \\ &+ \int_{\varphi_r^{cws}}^{\varphi_s^{cws}} \pi_r^{cws}(\varphi) \frac{g(\varphi)}{1 - G(\varphi_\epsilon^{cws})} d\varphi \\ &+ \int_{\varphi_s^{cws}}^{\infty} \pi_s^{cws}(\varphi) \frac{g(\varphi)}{1 - G(\varphi_\epsilon^{cws})} d\varphi. \end{aligned} \quad (17)$$

Substituting for $\pi_b^{cws}(\varphi)$, $\pi_r^{cws}(\varphi)$, and $\pi_s^{cws}(\varphi)$ and utilizing equations (11), (12) and (13), it is possible to show (again with considerable algebra)

$$\bar{\pi}^{cws} = \frac{[\sigma - 1]f}{k - \sigma + 1} \Psi_2. \quad (18)$$

where:

$$\Psi_2 = \left[1 + [1 + \tau\kappa]^k \left[\Delta_2^{\frac{k}{\sigma-1}} \left[\frac{f}{f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s - f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} \right] \right] > 0. \quad (19)$$

The exit cutoff φ_ϵ^{cws} can be obtained by substituting (18) into (16) and using $1 - G(\varphi) = \varphi^{-k}$

$$\varphi_\epsilon^{cws} = \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_2 \right]^{\frac{1}{k}}. \quad (20)$$

APPENDIX

Having determined φ_ϵ^{cws} , it is again possible to obtain expressions for φ_r^{cws} and φ_s^{cws} , and the price index, P^{cws} . The retrofitting cutoff can be obtained by substituting equation (20) into equation (14)

$$\varphi_r^{cws} = \frac{1}{1 + \tau\kappa} \left[\frac{f_r}{\delta_2} \right] \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_2 \right]^{\frac{1}{k}}. \quad (21)$$

The technology cutoff, on the other hand, can be obtained by substituting equation (20) into equation (15)

$$\varphi_s^{cws} = \frac{1}{1 + \tau\kappa} \left[\frac{f_s - f_r}{\delta_1 f} \right]^{\frac{1}{\sigma-1}} \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_2 \right]^{\frac{1}{k}}. \quad (22)$$

An expression for the price index can be obtained by substituting equation (20) into equation (11) and noting that given our assumption that environmental tax revenues are not returned to consumers¹ $I = L$. This yields

$$P^{cws} = \left[\frac{\sigma f}{L} \right]^{\frac{1}{\sigma-1}} \left[\rho \left[\left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{f}{\delta f_\epsilon} \right] \Psi_2 \right]^{\frac{1}{k}} \right]^{-1}. \quad (23)$$

A.3 The Effects of Environmental Regulation

We now turn to examine the effects of environmental regulation on entry and exit, revenues and emission intensity by comparing equilibrium outcomes under the CWS to those under the no-regulation regime.

A.3.1 Entry and Exit

We begin by examining how regulation affects entry and exit. The effects of regulation on entry and exit can be determined by comparing φ_ϵ^{no} and φ_ϵ^{cws} .

¹As we note in the main text, we make this assumption to ensure our setup is consistent with the design of the CWS.

APPENDIX

Doing so yields

$$\frac{\varphi_\epsilon^{cws}}{\varphi_\epsilon^{no}} = \left[\frac{1 + [1 + \tau\kappa]^k \left[\Delta_2^{\frac{k}{\sigma-1}} \left[\frac{f}{f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s - f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} \right]}{1 + \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s} \right]^{\frac{k-\sigma+1}{\sigma-1}}} \right]^k. \quad (24)$$

A sufficient condition to ensure $\varphi_\epsilon^{cws}/\varphi_\epsilon^{no} > 1$ is

$$[1 + \tau\kappa]^k \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s - f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} > \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s} \right]^{\frac{k-\sigma+1}{\sigma-1}}. \quad (25)$$

This condition can be re-expressed as $[1 + \tau\kappa]^k [f_s/(f_s - f_r)]^{\frac{k-\sigma+1}{\sigma-1}} > 1$, which is satisfied given the assumptions of the model. Hence, CWS-style regulation causes some firms to exit the market.

A.3.2 Technology Upgrading

Next, we turn to examine how regulation affects firm technology choices, as these choices directly determine firm revenues and emission intensity, our key outcomes of interest.

The effects of environmental regulation on retrofitting are relatively straightforward. In the no-regulation regime, no facility retrofits. In the CWS regime, a facility with $\varphi \in [\varphi_r^{cws}, \varphi_s^{cws}]$ adopts the retrofit technology. If $f_s > \frac{\Delta_1 + \Delta_2}{\Delta_2} f_r$ and $f_r > [[1 + \tau\kappa]^{\sigma-1} - 1] f$, then a positive measure of firms retrofit in response to environmental regulation. These inequalities are satisfied by assumption, meaning CWS-style regulation causes some firms to retrofit in response to regulation.

