Essays on the Economics of Environmental Regulation

by

Nouri Najjar

B.Sc., University of Victoria, 2010

M.A., University of British Columbia, 2011

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Economics)

The University of British Columbia

(Vancouver)

August 2018

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

Essays on the Economics of Environmental Regulation

submitted by Nouri Najjar in partial fulfillment of the requirements for

the degree of Doctor of Philosophy

in Economics

Examining Committee:

Prof. Brian Copeland, Economics, University of British Columbia
Co-supervisor

Prof. Carol McAusland, Food and Resource Economics, University of British Columbia
Co-supervisor

Prof. Thomas Lemieux, Economics, University of British Columbia
Supervisory Committee Member

Prof. James Brander, Sauder School of Business, University of British Columbia
University Examiner

Prof. Nicole Fortin, Economics, University of British Columbia
University Examiner

Additional Supervisory Committee Members:

Supervisory Committee Member

Supervisory Committee Member

Abstract

Environmental regulations targeting producers are in place around much of the world. Yet, there is limited evidence of how firms are affected by these policies. This thesis provides new empirical and theoretical evidence on the effects of environmental regulation on producers.

The first two chapters of this thesis explore a trend underway in much of the industrialized world: pollution from manufacturing has been falling despite increased output. In the first chapter, we develop a theoretical model to show the channels through which regulation can contribute to an industry's "clean-up". This model highlights the role that fixed costs producers must pay to adopt cleaner production processes play in dictating these channels. We show that if these fixed costs are relatively low, the adoption of cleaner processes will be the primary regulatory channel of an industry's clean-up. However, if these fixed costs are relatively high, then plant exit and reductions in output from regulated plants will be the primary channels.

The second chapter provides the first estimates of the regulatory channels of the manufacturing clean-up. We estimate the share of the Canadian manufacturing clean-up explained by the adoption of cleaner production processes, the reallocation of output across producers, and producer entry and exit. To do this, we examine a major revision to Canadian environmental policy using a novel, confidential dataset containing information on the production decisions and pollution emissions of Canadian manufacturing plants. We find regulation explains, at most, 61% of the Canadian clean-up, but the underlying channels differ strikingly across pollutants.

A concern in debates over environmental regulation is a potential loss of international competitiveness among domestic producers. Despite its pervasiveness in policy discussions, evidence of these losses remains scarce. The third chapter of this thesis provides the first plant-level estimates of the effect of air pollution regulation on exporting. We study the effects of the Canada-Wide Standards for Particulate Matter and Ozone on the decision to export and export volumes of Canadian manufacturing facilities. We find evidence that environmental regulation caused relatively low-productivity exporters to leave the export market, and reduced the amount surviving exporters sold abroad.

Lay Summary

This thesis is intended to show some important dimensions through which environmental regulation affects industry, and the firms and facilities therein. The first two chapters consider the causes of an important trend in Canada: the manufacturing industry has become cleaner, in terms of air pollution, in recent decades. The first chapter develops an economic model of how environmental regulation affects firms, and how these effects can cause an industry to become cleaner. In the second chapter, we then ask how a major increase in the stringency of environmental regulation has contributed to this trend. We find regulation had considerable effects on how manufacturing firms operate, which contributed to the manufacturing industry's clean-up. The third chapter documents an additional margin through which environmental regulation affects firms: exports. We find regulation reduced the amount of exports from regulated manufacturing facilities, and reduces the likelihood a relatively low-productivity facility chooses to export.

Preface

All chapters in this thesis were coauthored with Dr. Jevan Cherniwchan (University of Alberta). I was the lead author of both of the first two chapters. In the first chapter, I developed the economic model, which is the core contribution of that chapter. In the second chapter, I discovered the policy experiment that we study, which forms the basis of the empirical strategy we use to identify the effect of regulation. Both Jevan and I evenly contributed to the third chapter.

In addition, I was a primary contributor to all additional stages of the research, including developing the research questions, preparing data, carrying out estimation, and organizing and presenting results.

Table of Contents

List of Tables

List of Figures

Acknowledgments

I wish to thank Brian Copeland and Carol McAusland for their guidance and mentorship throughout my time at UBC. I am deeply grateful for the many hours they contributed to my development as an economist, and will be forever indebted to them. Thank you to my other committee members, Thomas Lemieux and David Green, for invaluable advice, critiques, and contributions. I would also like to thank my coauthor, Jevan Cherniwchan, for the depth of his commitment to our collaboration. To my colleagues at UBC – Alastair, Brad, Iain, Tom, Jacob, Jose, Joao, Timea, Hugo, Gaelle, Oscar, and Nathan, to name just a few – thank you for your friendship and insights over the years.

For helpful comments and suggestions, Jevan Cherniwchan and I wish to thank seminar participants at UBC, Western University, the University of Carlos the III, the University of Mannheim, the CREE annual meeting, the EEPRN symposium, the CEA annual conference, the Young Economists Symposium at Yale, the CU Environmental and Resource Economics Workshop, and the Pacific Northwest Labor Day. We are especially thankful to Arik Levinson and Joe Shapiro for insightful comments on an earlier version of the second chapter. We are indebted to Michael Willox and Danny Leung at Statistics Canada's Economic Analysis Division for their efforts to make this project possible. We thank members of Environment and Climate Change Canada's (ECCC) Legislative and Regulatory Affairs Division for helpful discussion on the design and implementation of Canadian air quality regulation. Funding from the EEPRN, the Alberta School of Business Nova Faculty Fellowship, and the Centre for Applied Business Research in Energy and the Environment at the University of Alberta is gratefully acknowledged. The data used in this thesis was created through a collaboration between ECCC, Statistics Canada, and the Productivity Partnership, which is supported by the Social Sciences and Humanities Research Council of Canada.

I thank my parents for their unending support, advice, and love throughout my life. To my sisters and the rest of my family, thank you as well for all your support.

Lastly, I owe the deepest debt of gratitude to Rebecca. Thank you for years of discussions, all manner of support, and for sharing your intelligence and drive with me each day.

Introduction

Air pollution has numerous negative effects on society. To name a few, certain pollutants have been shown to increase rates of mortality and morbidity, particularly among young children and adults with respiratory illness (Chay and Greenstone, 2003; Currie and Neidell, 2005; Schlenker and Walker, 2016). Moreover, exposure to air pollution can cause a reduction in worker productivity and labour supply (Chang et al., 2016; Graff Zivin and Neidell, 2012; Hanna and Oliva, 2015). As a consequence of these negative health and economic effects, in recent decades governments around much of the world have enacted increasingly stringent air pollution regulation. Despite their prevalence, relatively little is known about the way in which firms respond to and are affected by these policies.

This thesis provides new evidence of the effects of air pollution regulation on manufacturing facilities. In the first two chapters, we determine the channels through which environmental regulation can cause an industry to become less emission intensive (that is, to emit less pollution per unit of output produced). In the first of these chapters, we present a theoretical framework to capture these channels. We show regulation can cause an industry to "clean up" through three channels: a selection channel, a reallocation channel, and a process channel. The selection channel captures changes in an industry's emission intensity that occurs through facility entry and exit. If, for example, regulation causes relatively dirty facilities to shut down, then this selection channel will cause the industry to clean up. The reallocation channel captures changes in an industry's emission intensity that occurs due to reductions in output at surviving regulated facilities. If these facilities are relatively dirty, then this channel will contribute to the industry's clean-up. Lastly, the process channel occurs if regulation causes surviving facilities to adopt cleaner production processes.

In the first chapter we also show that the relative magnitude of each of the three channels will depend on the size of the fixed cost plants need to pay to adopt cleaner processes. If these fixed process costs are very small, then the process channel will drive an industry's clean-up. Otherwise, the selection and reallocation channels will play an important role in an industry's clean-up.

In the second chapter, we use a novel dataset to estimate the regulatory channels involved in the clean up of the Canadian manufacturing sector. To determine causal estimates of regulation's effect on manufacturing plants, we exploit a change in Canadian federal environmental policy that has never before been studied. This policy implemented regional air quality standards for two pollutants in every major town and city of the country. In addition, only facilities in a select group of industries were regulated. We exploit this variation in regulatory stringency across cities, industries, and time using a triple difference research design. In essence, this research design compares changes in outcomes for plants in targeted industries and regions violating an air quality standard to changes in outcomes at other plants.

We show robust evidence that this policy contributed to the Canadian manufacturing clean-up through each of the selection, reallocation, and process channels. Summing across all three channels, we find this policy is responsible for 60% of the reduction in manufacturing nitrogen oxide emission intensity from 2004 to 2010. The policy is also responsible for 20% of the reduction in particulate matter emission intensity over this period. In addition, we show that the relative magnitude of these channels varied across the two pollutants regulated by this policy. We argue these results are consistent with the theory presented in the first chapter, and in particular with our claim that the size of fixed process costs drives the channels of the clean-up. Lastly, we show additional evidence consistent with this hypothesis and theory.

In the third chapter, we ask whether environmental regulation affects a firm's incentive to export. Debates over environmental regulation often center on their potential negative effect on the competitiveness of domestic firms in international markets (see, e.g, Levinson and Taylor (2008)).[1] Yet, to date, there have been no micro-level estimates of the effect of regulation on export behaviour. We exploit the same policy change studied in Chapter 2 to estimate the effect of regulation on a facility's decision to enter (or exit) export markets, and their export volumes conditional on exporting. Our results suggest regulation negatively affected the international competitiveness of Canadian manufacturing facilities. We find evidence this policy caused a negative effect on the international competitiveness of affected manufacturing plants. For the average affected plant, this policy caused a 22% reduction in export volumes. Moreover, regulation caused a 10% increase in the probability of exiting the export market for relatively low-productivity exporters.

[1]The concern in these debates is that regulation in one country may increase the cost of production for domestic producers relative to their unregulated counterparts in foreign countries.

Chapter 1

A Theory of Industry Clean-Up

1.1 Introduction

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

Recent evidence suggests environmental regulation has played a large role in this "clean-up" of manufacturing (Shapiro and Walker, 2015). In addition, this clean-up appears to be due to reductions in the emission intensity of individual industries (Brunel, 2016; Levinson, 2015), rather than changes in the composition of the manufacturing sector. Yet, at present, little theoretical work has been done to directly assess how regulations change an industry's emission intensity. In this chapter we present a theoretical model to capture the channels through which a particular type of environmental regulation can cause an industry to clean up – what we call a two-part regulatory structure. In this type of regulation, firms must either make technological changes to meet industry best practices, or are subject to a regulatory penalty. The theoretical work has yet to assess this type of policy, despite its common use in regulating air pollutants. For example, the National Ambient Air Quality Standards used as part of the US Clean Air Act require regulated facilities to adopt state-of-the-art abatement technology, and fines those that fail to do so (Greenstone, 2002). Similarly, the Canada-Wide Standards for Particulate Matter and Ozone require regulated facilities to use clean production processes, and imposes production constraints on those that fail to do so (see Chapter 2 for details).

Before presenting our model, we start by following a similar approach to Cherniwchan et al. (2017) to show that a change in an industry's emission intensity can be decomposed into three channels: a selection channel, a reallocation channel, and a process channel. The selection channel reflects the exit of plants in response to regulation. The reallocation channel reflects the reduction in output at surviving regulated plants in response to regulation. Lastly, the process channel captures the adoption of cleaner production processes at surviving regulated plants.