Regulation also affects the adoption of the state-of-the-art technology, however, its effects are ambiguous. To see this, note that the ratio of s technology adoption cut-offs under the CWS and no regulation regimes can be written as

$$\frac{\varphi_h^{cws}}{\varphi_h^n} = \frac{1 + [1 + \tau\kappa]^k \left[\Delta_2^{\frac{k}{\sigma-1}} \left[\frac{f}{f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}} \left(\frac{f}{f_s - f_r} \right)^{\frac{k-\sigma+1}{\sigma-1}} \right]}{[1 + \tau\kappa]^{\frac{1}{k}} \left[\frac{f_s - f_r}{f_s} \right]^{\frac{1}{k(\sigma-1)}} \left[1 + \Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s} \right]^{\frac{k-\sigma+1}{\sigma-1}} \right]}. \quad (26)$$

APPENDIX

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

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

Hence, the effects of the CWS on state-of-the-art technology adoption depend on f ; if f is relatively small, then regulation increases the number of firms using the s technology, but f is large enough, then regulation reduces the number of firms using the s technology.

A.3.3 Revenues

We now turn to examine the effects of regulation on firm revenues. Before doing so, it is useful to note that in either regulatory regime, revenues for any facility (and thus profits) can be written as a monotonic function of the exit cut-off. To see this, note that in the no-regulation equilibrium, revenues at a facility using the business-as-usual technology are given by $r_b^{no}(\varphi) = [1/\rho\varphi]^{1-\sigma} IP^{\sigma-1}$, and revenues at a firm using the state-of-the-art technology are given by $r_s^{no}(\varphi) = [1/\rho\alpha\varphi]^{1-\sigma} IP^{\sigma-1}$. Using the fact that free entry implies $r_b^{no}(\varphi_\epsilon^{no}) = \sigma f$, we have

$$r_b^{no}(\varphi) = \left[\frac{\varphi}{\varphi_\epsilon^{no}} \right]^{\sigma-1} \sigma f \quad (27)$$

$$r_s^{no}(\varphi) = \left[\frac{\alpha\varphi}{\varphi_\epsilon^{no}} \right]^{\sigma-1} \sigma f. \quad (28)$$

Similarly, in the CWS equilibrium, revenues at a facility using the business-as-usual technology are given by $r_b^{cws}(\varphi) = [[1 + \kappa\tau]/\rho\varphi]^{1-\sigma} IP^{\sigma-1}$, revenues at a firm using a retrofitted technology are given by $r_r^{cws}(\varphi) = [1/\rho\varphi]^{1-\sigma} IP^{\sigma-1}$, and revenues at a firm using the state-of-the-art technology are given by $r_s^{cws}(\varphi) = [1/\rho\alpha\varphi]^{1-\sigma} IP^{\sigma-1}$. Again using the fact that free entry implies $r_b^{cws}(\varphi_\epsilon^{cws}) = \sigma f$,

APPENDIX

we have

$$r_b^{cws}(\varphi) = \left[\frac{\varphi}{\varphi_\epsilon^{cws}} \right]^{\sigma-1} \sigma f, \quad (29)$$

$$r_r^{cws}(\varphi) = \left[\frac{\varphi}{\varphi_\epsilon^{cws}} \right]^{\sigma-1} [1 + \kappa\tau]^{\sigma-1} \sigma f \quad (30)$$

$$r_s^{cws}(\varphi) = \left[\frac{\alpha\varphi}{\varphi_\epsilon^{cws}} \right]^{\sigma-1} [1 + \kappa\tau]^{\sigma-1} \sigma f. \quad (31)$$

We utilize equations (27)-(31) to determine the effects of regulation on firm revenue. The effects of regulation depend on the firms' technology choices in each regulatory regime; 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.

To start, we examine the effects of regulation at firms that use the business-as-usual technology in both the no regulation and CWS regimes. For these firms, the effect of regulation can be determined by comparing equations (27) with (29), yielding

$$\frac{r_b^{cws}(\varphi)}{r_b^{no}(\varphi)} = \left[\frac{\varphi_\epsilon^{no}}{\varphi_\epsilon^{cws}} \right]^{\sigma-1}. \quad (32)$$

Given $\varphi_\epsilon^{cws}/\varphi_\epsilon^n > 1$ (as shown above), it follows that $r_b^{cws}(\varphi) < r_b^{no}(\varphi)$, meaning CWS style regulation reduces revenue for firms that use the business-as-usual technology in both equilibria.

Next, we consider firms that retrofit from technology b to technology r . For these firms, the effect of regulation can be determined by comparing equations (27) with (30):

$$\frac{r_r^{cws}(\varphi)}{r_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^{no} [1 + \tau\kappa]}{\varphi_\epsilon^{cws}} \right]^{\sigma-1} \quad (33)$$

APPENDIX

Note that

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

which is negative if, and only if, the fixed cost of production f is large enough. That is, f must satisfy

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

Hence, regulation will reduce revenues at firms that retrofit the business-as-usual technology if the fixed cost of production are sufficiently high.

For facilities that use state-of-the-art technology in both regimes, the effects of regulation can be determined by comparing equations (28) and (31). This yields

$$\frac{r_s^{cws}(\varphi)}{r_s^{no}(\varphi)} = \left[\frac{\varphi_\epsilon^{no} [1 + \tau\kappa]}{\varphi_\epsilon^{cws}} \right]^{\sigma-1}, \quad (34)$$

which is the same as the condition for retrofitting facilities given above. Thus, regulation reduces revenues for these facilities if and only if the fixed cost of production, f , is sufficiently high.

If $\varphi_s^{cws} > \varphi_s^{no}$, then regulation causes some facilities to switch from the business-as-usual technology to the state-of-the-art technology. In this case

$$\frac{r_s^{cws}(\varphi)}{r_b^{no}(\varphi)} = \left[\frac{\varphi_\epsilon^{no} [1 + \tau\kappa] \alpha}{\varphi_\epsilon^{cws}} \right]^{\sigma-1}. \quad (35)$$

APPENDIX

This is greater than one if and only if

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{\alpha^{\frac{1}{k}} [1 + \tau\kappa]^{\frac{1}{k}} - 1}{\alpha^{\frac{1}{k}} [1 + \tau\kappa]^{\frac{1}{k}}} \right] \left[\left[\frac{[[1 + \tau\kappa]^{\sigma-1} - 1]^{\frac{k}{\sigma-1}}}{[1 + \tau\kappa]^{\frac{1}{k}}} \right] \left[\frac{1}{f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} \frac{1}{\alpha^{\frac{1}{k}}} + \left[[1 + \tau\kappa]^{\frac{k^2-1}{k}} \left[\frac{f_s}{f_s - f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} - \alpha^{\frac{1}{k}} \right] \frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}} \frac{1}{\alpha^{\frac{1}{k}}} \right]^{-1}, \quad (36)$$

which must be satisfied if $\varphi_s^{cws} > \varphi_s^{no}$ (the only condition under which this scenario is plausible). Hence, if $\varphi_s^{cws} > \varphi_s^{no}$, revenues rise at firms that switch from the business-as-usual technology to the state-of-the-art technology

If $\varphi_s^{cws} > \varphi_s^{no}$, then regulation causes some facilities to downgrade from state-of-the-art to the retrofitted technology. For these firms

$$\frac{r_r^{cws}(\varphi)}{r_b^n(\varphi)} = \left[\frac{\varphi_\epsilon^n [1 + \tau\kappa]}{\varphi_\epsilon^{cws}} \alpha \right]^{\sigma-1}, \quad (37)$$

which is less than one if and only if

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{[1 + \tau\kappa]^{\frac{1}{k}} - \alpha^{\frac{1}{k}}}{[1 + \tau\kappa]^{\frac{1}{k}}} \right] \left[\left[\frac{[[1 + \tau\kappa]^{\sigma-1} - 1]^{\frac{k}{\sigma-1}}}{[1 + \tau\kappa]^{\frac{1}{k}}} \right] \left[\frac{1}{f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} \alpha^{\frac{1}{k}} + \left[[1 + \tau\kappa]^{\frac{k^2-1}{k}} \left[\frac{f_s}{f_s - f_r} \right]^{\frac{k-\sigma+1}{\sigma-1}} - \frac{1}{\alpha^{\frac{1}{k}}} \right] \frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}} \alpha^{\frac{1}{k}} \right]^{-1}. \quad (38)$$

Notice that this cut-off value for f is lower than that required to ensure revenues for retrofitters falls. As such, there is a range of values for f for which facilities that retrofit business-as-usual technology experience an increase in revenue while those that switch from state-of-the-art technology to the retrofitted technology experience a reduction in revenue. Note also that imposing $\alpha > [1 + \tau\kappa]$ is sufficient to guarantee $r_r^{cws}(\varphi)/r_s^n(\varphi) < 1$.

The above suggests that CWS style regulation has heterogeneous effects on firm revenue depending on firm productivity. We now turn to consider the effects on revenues for the average plant. Average revenues under regime *no*

APPENDIX

can be written as

$$\bar{r}^{no} = \int_{\varphi_\epsilon^{no}}^{\varphi_s^{no}} r_b^{no}(\varphi)g(\varphi)d\varphi + \int_{\varphi_s^{no}} r_s^{no}(\varphi)g(\varphi)d\varphi, \quad (39)$$

where $r_b^{no}(\varphi)$ is the revenue for firm using the business-as-usual technology and $r_s^{no}(\varphi)$ is the revenue for a firm using the state-of-the-art technology. Substituting for $r_b^{no}(\varphi)$ and $r_s^{no}(\varphi)$ and rearranging yields

$$\bar{r}^{no} = \sigma f \frac{k}{k - \sigma + 1} \frac{\Psi_1}{[\varphi_\epsilon^{no}]^k}. \quad (40)$$

Similarly, average revenues under regime *cws* can be written as

$$\bar{r}^{cws} = \sigma f \frac{k}{k - \sigma + 1} \frac{\Psi_2}{[\varphi_\epsilon^{cws}]^k}. \quad (41)$$

Taking the ratio of revenues under the *cws* regime to those under regime *no*, and using the definitions for the exit cut-offs gives

$$\frac{\bar{r}^{cws}}{\bar{r}^{no}} = \left[\frac{\Psi_1}{\Psi_2} \right]^{k^2-1}. \quad (42)$$

This expression is less than one, as $\Psi_1 < \Psi_2$ and $k > \sigma - 1 > 1$.