We next present the theoretical model we use to show the mechanisms driving these channels. This model is based on a closed-economy variant of the Melitz (2003) model in which pollution from heterogeneous firms is regulated. This model has three key features. First, it allows for firm 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 firms to capture the fact that leading technologies are often used as a benchmark for the technical changes required under regulation. Third, it allows for differences across pollutants in the cost of adopting less-polluting production processes, which we call process costs. This feature is important because pollutants often feature different process costs.[1]

This chapter relates most closely to the recent body of theoretical work studying environmental regulations and pollution in the presence of firm heterogeneity. Our model makes three primary contributions to this literature. First, as we discussed above, we consider a regulation that features a two-part regulatory structure. This is in contrast to the existing literature, which focuses on uniform pollution taxes (Andersen, 2018; Cao et al., 2016; Forslid et al., 2014; Li and Shi, 2017; Li and Sun, 2015), pollution permit trading (Anoulies, 2017; Cui et al., 2015; Konishi and Tarui, 2015), or pollution intensity standards (Li and Shi, 2017; Li and Sun, 2015). We make this departure to capture a common feature of environmental policy that has not yet been highlighted in the theoretical literature.

Our second contribution is to make explicit the connection between environmental regulation and the channels through which an industry becomes less pollution intensive. While not explicitly stated as such, the majority of the existing literature has focused on the selection and

[1] For example, nitrogen oxide (NO_X) process costs are relatively low, while $PM_{2.5}$ process costs are typically relatively high (Canadian Council of Ministers of the Environment, 1998b; Environment Canada, 2002).

reallocation effects induced by policy (e.g. Konishi and Tarui (2015) or Anoulies (2017)). By allowing for process effects, our work is more closely related to that of Cao et al. (2016), who study the effects of a uniform pollution tax on investment in abatement technology. Unlike Cao et al., we rely on constant elasticity of substitution preferences, rather than quasi-linear preferences, and allow for a different form of technology adoption (discussed below).

Our third contribution is to focus on discrete technology choices, following an approach used by Bustos (2011), rather than consider continuous abatement investments. This is in contrast to two alternative approaches. The first is the canonical approach used in environmental economics in which facilities can make incremental process changes, which increase their variable production costs, but do not require additional fixed costs (see, e.g., Antweiler et al. (2001); Shapiro and Walker (2015)). The second is to allow for smooth changes in a fixed-cost technology (see, e.g., Cao et al. (2016)). While a number of papers also present models based on Bustos (2011) featuring heterogeneous polluting firms that make endogenous abatement decisions (Batrakova and Davies, 2012; Cui et al., 2012; Forslid et al., 2014), these studies focus on the effects of international trade. Instead, we show how environmental regulation influences firms to make discrete changes in technology. As we show, the fixed costs required to make these changes play an important role in dictating the channels involved in a clean-up, particularly if facilities can effectively avoid regulation by adopting process changes. In this case, the size of the process cost directly affects the relative magnitudes of the channels involved in an industry's clean-up.

The remainder of this chapter is as follows. In Section 1.2 we present the results of our decomposition exercise showing the potential channels of a clean-up. In Section 1.3, we present our theoretical model. In Section 1.4, we discuss the model's main implications for the channels of a clean-up. Lastly, a short conclusion summarizes.

1.2 Channels of the Clean-Up

Before presenting our theoretical model, we first use a decomposition exercise to derive the channels through which plant-level changes in response to environmental regulation can cause an industry to clean-up. For this exercise, we follow an approach similar to that introduced by Cherniwchan et al. (2017), which itself extends the decomposition presented by Levinson (2009).

To that end, let output and pollution from manufacturing industry i be given by X_i and Z_i, respectively. We define an industry's pollution intensity as the amount of pollution emitted per unit of output produced, and let this be 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 $\lambda_i(n) = x_i(n)/X_i$ 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.[2] In this case, the emission intensity of industry i can be expressed as a weighted average of the plant emission intensities

$$E_i = \int_0^{n_i} e_i(n)\lambda_i(n)dn. \tag{1.1}$$

Totally differentiating Equation (1.1) gives the change in emission intensity of any industry i as

$$dE_i = \int_0^{n_i} de_i(n)\lambda_i(n)dn + \int_0^{n_i} e_i(n)d\lambda_i(n)dn + e_i(n_i)\lambda_i(n_i)dn_i. \tag{1.2}$$

We call the first term on the right-hand side of equation (1.2) 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 (1.2), 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 (1.2) 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 a theoretical framework to capture the regulatory process, reallocation, and selection effects induced by a regulatory change.

[2] As in Cherniwchan et al. (2017), we assume plants are ranked in reverse order of productivity. Consequently, selection removes the least productive plants.

1.3 Theory of a Clean-Up

1.3.1 Model Setup

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

The representative consumer derives utility from the consumption of goods and disutility from aggregate pollution according to $U = [\int_0^M q(\omega)^\rho d\omega]^{1/\rho} - h(Z)$, where $q(\omega)$ denotes consumption of good ω, and M denotes the measure of varieties available in the economy. It is assumed consumers ignore pollution when making their consumption decisions. As a result, the demand for variety ω is given by $q(\omega) = IP^{\sigma-1}p(\omega)^{-\sigma}$, where I denotes consumer income, $P = [\int_0^M p(\omega)^{1-\sigma} d\omega]^{1/(1-\sigma)}$ is the economy's price index, and $\sigma = 1/[1-\rho] > 1$ is the elasticity of substitution between goods.

The supply side of the economy features monopolistic competition and free entry, meaning each firm in the economy produces a unique variety. To enter, firms pay a fixed entry cost f_ε, and upon entry, draw a productivity level φ from a common distribution $G(\varphi)$.[3] Based on the realization of φ, firms decide whether to exit or stay in the market, and conditional on staying, how much to produce and what technology to use in production.

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

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

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

state-of-the-art technology are given by $C_s = c_s^l(\varphi)x + f + f_s$, where $c_s^l(\varphi)$ is the marginal cost of producing x with the state-of-the-art technology in regime l.

We view state-of-the-art technology as reflecting a suite of industry-leading processes and technologies. These reduce labor costs and pollution intensity jointly for at least two potential reasons. First, some of these processes may make production cleaner by reducing the intensity of certain dirty inputs, such as fossil fuels.[4] Second, this serves as a simple method to reflect that some firms may be forward-looking, and believe that future regulation will be tightened. Consequently, when building a facility, they may choose to do so with clean processes in anticipation of future regulation, even if the precise nature of that regulation is unknown.[5]

Instead of becoming state-of-the-art, firms may retrofit their business-as-usual technology so that it has the same emission intensity as the state-of-the-art technology. As such, the emission intensity of a retrofitted plant (e_r) is also $\kappa/[\gamma\varphi]$, meaning the total level of pollution generated by production is $z_r(\varphi) = [\kappa x]/[\gamma\varphi]$. Retrofitting also requires firms to pay a fixed cost (f_r). However, retrofitting does not affect labor productivity, meaning it is less costly than adopting the state-of-the-art technology, so $f_r < f_s$. The total costs of production for a retrofitted plant are given by $C_r = c_r^l(\varphi)x + f + f_r$, where $c_r^l(\varphi)$ is the marginal cost of producing x in regulatory regime l with the retrofitted technology.

1.3.2 The No-Regulation Equilibrium

Our interest is in understanding the effects of environmental regulation that requires firms to use clean production processes, and penalizes those that fail to do so. We call this targeted regulation. We first consider a no regulation regime (labeled with superscript *no*) in which pollution is not regulated. This means labor costs are the only variable costs of production, so $c_b^{no}(\varphi) = c_r^{no}(\varphi) = 1/\varphi$ and $c_s^{no}(\varphi) = 1/[\alpha\varphi]$.

A firm that has drawn a productivity level maximizes profits by deciding whether to stay in the market, and if they stay, choosing how much to produce and what technology to use. Given the structure of consumer preferences, this implies that producing firms set prices at a constant mark-up over marginal costs. Hence, in the absence of regulation, firms that employ business-as-usual and retrofitted technologies charge the same price: $p_b^{no}(\varphi) = p_r^{no}(\varphi) = 1/[\rho\varphi]$. If, instead, a firm employs the state-of-the-art technology, it charges $p_s^{no}(\varphi) = 1/[\rho\alpha\varphi]$.

Firms choose between the three available technologies to maximize profits. If firms employ the business-as-usual technology, profits are given by $\pi_b^{no} = \frac{1}{\sigma} I [P\rho]^{\sigma-1} \varphi^{\sigma-1} - f$. Profits

[4]For example, by increasing thermal efficiency, as can be the case with low-NO_X burners (Applied Technologies of New York Inc., 2018). In general, burning less fossil fuels on-site, all else equal, has the potential to reduces a firm's emissions of many air pollutants.

[5]As our model is static, we abstract away from the intertemporal nature of this decision.

from employing the retrofitted technology are $\pi_r^{no} = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} - [f + f_r]$. Finally, profits from choosing the state-of-the-art technology are given by $\pi_s^{no} = \frac{1}{\sigma} I \left[P\rho\right]^{\sigma-1} \varphi^{\sigma-1} \alpha^{\sigma-1} - [f + f_s]$.

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

In addition, we assume $f_s > [\alpha^{\sigma-1} - 1] f$, so that the marginal surviving firm uses business-as-usual technology. In this case, the marginal producer's productivity cutoff φ_ε^{no} can be determined by noting that $\pi_b^{no}(\varphi_\varepsilon^{no}) = 0$. Substituting for $\pi_b^{no}(\varphi_\varepsilon^{no})$ and rearranging yields

$$\varphi_\varepsilon^{no} = \left[\frac{\sigma f}{I}\right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{no}}. \tag{1.3}$$

That is, firms that draw a productivity level below φ_ε^{no} exit the market, as they would not be profitable enough to pay the fixed cost of production (f).

Similarly, firms upgrade to the state-of-the-art technology s when it is profit maximizing to do so. Firms with low productivity levels choose to use the business-as-usual technology, while relatively productive firms will adopt state-of-the-art technology. 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

$$\varphi_s^{no} = \left[\frac{\sigma f_s}{\Delta_1 I}\right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{no}}, \tag{1.4}$$

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

By combining equations (1.3) and (1.4), we can can write the technology upgrading cutoff as a function of the exit cutoff

$$\varphi_s^{no} = \varphi_\varepsilon^{no} \left[\frac{f_s}{\Delta_1 f}\right]^{\frac{1}{\sigma-1}}. \tag{1.5}$$

Hence, all equilibrium values can be defined by the exit cutoff.