A.3.4 Emission Intensity

Next, we examine the effects of environmental regulation on the emission intensity of firm production. The effect of regulation on firm emission intensity is determined by the adoption of new technology. Given our setup, the emission intensity of facilities that use the business-as-usual technology in both regimes is unaffected. Facilities that change from business-as-usual to either the retrofitted or state-of-the-art technology experience a decline in emission intensity, given by $\frac{e_r^{cws}(\varphi)}{e_b^{no}(\varphi)} = \frac{e_s^{cws}(\varphi)}{e_b^{no}(\varphi)} = \frac{1}{\gamma} < 1$. Lastly, there is no change in the emission intensity of facilities that downgrade from state-of-the-art to retrofitted technology, or use state-of-the-art technology in both regimes. Altogether, this implies that the emission intensity of the average plant falls in response to environmental regulation.

APPENDIX

A.3.5 Changing The Fixed Costs of Retrofitting

As a last step, we examine how the magnitude of the fixed cost of retrofitting, f_r , affects firm decisions to retrofit the baseline technology. Lowering f_r lowers the retrofitting cut-off productivity level. As a result, under the CWS regime, a lower value of 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

$$\begin{aligned} \frac{\partial \varphi_r^{cws}}{\partial f_r} &= \left[\frac{1}{1 + \tau\kappa} \right] \left[\frac{f_r}{\Delta_2 f} \right]^{\frac{1}{\sigma-1}} \left[\frac{\partial \varphi_\epsilon^{cws}}{\partial f_r} + \left[\frac{1}{\sigma-1} \right] \frac{\varphi_\epsilon^{cws}}{f_r} \right] \\ &= \left[\frac{1}{1 + \tau\kappa} \right] \left[\frac{f_r}{\Delta_2 f} \right]^{\frac{1}{\sigma-1}} \left[\frac{1}{k[\varphi_\epsilon^{cws}]^k} \frac{\partial [\varphi_\epsilon^{cws}]^k}{\partial f_r} + \left[\frac{1}{\sigma-1} \right] \frac{\varphi_\epsilon^{cws}}{f_r} \right]. \end{aligned} \quad (43)$$

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

$$\frac{\partial [\varphi_\epsilon^{cws}]^k}{\partial f_r} > - \left[\frac{k}{\sigma-1} \right] \frac{[\varphi_\epsilon^{cws}]^k}{f_r}, \quad (44)$$

where $\frac{\partial [\varphi_\epsilon^{cws}]^k}{\partial f_r} = \left[\frac{f}{\delta f} \right] [1 + \tau\kappa]^k f^{\frac{k-\sigma+1}{\sigma-1}} \left[\left[\frac{\Delta_1}{f_s - f_r} \right]^{\frac{k}{\sigma-1}} - \left[\frac{\Delta_2}{f_r} \right]^{\frac{k}{\sigma-1}} \right]$. With some algebra, one can show that Equation (44) reduces to

$$\begin{aligned} \left[\frac{k}{k - \sigma - 1} \right] \frac{1}{f_r} \left[\frac{1}{f} \right]^{\frac{k-\sigma+1}{\sigma-1}} \left[\frac{1}{1 + \tau\kappa} \right]^{\sigma-1} + \left[1 + \left[\frac{k}{k - \sigma - 1} \right] \frac{1}{f_r} \frac{1}{f_s - f_r} \right] \left[\frac{\Delta_1}{f_s - f_r} \right]^{\frac{k}{\sigma-1}} > \\ - \left[\frac{\sigma - 1}{k - \sigma + 1} \right] \left[\frac{\Delta_2}{f_r} \right]^{\frac{k}{\sigma-1}}, \end{aligned} \quad (45)$$

which is always satisfied.

In addition, lowering f_r lowers the exit cut-off under the CWS regime if f_s

APPENDIX

isn't too large. Differentiating φ_ϵ^{cws} with respect to f_r gives

$$\frac{\partial \varphi_\epsilon^{cws}}{\partial f_r} = \left[\frac{\sigma - 1}{k - \sigma + 1} \right]^{\frac{1}{k}} \left[\frac{f}{\delta f_\epsilon} \right]^{\frac{1}{k}} [\Lambda^{cws}]^{\frac{1-k}{k}} [1 + \tau\kappa]^k \left[\frac{k - \sigma + 1}{\sigma - 1} \right] f^{\frac{k-\sigma+1}{\sigma-1}} \left[\Delta_2^{\frac{k-\sigma+1}{\sigma-1}} \left[\frac{1}{f_s - f_r} \right]^{\frac{k-2[\sigma-1]}{\sigma-1}} - \Delta_1^{\frac{k-\sigma+1}{\sigma-1}} \left[\frac{1}{f_r} \right]^{\frac{k-2(\sigma-1)}{\sigma-1}} \right], \quad (46)$$

which is greater than zero if and only if $f_s < \left[1 + \left[\frac{\Delta_1}{\Delta_2} \right]^{\frac{k-[\sigma-1]}{k-2[\sigma-1]}} \right] f_r$. 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 $f_s > \left[\frac{\Delta_1 + \Delta_2}{\Delta_2} \right] f_r$ 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, 2016b). 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 \$12,000,000.² 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.³

For each pollutant, plants are required to report their releases by medium (to air, water, and land), quantities sent for disposal and recycling, and methods used to compute releases. 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 geographic coordinates.

The ASM was used as Statistics Canada’s manufacturing census until 2012,

²For details, see sections 272 and 273 of the Canadian Environmental Protection Act.

³These 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>

APPENDIX

and provides longitudinal information for the majority of manufacturing plants in Canada.⁴ 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 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. Instead, 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 effects of the CWS. This issue arises if match probability is correlated

⁴The ASM was discontinued in 2012 and was replaced with a repeated cross-section survey.

APPENDIX

with treatment, meaning that the match is correlated with whether or not the plant is regulated under the CWS. If this is the case, our estimates of the effects of the CWS will be biased by the fact that matched plants are more likely to be regulated. Indeed, this is a potentially important issue in our context; as we show in the main body of the paper, the effects of the CWS vary across plants of different productivities. 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, it appears match probability is potentially correlated with regulation. To correct for any resulting bias, we use a simple weighting strategy.

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

$$\beta = Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2}, \quad (47)$$

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

The effect of regulation in the matched sample is given by

$$\begin{aligned} \beta^{match} &= Pr(g_1|match)\beta^{g_1} + Pr(g_2|match)\beta^{g_2} \\ &= \frac{Pr(match|g_1)Pr(g_1)}{Pr(match)}\beta^{g_1} + \frac{Pr(match|g_2)Pr(g_2)}{Pr(match)}\beta^{g_2}, \end{aligned} \quad (48)$$

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

APPENDIX

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 effect of regulation. The weighted effect from the matched data would be

$$\begin{aligned}\beta^{match,weighted} &= \omega_{g1}Pr(g_1|match)\beta^{g1} + \omega_{g2}Pr(g_2|match)\beta^{g2} \\ &= Pr(g_1)\beta^{g1} + Pr(g_2)\beta^{g2},\end{aligned}\tag{49}$$

which is the true effect of regulation, 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 pollutant sample in our analysis.

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-

APPENDIX

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.

ASM matched data than in the universe of polluters. Weighting reduces considerably reduces this over-estimate, 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

As we discuss in the main text, we determine regulatory status in each CMA using data from Canada’s National Air Pollution Surveillance Program (NAPS). The NAPS is a network of 286 air quality monitoring stations located across Canada, and is Canada’s main source for ambient 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 98th percentile of each CMA’s 24-hour concentration in a given year.⁵ For O₃, we construct the 4th highest 8-hour concentration

⁵The 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.

APPENDIX

reported in a CMA in a given year.⁶ 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 O₃ measurements (Canadian Council of Ministers of the Environment, 2002, p. 12).

⁶For 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.

B.2 Robustness Checks

As we described in the main text, we engage in several exercises to examine the robustness of our main findings. The results of these exercises are described below. 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.⁷

We begin by examining whether our results are capturing the effects of a non-linear relationship between CMA air quality and the production choices of plants therein.⁸ 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 have 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 a bin according to the relevant CMA's air quality in that year, and then estimate a version of our main specification 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.

⁷These regression results are available upon request.

⁸Such 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.

APPENDIX

This specification is given by:

$$\begin{aligned}
 Y_{pict} = & \sum_b \beta_{PM}^b [K_i \times I(\underline{A}_b^{PM} \leq a_{ct}^{PM} < \overline{A}_b^{PM})] \\
 & + \sum_b \beta_{O_3}^b [K_i \times I(\underline{A}_b^{O_3} \leq a_{ct}^{O_3} < \overline{A}_b^{O_3})] + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict},
 \end{aligned} \tag{50}$$

where b indexes air quality bin numbers, K_i selects all industries targeted by the CWS, a_{ct}^j is the air quality measured in CMA c for pollutant j in year t , \underline{A}_b^j is the air quality lower bound for bin b for pollutant j , \overline{A}_b^j is the air quality upper bound for bin b for pollutant j , and $I(\underline{A}_b^j \leq a_{ct}^j < \overline{A}_b^j)$ is an indicator for all CMA-years with air quality that corresponds to bin b for pollutant j .⁹ The coefficient β_j^b gives the effects of standard j in air quality bin b .

In estimating Equation (50), we omit the “cleanest” air quality bin for each standard. For the PM_{2.5} standard, we break the air quality distribution into seven equal-sized bins from 18 to 36 $\mu\text{g}/\text{m}^3$. For the O₃ standard, we break the air quality distribution into six equal-sized bins from 57 to 77 ppb.¹⁰

The results of the estimating of Equation (50) using the full sample of pollutants from the NPRI are displayed in Figure 1. Figure 1 displays the coefficients and confidence intervals for PM_{2.5} and NO_x emissions, respectively. Only the coefficients for the PM_{2.5} standard are shown for PM_{2.5} emissions, and the O₃ standard for NO_x emissions. Each figure also displays the fraction of observations in each bin that are treated over the sample, to show that there are treated plants 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.