To solve for the industry equilibrium, as in Melitz (2003), we exploit the fact that because of free entry, in expectation, firms earn zero profits. Using this condition solves for all endogenous variables, including industry prices, firm exit, and technology choices. In equilibrium, the fixed entry cost f_ε, must equal expected profits. Allowing for some exogenous probability of exit, given by δ, then this free entry condition is

$$f_\varepsilon = \frac{1-G(\varphi_\varepsilon^{no})}{\delta}\bar{\pi}^{no}, \tag{1.6}$$

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

$$\bar{\pi}^{no} = \int_{\varphi_\varepsilon^{no}}^{\varphi_s^{no}} \pi_b^{no}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\varepsilon^{no})}d\varphi + \int_{\varphi_s^{no}}^{\infty} \pi_s^{no}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\varepsilon^{no})}d\varphi. \tag{1.7}$$

After substituting for $\pi_b^{no}(\varphi)$ and $\pi_s^{no}(\varphi)$ and exploiting equations (1.3) and (1.5), it is possible to show

$$\bar{\pi}^{no} = \frac{[\sigma-1]f}{k-\sigma+1}\Psi_1, \tag{1.8}$$

where k is the shape parameter of the Pareto productivity distribution, and

$$\Psi_1 = \left[1+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]. \tag{1.9}$$

Hence, the exit cutoff φ_ε^{no} can be obtained by substituting equation (1.8) into equation equation (1.6) and noting $1-G(\varphi)=\varphi^{-k}$. Doing so yields the following expression for the exit cutoff

$$\varphi_\varepsilon^{no} = \left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_1\right]^{\frac{1}{k}}. \tag{1.10}$$

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 (1.10) into equation (1.5)

$$\varphi_s^{no} = \left[\frac{f_s}{\Delta_1 f}\right]^{\frac{1}{\sigma-1}}\left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_1\right]^{\frac{1}{k}}. \tag{1.11}$$

The price index can be obtained by first noting that $I=L$, and then substituting equa-

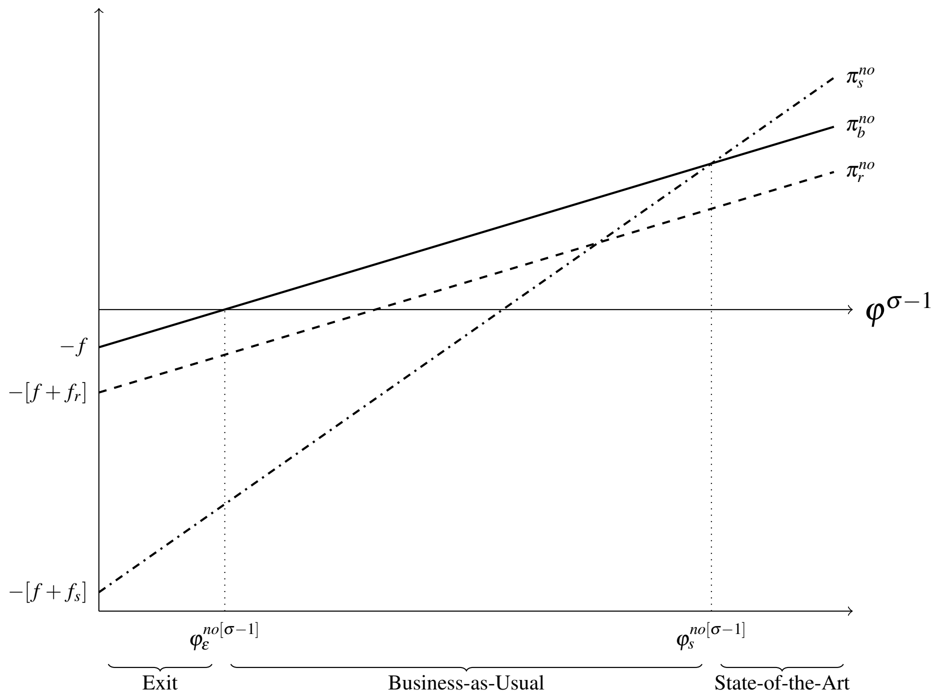

Figure 1.1: Technology Choices without Environmental Regulation

tion (1.10) into equation (1.3) and rearranging to obtain

$$P^{no} = \left[\frac{\sigma f}{L}\right]^{\frac{1}{\sigma-1}} \left[\rho \left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_1\right]^{\frac{1}{k}}\right]^{-1}. \tag{1.12}$$

The exit and technology choices made by firms are highlighted in Figure 3.1, which depicts the profits associated with adopting each technology as a function of firm productivity.[6] 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 $\varphi \in \{\varphi_\varepsilon^{no}, \varphi_s^{no}\}$, then it will produce using the business-as-usual technology. However, if a firm has a productivity level $\varphi > \varphi_s^{no}$, then the reduction in variable cost created by adopting the state-of-the-art technology is great enough to justify the fixed cost of adoption, meaning that these firms adopt the state-of-the-art technology.

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

1.3.3 The Partial Equilibrium Effects of Regulation

We now consider the effects of adopting a targeted form of environmental regulation in which firms are penalized for failing to use a clean production process. This regime can be thought of as a weak technology (or process) standard in which firms need to either adopt a specific form of production technology (or process), or face a constraint that limits the profitability of producing.[7] We first examine this policy in a partial equilibrium context in which industry prices are held fixed at the no-regulation level. This is a useful exercise as it underscores the intuition behind the selection, reallocation, and process effects of such a policy.

In this regime (labeled with superscript *tar*), the government regulates pollution using a two-part regulatory rule. If a firm uses a clean production process (either the state-of-the-art technology or the retrofitted technology), it is not subject to regulation because it is operating with the lowest emission intensity currently available. As a result, the marginal costs of production for these firms are unaffected by regulation. To make this explicit, firms that use the retrofitted technology, which we label with subscript r, face the same marginal costs with regulation (labeled with superscript *tar*) and without regulation (labeled with superscript *no*): $c_r^{tar}(\varphi) = c_r^{no}(\varphi) = 1/\varphi$. Similarly, the same is true for firms using state-of-the-art technology (labeled with subscript s): $c_s^{tar}(\varphi) = c_s^{no}(\varphi) = 1/[\alpha\varphi]$. In contrast, a firm that employs a dirty production process (the business-as-usual technology) is subject to a regulatory constraint in the form of a tax τ on each unit of pollution emitted.[8] Hence, regulation raises the marginal costs of production for these firms. Labeling business-as-usual production technology with subscript b, this means $c_b^{no}(\varphi) < c_b^{tar}(\varphi) = [1+\kappa\tau]/\varphi$. To abstract away from the redistributive aspects of environmental taxation, we assume tax revenue is not returned to consumers and is spent outside the model.

Given that firm prices feature a constant markup, this increase in marginal costs raises the price of output for firms producing with the business-as-usual technology. That is, $p_b^{no}(\varphi) < p_b^{tar}(\varphi) = [1+\kappa\tau]/[\rho\varphi]$, and profits are $\pi_b^{tar} = \frac{1}{\sigma} I \left[P^{no}\rho\right]^{\sigma-1} \varphi^{\sigma-1} [\frac{1}{1+\kappa\tau}]^{\sigma-1} - f$. This means, holding industry prices fixed, the profit from using the business-as-usual technology falls for any level of productivity φ.

This partial equilibrium outcome is depicted in Figure 3.2, which displays the technological choices made by firms when faced with targeted regulation holding industry prices (P) fixed. As the figure shows, a reduction in the profitability of using the business-as-usual technology

[7] The reduction in profitability could occur either because of increased production costs, say as the result of a tax, fine, or a production constraint.

[8] Alternatively, we could impose a production cap, without substantively affecting the intuition behind the results. We use a tax for analytic tractability.

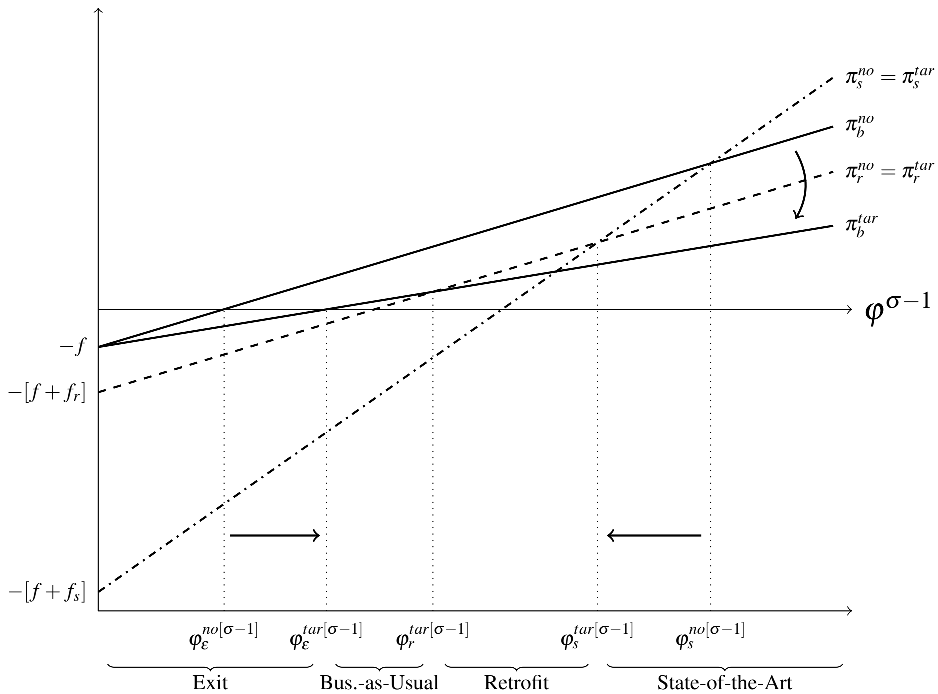

Figure 1.2: Technology Choices with Targeted Environmental Regulation

increases the productivity level for which it is unprofitable to enter the market from φ_ε^{no} to $\varphi_\varepsilon^{tar}$. As such, firms with $\varphi \in \{\varphi_\varepsilon^{no}, \varphi_\varepsilon^{tar}\}$ exit in response to regulation. Moreover, given the design of regulation, profits from using the retrofitted or state-of-the-art technology do not change. This means the increase in the variable cost of the business-as-usual technology makes technology upgrading a profitable alternative for some firms. As depicted, it is profit maximizing for firms with productivity $\varphi \in \{\varphi_r^{tar}, \varphi_s^{no}\}$ to retrofit their technology in response to regulation. For these firms, the benefit of avoided tax payments outweighs the increase in fixed production costs. Similarly, firms with productivity $\varphi \in \{\varphi_c^{tar}, \varphi_s^{no}\}$ adopt the state-of-the-art technology in response to regulation because it is now profit maximizing to do so.

While Figure 3.2 clearly highlights how environmental regulations create selection effects by causing firms to exit in response to regulation, the reallocation and process effects are not readily apparent from the figure. As such, we further explore how regulations affect firm revenues and emission intensities to make these additional effects clear.