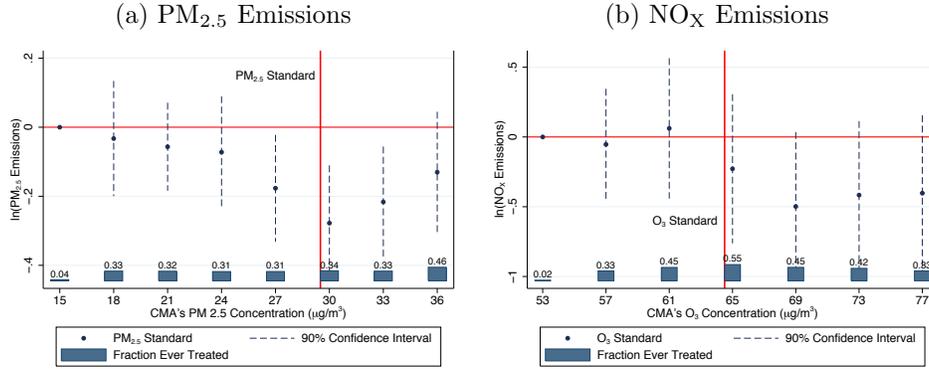
The results presented in Figure 1 show that a break that occurs just below

⁹For example, suppose PM_{2.5} air quality ranged from 20 to 40 $\mu\text{g}/\text{m}^3$, and we split this into two equal-sized bins. The upper and lower bounds for bin one would be $\overline{A}_1^{PM} = 30$ and $\underline{A}_1^{PM} = 20$, respectively. The upper and lower bounds for bin two would be $\overline{A}_2^{PM} = 40$ and $\underline{A}_2^{PM} = 30$, respectively. Bin one would select all plants in CMAs with air quality below 30 $\mu\text{g}/\text{m}^3$, and bin two would select all plants in CMAs with air quality above 30 $\mu\text{g}/\text{m}^3$.

¹⁰For the PM_{2.5} regulation we include all CMA-years with air quality above 36 $\mu\text{g}/\text{m}^3$ in the top bin. For the O₃ regulation we include all CMA-years with air quality above 77 ppb in the top bin.

APPENDIX

Figure 1: Mean Pollution Concentrations by Year



Notes: Figure displays estimates from a flexible DDD estimation of the $\text{PM}_{2.5}$ standard’s effect on $\text{PM}_{2.5}$ emissions and the O_3 standard’s effects on NO_x emissions allowing the effects of regulation to vary by CMA air quality. In each panel, diamonds reflect the point estimates for each CMA air quality bin, while the dashed line displays the associated 90% confidence interval. These coefficients are measured relative to the excluded group (air quality below 18 $\mu\text{g}/\text{m}^3$ for $\text{PM}_{2.5}$ 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.

the $\text{PM}_{2.5}$ standard’s threshold for $\text{PM}_{2.5}$ emissions and at the precise level of the O_3 standard’s threshold for NO_x emissions. This suggests that there are no significant differences in the trends of 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 the main body of the paper 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.

We next 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 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

APPENDIX

future.

This is particularly demanding in this setting because the majority of treated CMAs begin the sample period under treatment, particularly for the O₃ standard. 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 still treated). This regression is estimated by fitting the following generalized triple-difference estimator to the data

$$Y_{pict} = \sum_{k=-3} \beta_{PM}^k T_{ick}^{PM} + \beta_{O_3} T_{ict}^{O_3} + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict}, \quad (51)$$

where T_{ick}^{PM} is an indicator for the years before ($k < 0$) or after ($k \geq 0$) a plant is treated for standard j , and $T_{ict}^{O_3}$ captures the average effect of the O₃ standard. We exclude the year prior to treatment for the PM_{2.5} standard ($k = -1$), so the coefficients of interest (β_{PM}^k) report the semi-elasticity of treatment k years before or after treatment relative to the year before treatment. In other words, β_{PM}^k is the triple-difference coefficients relative to the year before a plant is first treated by the standard.¹¹

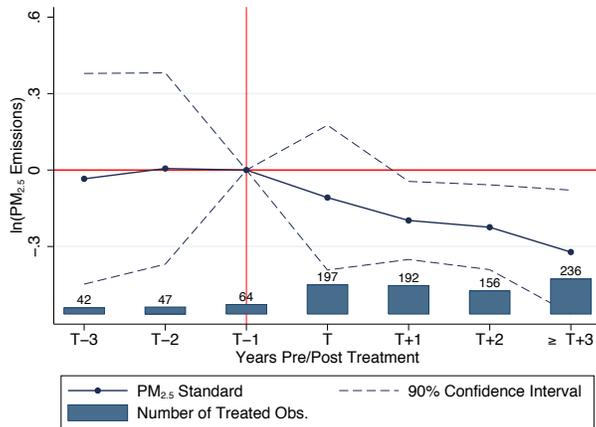
We estimate Equation (51) from three periods before a plant is treated onward. Separate coefficients are estimated up to three periods post treatment, and all periods greater than three years after treatment are pooled. We drop all observations that occur prior to three periods before a plant is treated. 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

¹¹Note that in our main specification 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.

APPENDIX

Figure 2: 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.

shown in Figure 2. The dependent variable is the natural logarithm of PM_{2.5} emissions and standard errors are clustered by CMA-industry.

Figure 2 shows strong evidence that there was no significant difference in pre-regulation trends for our treatment and control groups for the PM standard, with the pre-regulation coefficients hovering tightly around zero. In addition, there was a clear break in PM_{2.5} emissions starting in the year of regulation and persisting following treatment.