These effects for firms that survive regulation (those with $\varphi > \varphi_\varepsilon^{tar}$) are displayed in Figure 1.3 and Figure 1.4. These figures depict the effects of environmental regulation on firm revenues (Figure 1.3) and emission intensity (Figure 1.4) holding industry prices fixed. Both figures show that the most productive firms, with productivity $\varphi > \varphi_s^{no}$, are unaffected by

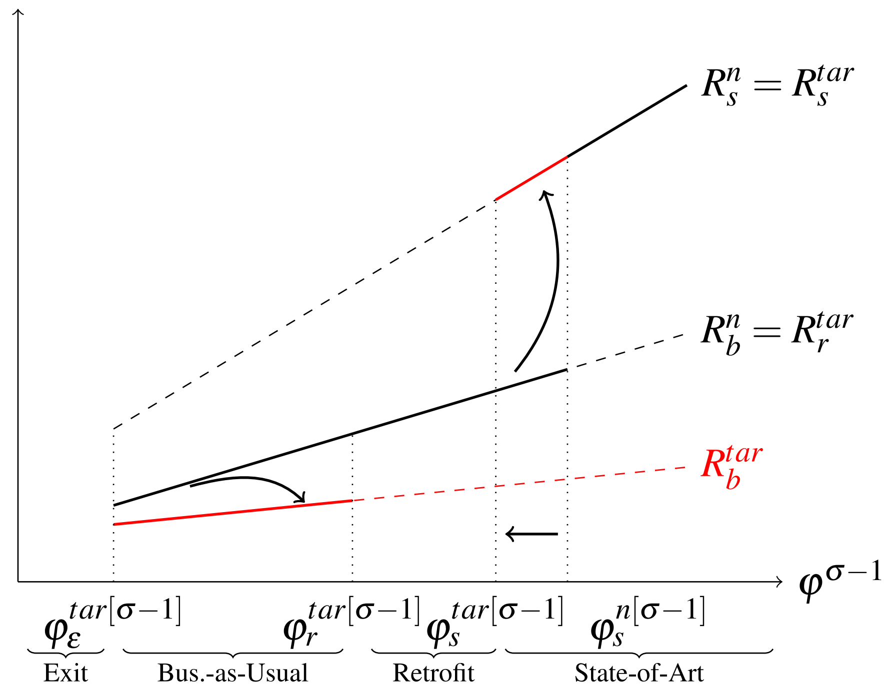

Figure 1.3: Revenues for Surviving Firms with Targeted Environmental Regulation

regulation, as they use a clean production technology in either regime. In contrast, regulation causes the least productive firms, with productivity $\varphi < \varphi_r^{tar}$, to produce less, but with the same pollution intensity. This is because they use the business-as-usual technology under either regime, and variable costs rise under regulation. Lastly, pollution intensity falls for the firms in the middle of the productivity distribution, with productivity $\varphi \in \{\varphi_r^{tar}, \varphi_s^{no}\}$. This occurs because they either retrofit or adopt state-of-the-art technology. The retrofitting firms experience no change in output, as their variable costs do not change relative to business-as-usual. However, output increases for the new state-of-the-art adopters, as both their pollution intensity and variable costs fall.

1.3.4 The Equilibrium Effects of Regulation

The discussion in the preceding section illustrated how targeted environmental regulation causes an industry to clean up through selection, reallocation, and process effects, but did so in a partial equilibrium setting. In this section, we show that similar results hold in equilibrium. In particular, we show that in equilibrium:

1. Regulation causes some firms to exit.

2. The effects of regulation vary across the productivity distribution. Revenues fall for the

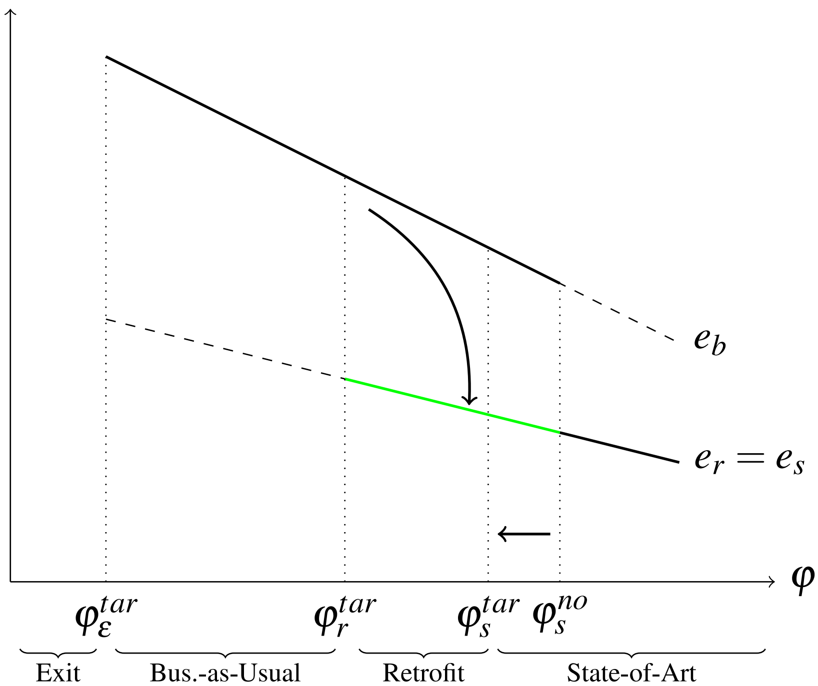

Figure 1.4: Pollution Intensity for Surviving Firms with Targeted Environmental Regulation

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.

3. 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.

Closing the Model Under Targeted Environmental Regulation

To show these equilibrium effects, we first solve the model under the targeted environmental regulation regime, and then show the relevant comparative statics. As in Section 1.3.2, we again solve the model by first deriving productivity cutoffs; however, in this case, we now also consider the cut-off for retrofitting in addition to the cut-offs for firm exit, and technology upgrading.

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

and rearranging yields

$$\varphi_{\varepsilon}^{tar} = \left[\frac{\sigma f}{I}\right]^{\frac{1}{\sigma-1}} \left[\frac{1+\tau\kappa}{\rho P^{tar}}\right]. \tag{1.13}$$

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^{tar} is defined by $\pi_b^{tar}(\varphi_b^{tar}) = \pi_r^{tar}(\varphi_r^{tar})$. Substituting and rearranging yields

$$\varphi_r^{tar} = \left[\frac{\sigma f_r}{I\Delta_2}\right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{tar}}, \tag{1.14}$$

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

$$\varphi_s^{tar} = \left[\frac{\sigma[f_s - f_r]}{\Delta_1 I}\right]^{\frac{1}{\sigma-1}} \frac{1}{\rho P^{tar}}. \tag{1.15}$$

Both φ_r^{tar} and φ_s^{tar} can be expressed as functions of the exit cutoff productivity φ_b^{tar}. The expression for the retrofitting cutoff can be obtained by combining equation (1.14) with (1.13) to get

$$\varphi_r^{tar} = \frac{\varphi_{\varepsilon}^{tar}}{1+\tau\kappa} \left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}. \tag{1.16}$$

Similarly, an expression for the upgrading cutoff can be obtained by combining equation (1.15) with (1.13) to get

$$\varphi_s^{tar} = \frac{\varphi_{\varepsilon}^{tar}}{1+\tau\kappa} \left[\frac{f_s - f_r}{\Delta_1 f}\right]^{\frac{1}{\sigma-1}}. \tag{1.17}$$

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 regulated equilibrium.

To solve for the industry equilibrium, we again exploit the fact that in expectation, firms earn zero profits due to free entry. Hence, in the regulated equilibrium, the fixed entry cost f_{ε} must equal expected profits

$$f_{\varepsilon} = \frac{1 - G(\varphi_{\varepsilon}^{tar})}{\delta} \bar{\pi}^{tar}, \tag{1.18}$$

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

$$\begin{aligned}\bar{\pi}^{tar} = &\int_{\varphi_\varepsilon^{tar}}^{\varphi_r^{tar}} \pi_b^{tar}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\varepsilon^{tar})}d\varphi \\ &+\int_{\varphi_r^{tar}}^{\varphi_s^{tar}} \pi_r^{tar}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\varepsilon^{tar})}d\varphi \\ &+\int_{\varphi_s^{tar}}^{\infty} \pi_s^{tar}(\varphi)\frac{g(\varphi)}{1-G(\varphi_\varepsilon^{tar})}d\varphi. \end{aligned} \tag{1.19}$$

Substituting for $\pi_b^{tar}(\varphi)$, $\pi_r^{tar}(\varphi)$, and $\pi_s^{tar}(\varphi)$ and utilizing equations (1.13), (1.14) and (1.15), it is possible to show

$$\bar{\pi}^{tar} = \frac{[\sigma-1]f}{k-\sigma+1}\Psi_2. \tag{1.20}$$

where:

$$\Psi_2 = \left[1+[1+\tau\kappa]^k\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]\right] > 0. \tag{1.21}$$

The exit cutoff $\varphi_\varepsilon^{tar}$ can be obtained by substituting (1.20) into (1.18) and using $1-G(\varphi) = \varphi^{-k}$

$$\varphi_\varepsilon^{tar} = \left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_2\right]^{\frac{1}{k}}. \tag{1.22}$$

Having determined $\varphi_\varepsilon^{tar}$, it is again possible to obtain expressions for φ_r^{tar} and φ_s^{tar}, and the price index, P^{tar}. The retrofitting cutoff can be obtained by substituting equation (1.22) into equation (1.16)

$$\varphi_r^{tar} = \frac{1}{1+\tau\kappa}\left[\frac{f_r}{\Delta_2}\right]\left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_2\right]^{\frac{1}{k}}. \tag{1.23}$$

The technology cutoff, on the other hand, can be obtained by substituting equation (1.22) into equation (1.17)

$$\varphi_s^{tar} = \frac{1}{1+\tau\kappa}\left[\frac{f_s-f_r}{\Delta_1 f}\right]^{\frac{1}{\sigma-1}}\left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_2\right]^{\frac{1}{k}}. \tag{1.24}$$

An expression for the price index can be obtained by substituting equation (1.22) into equation (1.13) and noting that given our assumption that environmental tax revenues are not returned to consumers $I = L$. This yields

$$P^{tar} = \left[\frac{\sigma f}{L}\right]^{\frac{1}{\sigma-1}} \left[\rho\left[\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{f}{\delta f_\varepsilon}\right]\Psi_2\right]^{\frac{1}{k}}\right]^{-1}. \tag{1.25}$$

Comparative Statics of Targeted Environmental Regulation

We begin our comparative statics exercise by examining how regulation affects firm exit. In all comparative static exercises that follow, we assume $f_r > 0$, as well as the parameter restrictions imposed in previous sections.

Proposition 1. *Targeted environmental regulation causes firms to exit.*

Proof. This claim can be proved by comparing φ_ε^{no} and $\varphi_\varepsilon^{tar}$. Doing so yields

$$\frac{\varphi_\varepsilon^{tar}}{\varphi_\varepsilon^{no}} = \left[\frac{1+[1+\tau\kappa]^k\left[\Delta_2^{\frac{k}{\sigma-1}}\left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}{1+\Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}}\right]^k. \tag{1.26}$$

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

$$[1+\tau\kappa]^k \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} > \Delta_1^{\frac{k}{\sigma-1}}\left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}, \tag{1.27}$$

which can be re-expressed as $[1+\tau\kappa]^k\,[f_s/[f_s-f_r]]^{\frac{k-\sigma+1}{\sigma-1}} > 1$. This is satisfied given the assumptions of the model. □

Proposition 1 shows that if more productive firms are less pollution intensive, then the selection channel is present for targeted regulation. In this case, targeted regulation causes the least productive firms to exit, which contributes to an industry's clean-up.