Having established that our main results are not due to differential trends across groups, we now examine the possibility that plants are systematically making pre-emptive production changes to avoid regulation, as this would lead to biased estimates of the effects of the CWS.

We first examine when plants are regulated by the CWS. If the majority of plants are regulated near the end of the CWS period, then there is a strong possibility that plants may have been able to respond pre-emptively in anticipation of future regulation.

APPENDIX

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 O₃ standard.

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 O₃ 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 the O₃ standard reduced emissions of NO_x by 56%.¹² 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.

These results suggest that the baseline estimates presented in our main

¹²Note that we estimate all robustness checks using the publicly available NPRI data, rather than the matched data, so as to reduce the number of estimates requiring vetting by Statistics Canada. As a result, the number of observations differ between the robustness checks and the main analysis. The results using the matched sample are very similar, and can be provided upon request.

APPENDIX

Table A3: CWS Effect on Emissions for Initial Treatment Cohort

	(1)	(2)
	PM _{2.5}	NO _x
PM _{2.5} Standard	-0.169 (0.087)	0.0132 (0.072)
O ₃ Standard	-0.059 (0.082)	-0.560 (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 O₃ standard. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry.

analysis are not driven by preemptive changes to avoid regulation. 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 drop plants that emit a large fraction of their CMA’s emissions. Dropping large plants lowers the potential for bias by removing plants who 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.¹³ We report the results for emissions in

¹³For reference, the average plant fraction of city emissions is: 7% for PM_{2.5} and 10% for NO_x.

APPENDIX

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.

Our estimates of the effects of the CWS by plant productivity level 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 possibility of differential trends across large and small emitters, we estimate a version of our main specification that allows for separate CMA-year fixed effects for relatively large and relatively small 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,

APPENDIX

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 (0.0651)	-0.205 (0.0698)	-0.203 (0.0750)	-0.133 (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)
O ₃ Standard	-0.273 (0.115)	-0.205 (0.129)	-0.219 (0.134)	-0.0696 (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 O₃ standard. The effect of the PM_{2.5} standard is shown for PM emitters, and the O₃ standard is shown for O₃ NO_x emitters. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry.

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 our estimates of the effects of the CWS. PM_{2.5} regulation significantly reduced PM emissions from affected plants, and O₃ regulation significantly reduced NO_x emissions from affected plants. Consequently, we conclude our results are unlikely to be reflective of differential trends across large and small emitters.

APPENDIX

Table A5: CWS Effect on Emissions with Emitter Size Trends

	(1) PM _{2.5}	(2) NO _x
PM _{2.5} Standard	-0.128 (0.0590)	0.0573 (0.0841)
O ₃ Standard	-0.0644 (0.0776)	-0.277 (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 O₃ standard. All regressions include plant, industry-year and emitter size-by-CMA-by-year fixed effects. Standard errors are clustered by CMA-industry.

Finally, we turn to address the possibility that our results are capturing the effects of firm ownership. While we treat each plant in our analysis as an independent agent, 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, leading to a violation of the Stable Unit Treatment Value Assumption (SUTVA) that is implicit in our analysis. 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 multi- and single-plant firms.^{14,15}

We use the parent company name information reported in the NPRI to

¹⁴An alternative approach is to simply drop all multi-plant firms. Doing this produces similar results.

¹⁵Our 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.

APPENDIX

identify multi-plant firms. This information is entered as a text string, which is 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.¹⁶ 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 $\text{PM}_{2.5}$ 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). 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 $\text{PM}_{2.5}$ emitters we estimate the $\text{PM}_{2.5}$ standard's effect on plant emissions, and for NO_x emitters we estimate the O_3 standard's effect on plant emissions. These results are reported in Table A6. As can be seen from the table, in all specifications there is no significant difference in the estimated effect of the CWS for plants owned by single-plant firms and those owned by multi-plant firms. As a result, it appears the potential failure of SUTVA through the common-ownership channel does not appear to be an issue for our analysis.

¹⁶For details on the Levenshtein Distance measure, see Yujian and Bo (2007).

APPENDIX

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

	<u>Panel A: PM_{2.5}</u>			
	<u>Same Pollutant</u>		<u>Any Pollutant</u>	
	Coarse	Fine	Coarse	Fine
	Matching	Matching	Matching	Matching
	(1)	(2)	(3)	(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
	<u>Panel B: NO_x</u>			
	<u>Same Pollutant</u>		<u>Any Pollutant</u>	
	Coarse	Fine	Coarse	Fine
	Matching	Matching	Matching	Matching
	(5)	(6)	(7)	(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 O₃ standard. The effect of the PM_{2.5} standard is shown for PM emitters, and the O₃ standard is shown for O₃ NO_x emitters. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry.

B.3 Decomposition

Here we derive the empirical analogue of the industry decomposition given by Equation (2) in the main text that we use to construct estimates of the process, reallocation and selection effects.

Letting $E_{it} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn$ denote the emission intensity of industry i at time t , where $e_{it}(n)$ denotes plant emission intensity, $\lambda_{it}(n)$ is a plant's share of industry output, and n_{it} is the marginal surviving plant, we can write the change in an industry's pollution intensity between $t - 1$ and t as:

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

This can be rewritten as:

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

With some algebra, this reduces to

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

Dividing by E_{it-1} yields our empirical decomposition, given by equation (5) in the main text.

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

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

APPENDIX

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 yields:

$$\frac{X_{it}}{X_{it-1}} = 1 - s_{xt-1}^{Exit} + \beta_x s_{xt-1}^{Treat}. \quad (52)$$

With some algebra it can be shown that $\hat{\lambda}_{it}(n)$ is as in the text.

Appendix C Other Margins of Plant Adjustment

As we discussed in the main text, there are other potential explanations for our findings, including changes in primary inputs, intermediate inputs, and productivity. Here, we examine the support for these alternative explanations.

In our main analysis, we hypothesize that the differences in outcomes for $PM_{2.5}$ and NO_X is due to differences in the fixed cost of process changes, which affects a plant's willingness to adopt cleaner production processes. An 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 assess the above hypothesis. If this hypothesis were true, then the CWS should have caused an increase in spending on inputs for NO_X emitters.¹⁷ 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.¹⁸ 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 the main body of the paper. Examining the effect of the CWS on plant productivity allows us to test this hypothesis.

¹⁷Here we have assumed plants would use the cheapest input in the absence of regulation.

¹⁸For a recent review of this literature, see Ambec et al. (2013)

APPENDIX

We examine these alternative hypotheses using data on the total number of plant employees¹⁹, 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 A7. Panel A shows estimates of our main specification for PM_{2.5} emitters and Panel B shows estimates for NO_x emitters. In each panel, we report estimates 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 O₃ 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 A8.²⁰ 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.

¹⁹Although 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).

²⁰The effects on R&D are omitted, but are available upon request.

APPENDIX

Table A7: Other Margins of Plant Adjustment

	Panel A: PM _{2.5}				
	<u>Prim. Inputs</u>	<u>Inter. Inputs</u>		<u>Productivity</u>	
	(1)	(2)	(3)	(4)	(5)
	Employment	Materials	Energy	VA/Worker	Pr(R&D)
PM 2.5 Std.	-0.040 (0.064)	-0.119 (0.064)	-0.086 (0.056)	-0.098 (0.073)	0.033 (0.040)
O ₃ Std.	0.071 (0.068)	-0.008 (0.071)	0.224 (0.108)	0.039 (0.060)	-0.086 (0.060)
R^2	0.188	0.218	0.151	0.185	0.155
N	6501	6499	6478	6501	6501

	Panel B: NO _x				
	<u>Prim. Inputs</u>	<u>Inter. Inputs</u>		<u>Productivity</u>	
	(1)	(2)	(3)	(4)	(5)
	Employment	Materials	Energy	VA/Worker	Pr(R&D)
PM 2.5 Std.	0.003 (0.069)	0.039 (0.077)	-0.094 (0.093)	-0.231 (0.085)	0.061 (0.060)
O ₃ Std.	-0.064 (0.157)	-0.069 (0.154)	0.085 (0.264)	-0.062 (0.117)	-0.143 (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 O₃ 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.

As the estimates reported in Table A7 and Table A8 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

Table A8: 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 (0.113)	-0.194 (0.094)	-0.125 (0.074)	-0.247 (0.119)	0.141 (0.097)	0.165 (0.099)	-0.080 (0.122)	-0.418 (0.124)
x Q2	-0.093 (0.055)	-0.044 (0.070)	-0.051 (0.073)	0.058 (0.061)	-0.031 (0.079)	-0.007 (0.094)	-0.189 (0.146)	-0.188 (0.092)
x Q3	-0.065 (0.072)	-0.049 (0.095)	-0.041 (0.098)	0.027 (0.089)	-0.116 (0.100)	-0.049 (0.122)	-0.044 (0.102)	-0.076 (0.111)
O ₃ Std.								
x Q1	0.131 (0.086)	0.004 (0.092)	0.311 (0.128)	-0.058 (0.078)	-0.079 (0.200)	-0.004 (0.194)	-0.181 (0.304)	-0.177 (0.156)
x Q2	0.024 (0.128)	-0.031 (0.149)	0.252 (0.243)	0.014 (0.099)	-0.057 (0.163)	-0.108 (0.166)	0.190 (0.257)	0.011 (0.157)
x Q3	0.047 (0.076)	-0.016 (0.082)	0.109 (0.142)	0.136 (0.092)	-0.010 (0.170)	-0.053 (0.159)	0.237 (0.279)	-0.085 (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 O₃ 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.

APPENDIX

levels), and reduced labor productivity (also not significant at conventional levels). PM_{2.5} regulation also caused a significant reduction in labor productivity among 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 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 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 A7 and Table A8 also suggest O₃ 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 process changes.

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.
- 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.
- Environment and Climate Change Canada (2013). National Air Pollution Surveillance Program (NAPS). <http://www.ec.gc.ca/rnspa-naps/>.
- 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>.
- Environment and Climate Change Canada (2016b). Summary of National Pollutant Release Inventory Reporting requirements. <https://www.ec.gc.ca/inrp-npri/default.asp?lang=en&n=629573FE-1>.
- 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.
- 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.
- Yujian, L. and L. Bo (2007). A normalized Levenshtein distance metric. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(6), 1091–1095.