Targeted regulation results in firms exiting because the industry's price index does not rise enough to offset the increased cost of production for the least productive firms. Note that this effect need not always occur as a result of regulation. For example, Anoulies (2017) finds that, for a given initial permit allocation rule, tightening the emissions cap used in a permit trading system does not cause firms to exit. This is because the marginal producer is perfectly offset by the increase in industry prices.[9,10]

We now turn to examine the effects of targeted regulation on firm revenues. Before doing so, it is useful to note that in either regulatory regime, revenues for any firm (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 firm 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 $r_b^{no}(\varphi_\varepsilon^{no}) = \sigma f$, we have

$$r_b^{no}(\varphi) = \left[\frac{\varphi}{\varphi_\varepsilon^{no}}\right]^{\sigma-1} \sigma f \tag{1.28}$$

$$r_s^{no}(\varphi) = \left[\frac{\alpha\varphi}{\varphi_\varepsilon^{no}}\right]^{\sigma-1} \sigma f. \tag{1.29}$$

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

$$r_b^{tar}(\varphi) = \left[\frac{\varphi}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1} \sigma f, \tag{1.30}$$

$$r_r^{tar}(\varphi) = \left[\frac{\varphi}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1} [1+\kappa\tau]^{\sigma-1} \sigma f \tag{1.31}$$

$$r_s^{tar}(\varphi) = \left[\frac{\alpha\varphi}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1} [1+\kappa\tau]^{\sigma-1} \sigma f. \tag{1.32}$$

Comparing revenues from one regime to another leads to the next result.

Proposition 2. *Targeted regulation reduces revenues for surviving firms with $\varphi < \varphi_r^{tar}$.*

Proof. By definition, surviving firms with $\varphi < \varphi_r^{tar}$ use the business-as-usual technology in both the no regulation and regulation regimes. For these firms, the effect of regulation can be

[9] Anoulies, however, also shows that the initial permit allocation rule does affect exit.

[10] In contrast, Andersen (2018) finds that tightening an emission tax causes firms to exit. This result is driven by Andersen's assumption on the abatement process used by firms, as this results in the pollution tax affecting the fixed cost of production.

determined by comparing equations (1.28) with (1.30), yielding

$$\frac{r_b^{tar}(\varphi)}{r_b^{no}(\varphi)} = \left[\frac{\varphi_\varepsilon^{no}}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1}. \tag{1.33}$$

Given $\varphi_\varepsilon^{tar}/\varphi_\varepsilon^{n} > 1$ (from Proposition 1), it follows that $r_b^{tar}(\varphi) < r_b^{no}(\varphi)$. □

Proposition 3. *Targeted regulation has an ambiguous effect on revenues for surviving firms with $\varphi \geq \varphi_r^{tar}$.*

Proof. See Appendix A.1. □

Corollary 1. *If $[1+\tau\kappa] > \alpha$, then targeted regulation causes the largest reduction in revenue for firms with $\varphi < \varphi_r^{tar}$.*

Proof. See Appendix A.1. □

Proposition 2 and Proposition 3 show that targeted regulation has differential effects on revenues for firms of different productivity levels. Corollary 1 makes it clear that, with the exception of one special case[11], the largest reduction in revenues occurs for the least productive surviving firms in an industry. This is intuitive, as the policy's design results in increased variable costs only for those firms that fail to adopt cleaner processes. As the adoption of cleaner processes requires paying a fixed cost, it is the least productive firms in an industry that face an increase in variable costs directly as a result of regulation. Revenues for these firms must fall as a result of regulation.

We next turn to examine the effects of targeted regulation on the emission intensity of a firm's production. Recall that the effect of regulation on firm emission intensities is determined by the adoption of new technology. As the adoption of new technology requires paying a fixed cost, similar to targeted regulation's effect on firm revenues, its effect on a firm's emission intensity also depends on the firm's productivity.

Proposition 4. *Targeted regulation reduces the emission intensity of a firm with $\varphi \in [\varphi_r^{tar}, \varphi_s^{no}]$, and does not change the emission intensity of other firms.*

Proof. In the no-regulation regime, no firm retrofits. However, in the regulation regime, a firm with $\varphi \in [\varphi_r^{tar}, \varphi_s^{tar}\}$ adopts the retrofitted technology. By construction, the emission intensity of these firms falls, as $\frac{e_r^{tar}(\varphi)}{e_b^{n}(\varphi)} = \frac{e_s^{tar}(\varphi)}{e_b^{n}(\varphi)} = \frac{1}{\gamma} < 1$. If $f_s > \frac{\Delta_1+\Delta_2}{\Delta_2} f_r$ and $f_r > \left[[1+\tau\kappa]^{\sigma-1} - 1\right] f$, then a positive measure of firms retrofit in response to environmental regulation. These

[11]This special case is discussed in the appendix.

inequalities are satisfied by assumption, meaning targeted regulation causes some firms to retrofit in response to regulation.

Given our setup, the emission intensity of firms that use the business-as-usual technology in both regimes is unaffected. Similarly, there is no change in emission intensity form firms that downgrade from state-of-the-art to retrofitted technology, or use state-of-the-art technology in both regimes. □

Proposition 4 shows the process effect of regulation. For a targeted regulation, this process effect is driven by firms in the middle of the productivity distribution, as they respond to regulation by adopting process changes. The least productive firms, in contrast, may contribute to a clean-up through the reallocation effect, as is made clear in Proposition 2.[12]

We are not the first to show regulation may have differential effects on firms of different productivity levels. Cao et al. (2016), for example, find low productivity firms respond to a uniform pollution tax by allocating inputs that could have been used in production towards abatement. This relates to Corollary 1 above, as given the same amount of inputs, this reallocation of inputs would serve to reduce output. Cao et al. also find that more productive firms adopt less polluting technology, which is similar to our result presented in Proposition 4.[13] Our results differ from Cao et al. in two important ways. First, Cao et al.'s results require imposing a quasi-linear preference system, and does not hold if preferences are of the constant elasticity of substitution form, as is commonly assumed in the heterogeneous firms literature. Second, we show that under targeted regulation, the most productive firms in an industry are not directly affected by regulation because they use relatively clean processes in the absence of regulation. This is an important distinction between targeted regulation and regulation that directly affects the variable costs of all producers, such as a uniform tax or a cap-and-trade system.

Given the connection between technology adoption and the clean-up channels, a natural question is to ask how changing the fixed costs of retrofitting affects these channels?

Proposition 5. *Reducing f_r lowers φ_r^{tar}, and if f_s is not too large, also lowers $\varphi_\varepsilon^{tar}$.*

Proof. See Appendix A.1. □

Corollary 2. *Reducing f_r lowers the measure of firms that use the business-as-usual technology.*

[12]If more productive firms are less pollution intensive, then Proposition 2 and Corollary 1 imply a reallocation effect.

[13]For details, see Proposition 5 in Cao et al. (2016).

Proof. See Appendix A.1. □

Proposition 5 implies that a regulated equilibrium with a lower fixed cost of retrofitting will have a larger mass of firms using clean processes (either the retrofitted or state-of-the-art technology), relative to a regulated equilibrium with a relatively high retrofitting cost. In addition, the low retrofitting cost equilibrium will also have less exit, if the fixed cost of state-of-the-art technology is not too large. Consequently, lowering the fixed cost of retrofitting increases the magnitude of the process effect relative to the selection effect. Corollary 2 shows that, in addition, lowering the retrofitting cost leads to a smaller measure of firms that use the business-as-usual technology following regulation, who are the drivers of the reallocation effect.

1.4 Discussion

This model makes a number of clear predictions about the channels through which targeted regulation may cause an industry to clean-up. First, regulation will cause the least productive plants in an industry to exit. This occurs because in equilibrium, relatively low-productivity firms become less profitable, which reduces the measure of firms for whom it is worth paying the fixed cost of production. If, as we have assumed, more productive plants are less pollution intensive, then these plants exiting will reduce the industry's pollution intensity. That is, a selection effect induced by regulation will contribute to the industry's clean-up.

Second, the process effect arises from firms choosing to adopt cleaner production processes in response to regulation, which makes them less pollution intensive. However, only firms in the middle of the industry's productivity distribution should be induced by regulation to adopt cleaner processes. As the adoption of cleaner processes require paying a fixed cost, relatively low-productivity firms will not be profitable enough to pay these fixed costs. Moreover, the most productive firms in an industry use industry-leading processes and technologies even in the absence of regulation. That is, regulation does not affect their incentive to use the industry-leading technologies, and does not cause them to become cleaner. Firms in the middle of the productivity distribution, however, are productive enough to warrant adopting cleaner processes, but not productive enough to do so absent regulation.

Finally, the equilibrium firm size falls for the least productive firms that remain in the industry. This occurs because adopting a cleaner production process, which allows the firm to avoid the pollution tax, carries a fixed cost. As a result, some low-productivity plants will find it optimal to produce, but will not be profitable enough to pay the fixed cost to change processes. Hence, in equilibrium, variable production costs rise for the least productive plants, and their optimal size shrinks. Again, if more productive plants are less pollution intensive, then the reduction in output from these plants will reduce the industry's pollution intensity. This is the reallocation effect through which regulation contributes to an industry's clean-up.

In addition to clarifying the channels involved in a clean-up, this model shows the important role the fixed cost of adopting cleaner processes play in a clean-up. These fixed costs determine the number of firms (or measure of firms, to be more precise in this context) that adopt cleaner processes in response to regulation. If these fixed costs are very low, then a large number of firms should choose to adopt cleaner processes in response to regulation. This, in turn, reduces the number of firms that face higher production costs from regulation, and leads to less exit and a smaller reduction in output from surviving low-productivity firms. As a result, the fixed cost of process changes directly affects the magnitude of the process effect relative to the reallocation and selection effects.

1.5 Conclusion

In this chapter, we present a new theoretical framework to capture the channels through which targeted environmental regulation – regulation that requires firms to adopt clean processes, and penalizes those that fail to do so – causes an industry to clean up. We start by replicating a decomposition exercise by Cherniwchan et al. (2017) to show there are three distinct channels through which environmental regulation can cause a reduction in an industry's pollution intensity. The first channel, which we call the selection channel, captures the exit of plants in response to regulation. The second channel, which we call the reallocation channel, captures the reduction in output at surviving regulated plants in response to regulation. Lastly, the process channel captures the adoption of cleaner production processes at surviving regulated plants.

We then present a theoretical model to show how these channels may arise in response to targeted environmental regulation. We study this type of regulation because it is a common form of air pollutant regulation (for example, the US Clean Air Act contains this type of regulation), and has not yet been studied in the theoretical literature on the economic consequences of environmental regulation.

Our model is based on a closed-economy variant of the Melitz (2003) model in which pollution from heterogeneous firms is regulated. In this model, firms choose whether to produce, how much to produce, and what type of production process or technology to use. Unlike the majority of the relevant environmental economics literature, we assume firms upgrade to a cleaner production technology by paying a fixed cost, rather than facing higher variable costs.

We show that targeted regulation causes the least productive firms in an industry to exit, as well as relatively low-productivity surviving firms to shrink. If these firms are the most pollution intensive in the industry, then these results imply that targeted regulation will create both selection and reallocation effects that serve to clean up the industry. In addition, we find moderately productive surviving firms adopt cleaner processes in response to targeted regulation, which means the process effect also contributes to a clean-up. Finally, the most productive firms in an industry are not directly affected by targeted regulation, in terms of pollution intensity or production, as they use industry-leading technology with or without regulation.

Chapter 2

Estimating the Regulatory Channels of the Manufacturing Clean-Up

2.1 Introduction

Manufacturing pollution intensity – the amount of pollution emitted from the manufacturing sector per dollar of output shipped – has plummeted across much of the industrialized world in recent decades. From 1990 to 2008, for example, the pollution intensity of the US manufacturing sector fell by up to 77% for some pollutants (Levinson, 2015). Similar reductions have been shown for the European manufacturing sector (Brunel, 2016), and in several countries for sulphur dioxide emissions (Grether et al., 2009). Taken together, it appears the production of goods in much of the industrialized world is becoming cleaner.

Work on this manufacturing "clean-up", spurred by Levinson's (2009) seminal paper, has concluded that the source of this clean-up appears to be reductions in industry- (Brunel, 2016; Levinson, 2009), and even product-level (Shapiro and Walker, 2015) pollution intensity. That is, changes in the composition of industries and products produced in these countries do little to explain the clean-up. One obvious hypothesis as to why industries have become cleaner is that environmental regulations may have changed the way in which industries operate, thereby pushing them to become less pollution-intensive. Given the primary goal of regulation is to address pollution problems, this hypothesis seems plausible. Yet, there is little direct evidence of how regulation causes an industry to clean-up. The goal of this chapter is to contribute to this debate by asking how plant-level responses to environmental regulation have contributed to the manufacturing clean-up.

While there is little direct evidence of regulation's role in the clean-up, there is considerable indirect evidence that regulation is a potential, if not driving, cause. For example, Shapiro and Walker (2015) use a structural model to ask whether the clean-up of the U.S. manufacturing sector is consistent with changes in regulation, trade, productivity growth, or other economic factors. They conclude the clean-up would require a doubling of their model's shadow-price of emissions, which is consistent with regulation playing a large role in the clean-up. While this presents compelling evidence for the importance of regulation, it is not causal.

There is also indirect evidence on regulation's role in the clean-up from the literature focused on estimating the causal effects of regulation. This literature has shown, for example, that regulation may cause plants to exit (e.g. Becker and Henderson (2000); Henderson (1996); List et al. (2003)), and surviving plants to shrink (e.g. Greenstone (2002)). If these displaced and contracted producers were relatively pollution intensive, then these changes would serve to clean up an industry. Other work in this area has shown regulation may affect the production processes at regulated plants, by altering input-use and productivity (e.g. Berman and Bui (2001b), Greenstone et al. (2012), and Walker (2013)), for example. However, without connecting these changes to a plant's pollution intensity, it is unclear whether these would be viable mechanisms through which regulation could lead to a cleaner industry.

In this chapter, we present causal estimates of the clean-up caused by a particular regulatory change. To do this, we estimate the effect of a policy change on each of the three plant-level determinants of an industry's pollution intensity: the number of plants operating in an industry (which we call the "selection" effect), changes in output at regulated plants (the "reallocation" effect), and changes in plant emission intensity (the "process" effect). We then develop a simple framework to translate these causal micro-level estimates to estimates of the aggregate channels of the clean-up. Taken together, these estimates provide a complete characterization of how plant-level responses to regulation have contributed to the aggregate clean-up of manufacturing.

To obtain causal estimates of the regulatory channels of the clean-up requires observing longitudinal information on both pollution emissions and productive activities of individual manufacturing plants. We obtain this information from a newly created confidential Canadian dataset. This dataset is one of just a few such datasets in the world, and to the best of our knowledge, provides the most complete coverage of a nation's manufacturing sector of any of these datasets. We use this data to estimate the effect of a a major revision to Canadian air quality regulation - the implementation of the Canada Wide Standards for Particulate Matter and Ozone (CWS).

In addition to providing, potentially, the best data to answer our research question, Canada is a good setting for this study because the resulting estimates should be informative for understanding the mechanisms through which environmental regulations have contributed to the manufacturing clean-ups elsewhere. First, the policy we study is similar in design to the main air pollution regulation in the US, the Clean Air Act (CAA), and shares features of the air quality standards in place in Europe. While there is much evidence of the CAA's effects on American producers', we are the first to exploit the CWS to identify the effects of environmental regulation.[1] As well as having similar policy, Canada's clean-up appears to be very similar to those that have been 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.

Our goal is to estimate the effects of the CWS on the pollution intensity, output, and entry and exit decisions of affected Canadian manufacturing plants. To do so, we exploit variation in regulatory stringency created by the design of the CWS. The CWS was designed to ensure each region met a minimum level of air quality by establishing thresholds for the ambient concentration of $PM_{2.5}$ and 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 were explicitly focused on plants in "targeted industries" that were viewed as primary contributors to poor air quality. 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 find robust evidence that the CWS reduced pollution emissions from affected manufacturing plants. For the average $PM_{2.5}$ emitting plant, the CWS is associated with a 15% reduction in $PM_{2.5}$ emissions. Furthermore, the CWS is associated with a 33% reduction in NO_X emissions from the average NO_X emitting plant.

The theory outlined in Chapter 1 predicts that these reductions will be driven by different mechanisms. If the fixed costs of process changes are high on average, as we argue is the case for $PM_{2.5}$, 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

[1] The CWS has been discussed previously, however, in the environmental policy literature (see, e.g., Angle (2014)).

[2] The 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 2.3.

intensity of the average plant. If the fixed costs of process changes are low on average, as we argue is the case with NO_X, then even less productive plants should adopt cleaner production processes.[3] 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 little evidence that the CWS affected the emission intensity of the average regulated $PM_{2.5}$ emitting plant, but find evidence it 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 $PM_{2.5}$ 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 $PM_{2.5}$, 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 considerably to the clean-up of the Canadian manufacturing sector. To make this contribution explicit, we develop an empirical analogue of the industry decomposition suggested by Cherniwchan et al. (2017). This approach allows us to translate our point estimates into estimates of the process, reallocation and selection effects induced by environmental regulation. These estimates suggest that the effects of the CWS explain close to 21% of the reduction in the $PM_{2.5}$ intensity of Canadian manufacturing, but nearly 61% of the reduction in aggregate NO_X intensity. Moreover, the mechanisms driving these responses vary starkly across pollutants; the $PM_{2.5}$ 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.

The model presented in Chapter 1 suggests these differential responses to regulation are due to differences in the fixed cost of adopting cleaner production processes across pollutants. While we have focused on this channel given the available evidence documenting the substantial differences in the average costs of becoming less $PM_{2.5}$ and NO_X intensive, we do not observe these costs directly. Hence, to provide further evidence that our estimates are consistent with this mechanism, we also examine the heterogeneity in plant responses to regulation. The model suggests the effects of the CWS should only vary across plants of different productivity levels if the fixed costs of adopting cleaner processes 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 resulting estimates match our model's predictions. We find pollution from relatively

[3] We present details on the relative costs associated with process changes for these two pollutants in Section 2.3.

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 the fixed costs of adopting cleaner production processes.[4]

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.[5] Our work adds to this body of research in two ways. First, we show that the aggregate trends that have been documented in the US (Levinson, 2009, 2015) and Europe (Brunel, 2016) extend to Canada. Second, we provide causal evidence of how air quality regulations have affected these trends.[6]

Our work is also closely related to that of Martin et al. (2014) who estimate the effects of a carbon tax on the exit, sales, and energy intensity of UK manufacturing plants. While the potential clean-up of UK manufacturing is not the focus of their work, these factors serve as the determinants of industry carbon dioxide intensity and, as such, their estimates could be used to understand how carbon taxes have affected the carbon intensity of the UK manufacturing sector. Despite this, our work differs along three key dimensions. First, we observe plant-level pollution emissions, rather than energy use. Second, we show how plant level responses such as those estimated by Martin et al. can be used to obtain estimates of the process, reallocation and selection effects. Third, we use insights from a stylized model to understand the mechanisms that could be driving these effects.

This chapter also relates to the set of empirical studies 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

[4] We 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 below, we find little evidence to support these explanations.

[5] Others 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)).

[6] Our work is also related to that of Barrows and Ollivier (2018), who study how a potential mechanism that could be driving the process effect responds to changes in industry competition: within-plant changes in product mix. We are unable to explore this channel as our data does not contain information on pollution emitted by product line. Instead, we show direct evidence of a process effect in response to a regulatory change.

(Greenstone, 2003) and level (Gibson, 2016) of air pollutant emissions from manufacturing plants. This chapter complements this work by determining whether changes in plant pollution in response to regulation are due to changes in the level of output produced, or changes in the emission intensity of production.

Lastly, our work also relates to a large literature examining the effects of air quality regulation on various aspects of manufacturing plant operations. Our work is most closely related to the papers that have provided preliminary evidence of the importance of selection and reallocation effects by either examining the effects of regulation on plant entry and exit (e.g. Becker and Henderson (2000); Henderson (1996); List et al. (2003)) or plant output (e.g. Greenstone (2002)). We build on these earlier studies by also estimating the effects of air quality regulation on plant emission intensity, which allows us to provide the first estimates of the process effects induced by regulation.[7] We also build on this earlier work by showing that the effects of environmental regulation may vary across plants of different productivity levels.

The remainder of this chapter proceeds as follows. In Section 2.2, we document the clean-up of the Canadian manufacturing sector. Section 2.3 provides a brief overview of the CWS. Section 2.4 presents our data, outlines our research design and empirical specification, and presents our empirical results. Finally, Section 2.5 concludes.

2.2 The Clean-Up of Canadian Manufacturing

Our goal in this chapter 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 (2016)), it has yet to be documented in Canada. Hence, 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 2.1. The figure depicts changes in the aggregate emissions of four common pollutants from the Canadian manufacturing sector, as well as changes in aggregate manufacturing output. As it shows, 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

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

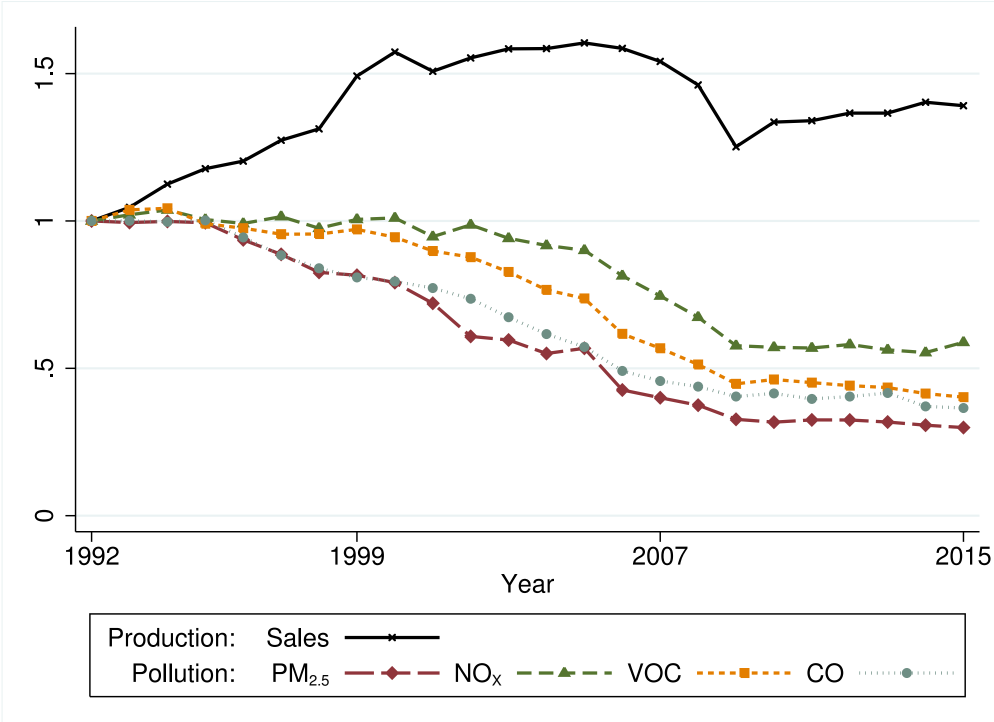

Figure 2.1: Output and Pollution from Canadian Manufacturing: 1992-2015

Notes: Figure depicts trends from 1992 to 2015 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. All series are expressed relative to their 1992 levels.

in the U.S. and Europe. For example, Levinson (2015) finds the emission intensity of US manufacturing fell by 3.6-4.3% annually from 1990 to 2008. Similarly, Brunel (2016) shows the emission intensity of European manufacturing fell by 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 intensity are driven by a "composition effect" created by a reallocation of economic activity from dirty emission-intensive industries to clean industries 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,

Table 2.1: Canadian Manufacturing Emission Intensity Decomposition: 92-15

Pollutant	Δ Emission Intensity (1)	Technique Effect (2)	Composition Effect (3)	Technique Share (4)
$PM_{2.5}$	-79	-78	-1	0.99
NO_X	-58	-52	-6	0.90
VOCs	-71	-67	-4	0.94
CO	-74	-73	-1	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 column reports the percentage change in emission intensity from the manufacturing sector. The second and third columns report the reduction in aggregate emission intensity due the technique and composition effects, respectively. The fourth column shows the fraction of column (1) attributable to changes in the technique effect, calculated as (column (2)/column (1)).

output, and pollution intensity of the manufacturing sector, respectively. Let Z_i, X_i, and E_i denote the same for individual manufacturing industries[8], 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

$$dE = \sum_i E_i d\theta_i + \sum_i \theta_i dE_i. \tag{2.1}$$

The first term of equation (2.1) is the aforementioned composition effect, while the second term is the technique effect.

We follow the approach taken by Levinson (2015) and take equation (2.1) 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 2.1. The first column reports the change in emission intensity that occurred for manufacturing as a whole. The second and third columns report the change in aggregate emission intensity attributable to the technique effect and composition effects, respectively.[9] The final column reports the share of the emission intensity change due to the technique effect.

The estimates reported in Table 2.1 suggest that the clean-up of the Canadian manufac-

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

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

turing sector can primarily be attributed to the technique effect. For example, the estimate reported in the first row indicates that during the 1992-2015 period, changes in industry emission intensity accounted for 99% of the reduction in manufacturing $PM_{2.5}$ intensity. This is further evidence that the Canadian clean-up is similar to those observed elsewhere; as shown by Levinson (2009, 2015) and Brunel (2016), the clean-ups of US and European manufacturing are also primarily due to the technique effect.[10] As a reference, the technique effect's share in Canada is generally higher than in the US, but lower than in Europe.

2.3 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.[11] Moreover, the design of the CWS makes it an attractive setting for studying the effects of environmental regulation.

First signed in 2000, the CWS was an agreement between the federal government of Canada and the various provincial environment ministries. The intent of the CWS was to improve air quality across the country by the end of 2010 by implementing two air quality standards – one for $PM_{2.5}$ and one for O_3 – that applied to each major town or city in Canada (we call these Census Metropolitan Areas or CMAs).[12] Much like the National Ambient Air Quality Standards at the centre of the U.S. Clean Air Act Amendments (CAAAs), these standards created a target level of air quality that were 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. The 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 98^{th} percentile of each region's 24-hour ambient concentration in a given year. The O_3 standard was applied as an 8-hour standard that required each CMA's O_3 concentration lie below 65 parts per billion (ppb). Achievement of the O_3 standard was based on the 4^{th} highest 8-hour concentration reported in a given year. Both the O_3 and $PM_{2.5}$ standards were

[10] In 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.

[11] It 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.

[12] The agreement defines a major town or city as a Census Agglomeration (CA) or Census Metropolitan Area (CMA). A CMA must have a total population of at least 100,000, while a CA must have a core population of at least 10,000. We use the term CMA for both.

determined over the calendar year.[13,14] These standards were intended to be distinct. That is, regions violating the $PM_{2.5}$ standard were required to address their $PM_{2.5}$ problem, and were not required to improve O_3 concentrations (unless they were also in violation of the O_3 standard).

In addition to differentiating between regions on the basis of air quality, the CWS explicitly designated a set of "targeted industries" that were to be the focus of more stringent regulation. These 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). 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.[15]

The key regulatory approach used by the CWS was to subject regulated plants – those in targeted industries and CMAs with ambient concentrations of either $PM_{2.5}$ or O_3 in excess of the relevant standard's threshold – to more stringent environmental regulation than other plants in the country. While in principal the agreement allowed provinces to choose from a number of different regulatory approaches, in practice provinces primarily used their annual operation permit systems to regulate plants. In general, these provincial systems require plants to prove compliance with certain environmental regulations in order to operate in any year (see, e.g. Canadian Council of Ministers of the Environment (2006); Environment Canada (2002); Environment Canada and Forest Products Association of Canada (2004)). To address the CWS, facilities could effectively follow one of two paths to meet the permitting requirements: either adopt technical changes recommended to their industry in the CWS (Government of Canada, 2003), or reduce activities contributing to the problematic pollutant. When local air quality was relatively clean (i.e. regions were in compliance with the CWS), the permitting constraints were laxer than when air quality was poor. Consequently, regulatory stringency facing a plant varied over time according to its region's air quality.

In essence, these regional air quality standards were used to trigger technology standards on regulated facilities, with the important caveat that facilities could choose to simply reduce their polluting activity rather than adopt cleaner technology.[16] This same approach was used

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

[14]For comparison, the National Ambient Air Quality Standards in the United States currently contain a 24-hour $PM_{2.5}$ standard set at $35\mu g/m^3$, and an 8-hour O_3 standard set at 70 pbb (Environmental Protection Agency, 2016).

[15]Some non-targeted industries were subject to other, more limited, forms of regulation. However, these were not explicitly part of the CWS and did not feature the same regional variation as the CWS air quality standards. We describe the other relevant regulations in Appendix A.2.5.

[16]In some instances explicit production constraints limiting the amount of polluting activity in a given calen-

for both the $PM_{2.5}$ and O_3 standards.

While the same regulatory approach was in place for both pollutants, due to technical constraints the options available for plants to adopt cleaner processes appear to have differed considerably across pollutants. While a full exploration of the myriad of process changes available to plants is beyond the scope of this paper, we present evidence of these differences as they pertain to manufacturing facilities, particularly in Canada. We present this evidence as both our theory and empirical results suggest the size of these process costs play an important role in dictating the channels through which an industry cleans up.

We first consider NO_X, the main ozone precursor targeted by the O_3 regulations, which is primarily caused by the combustion of fossil fuels. Industrial facilities can reduce NO_X emissions at a relatively low fixed-cost by adopting efficient combustion processes[17] or by adopting low-NO_x emissions burners (see, e.g. Environment Canada (2002), Canadian Council of Ministers of the Environment (1998b), or Environmental Protection Agency (1999a)). Alternatively, post-combustion processes, such as selective non-catalytic reduction and selective catalytic reduction, can be installed at a relatively high fixed-cost.[18] In the CWS context, the low cost process changes appear to have been sufficient to satisfy the policy's constraints. This is explicitly noted in Canada's federal emissions guidelines for both industrial boilers and heaters, and cement kilns, both of which were intended to provide the basis for process changes mandated by the CWS. The former, for example, states the guidelines for industrial boilers "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, 1998b). In addition, the document further claims that a post-combustion control technology would be required "only in isolated cases" (Canadian Council of Ministers of the Environment, 1998b). The cement kiln guidelines also note that combustion modifications can be achieved at lower costs than low-NO_X burners, and both of these are considerably less costly than post-combustion processes. While combustion modification costs are not provided, annualized costs for installing post-combustion processes are listed as four to sixty times the cost of low-NO_X burners (Canadian Council of Ministers of the Environment, 1998a).

In contrast, industrial $PM_{2.5}$ emissions are caused by a number of processes, several of which could potentially occur at the same facility. These processes include, for example, the combustion of fossil fuels, chemical reactions, wear and tear on machinery, and the processing of lumber. Similar to NO_X, there are both low- and high-cost approaches to reducing

dar year were imposed on facilities. These appear to have primarily required percentage reductions relative to base-year production levels, and as a consequence likely varied across facilities of different sizes.

[17]This may entail changing the temperature or fuel-oxygen ratio of combustion.

[18]Note that these typically produce larger emissions reductions than other alternatives.

$PM_{2.5}$ emissions. The low-cost approaches include fuel-switching, inertial separators, or wet scrubbers. These low-cost methods, however, can have limited applicability in industrial uses (World Bank Group, 1998a). Fuel switching primarily pertains to coal, and only affects the emissions from fuel use. Inertial separators are primarily intended for medium and coarse particulate matter, and are not particularly effective for $PM_{2.5}$. While wet scrubbers can accommodate $PM_{2.5}$ emissions, they are intended for production that involves a wet process. Instead, for the typical industrial facility, reducing $PM_{2.5}$ emissions requires installing a large filtration system, such as a baghouse or electrostatic precipitator, that carries a relatively large fixed-cost. This was also explicitly noted in the context of the CWS (see, e.g., Environment Canada (2002) and Environmental Protection Agency (1998, 1999b, 2002)).

The available evidence, then, suggests that Canadian industrial facilities could accomplish NO_X process changes at a much lower cost than the $PM_{2.5}$ process changes. As a reference, engineering abatement cost estimates for relevant process changes, adjusted to a per-ton basis, 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).[19,20]

2.3.1 Air Quality Improvements and the CWS

Before turning to our main empirical analysis, we first present descriptive evidence that the improvements in ambient air quality in Canada over the CWS period were consistent with the design of the CWS. In particular, as the CWS regulated facilities in the dirtiest regions of the country, air quality improvements should have been most pronounced in those particular regions.

To examine the changes in Canadian air quality over this period, we use ambient air quality data from Environment and Climate Change 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 air quality data. Each monitoring station is operated by a provincial authority, and the federal environment ministry oversees the network. Hourly monitor-level pollution concentration measures are available for ozone, most Criteria Air Contaminants, and some heavy metals (for data, see: Environment and Climate Change Canada

[19]No cost estimates are available for combustion modifications to reduce NO_X emissions, but these should be lower than that of the low-NO_X burners (Canadian Council of Ministers of the Environment, 1998a).

[20]It is worth noting that the process changes available to extremely large emitters, such as electric utilities, can differ considerably to the changes available to a typical industrial facility. Most notably, NO_X abatement is primarily achieved through high-cost post-combustion methods at these large utilities (World Bank Group, 1998b).

(2013)).

We use this data to construct the regional air quality measures used by the CWS. For $PM_{2.5}$, we construct the 98^{th} percentile of each CMA's 24-hour concentration in a given year.[21] For O_3, we construct the 4^{th} highest 8-hour concentration reported in a CMA in a given year.[22] For any CMA that contains more than one monitor, we follow the rule defined by the CWS and compute the average pollution concentration across all monitors for the $PM_{2.5}$ measurements and the maximum concentration for the O_3 measurements (Canadian Council of Ministers of the Environment, 2002, p. 12).

With this data, we then sort each of the CMAs into one of two groups for each standard: "clean" CMAs that never violated the relevant standard, and "dirty" CMAs that violated the standard at least once over the phase-in period. Doing this allows us to assess whether the changes in air quality across Canada matched with the design of the CWS.

Over the period 2000 to 2011, there was no significant change in mean $PM_{2.5}$ concentrations among the CMAs that never exceeded the $PM_{2.5}$ standard. Similarly, there was no significant changes in mean O_3 concentrations among the CMAs that never exceeded the O_3 standard. Mean $PM_{2.5}$ concentrations in the clean CMAs was approximately 15 $\mu g/m^3$ in each year; mean O_3 concentrations were between 55 and 58 ppb in each year. In contrast, mean $PM_{2.5}$ in the dirty CMAs fell from approximately 30 $\mu g/m^3$ in the beginning of the decade to approximately 22 $\mu g/m^3$ at the end. Similarly, mean O_3 in the dirty CMAs fell from approximately 80 ppb in the beginning of the decade to approximately 68 ppb at the end of the phase-in.[23] These changes in regional air quality are shown in Figure 2.4 and Figure 2.5, which plot the yearly mean pollution concentrations for the clean and dirty cities from 2000 to 2011, the entire period of the CWS.

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

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

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

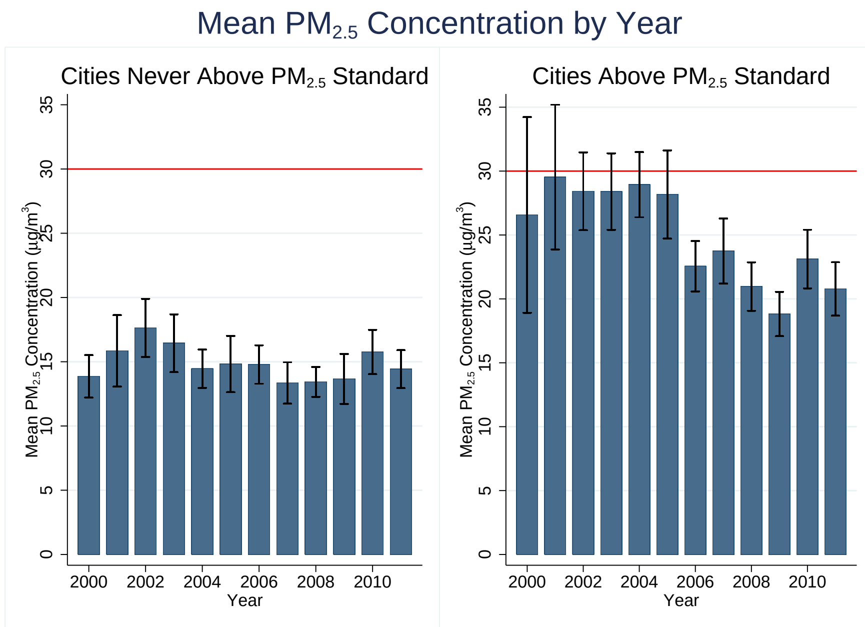

Figure 2.2: Mean $PM_{2.5}$ concentration by year with 95% confidence intervals. Panel A shows cities never above the $PM_{2.5}$ standard. Panel B shows cities above the $PM_{2.5}$ standard at least once. The red line represents the threshold for the $PM_{2.5}$ Standard. The air quality metric used is the 98^{th} percentile of each city's 24-hour concentration.

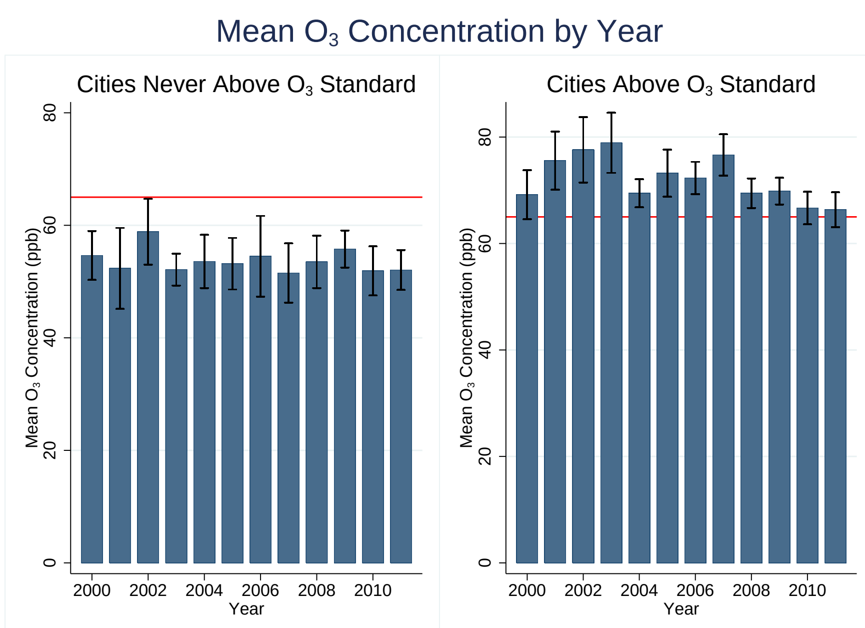

Figure 2.3: Mean O_3 concentration by year with 95% confidence intervals. Panel A shows cities never above the O_3 standard. Panel B shows cities above the O_3 standard at least once. The red line represents the threshold for the O_3 Standard. The air quality metric used is each city's 4^{th} highest 8-hour concentration.

As a further check, we also examine how the distributions of CMA air quality changed from the first half (2000-2005) to the second half (2006-2011) of the decade. As the intent of the CWS was to reduce extreme measures of air pollution, the largest change in air quality over this period should occur in the top of the air quality distribution. For this exercise, we again separate CMAs into two groups: (i) clean CMAs where ambient pollution concentrations were never above the CWS, and (ii) dirty CMAs that exceeded the CWS at least once. We estimate these distributions using kernel density estimation with a Gaussian kernel.

These distributions are depicted in Figure 2.4 and Figure 2.5. As the figures show, there was almost no change in either of the $PM_{2.5}$ or O_3 distributions for clean CMAs over the entire phase-in period. The same, however, cannot be said for the dirty CMAs. The $PM_{2.5}$ distribution shifted drastically from the beginning to the end of the phase-in, with almost all of the CMA-year observations lying below the CWS threshold in the second half of the phase-in. The right tail of the O_3 distribution shifted leftward, and the mass of CMA-years near the CWS threshold increased substantially. By the end of the phase-in period, most CMAs in Canada had met the $PM_{2.5}$ standard, and the dirty O_3 cities were moving towards compliance. This provides further evidence of changes in air quality consistent with the CWS.

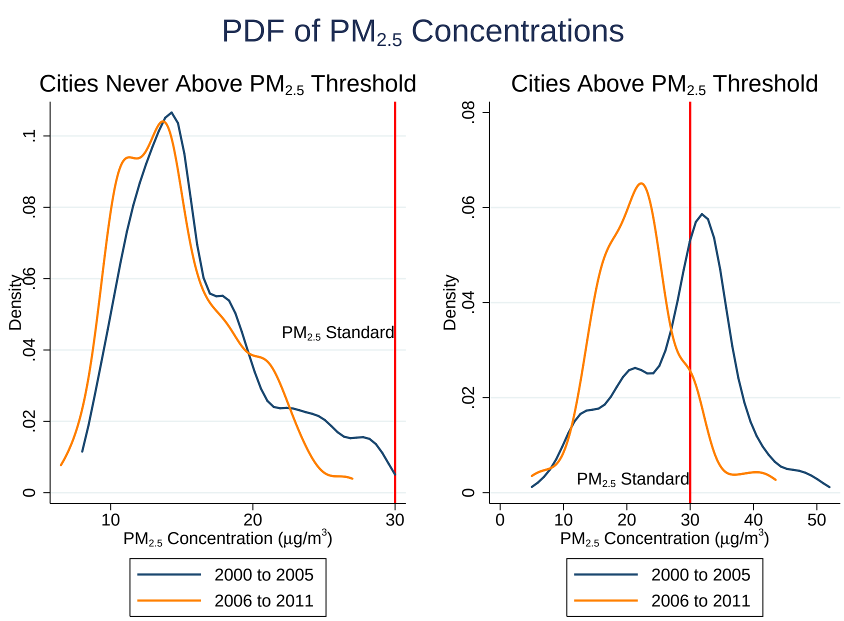

Figure 2.4: Kernel density estimates of the distribution of $PM_{2.5}$ concentrations across CMAs in the first half (2000-2005) and second half (2006-2011) of the CWS phase in period. The panel on the left displays pollution concentrations for CMAs that never exceeded the $PM_{2.5}$ standard. The right panel displays the pollution concentrations for CMAs that exceeded the $PM_{2.5}$ standard at least once. The vertical red lines represents the threshold used for the $PM_{2.5}$ standard.

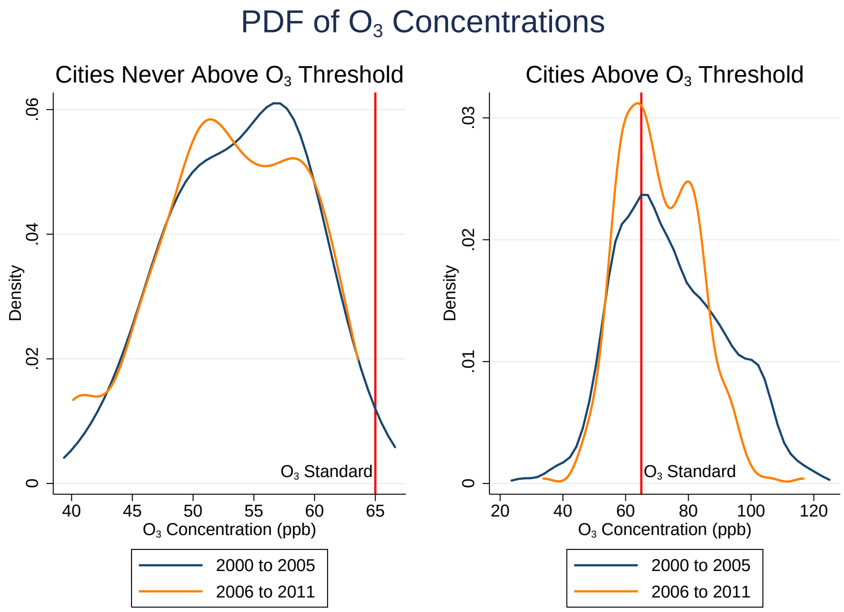

Figure 2.5: Kernel density estimates of the distribution of O_3 concentrations across CMAs in the first half (2000-2005) and second half (2006-2011) of the CWS phase in period. The panel on the left displays pollution concentrations for CMAs that never exceeded the O_3 standard. The right panel displays the pollution concentrations for CMAs that exceeded the O_3 standard at least once. The vertical red lines represents the threshold used for the O_3 standard.

A potential concern with these figures is that they could merely reflect different trends across regions owing to other factors beyond the CWS. A primary concern is that the CMAs exceeding one of these thresholds may be more heavily populated or industrialized than those below the threshold. To show this is not the driver of the documented change in air quality, in Figure 2.6 we show the change in mean $PM_{2.5}$ (panel (a)) and O_3 (panel (b)) concentrations for CMAs with a population of at least 300,000 people. As above, the panel on the left displays pollution concentrations of CMAs that never exceed the relevant standard, and the right shows concentrations of the dirty CMAs. The figure shows a pronounced drop in both $PM_{2.5}$ and O_3 air quality for the heavily populated regions that exceed the respective thresholds, and no change in air quality for the clean CMAs.

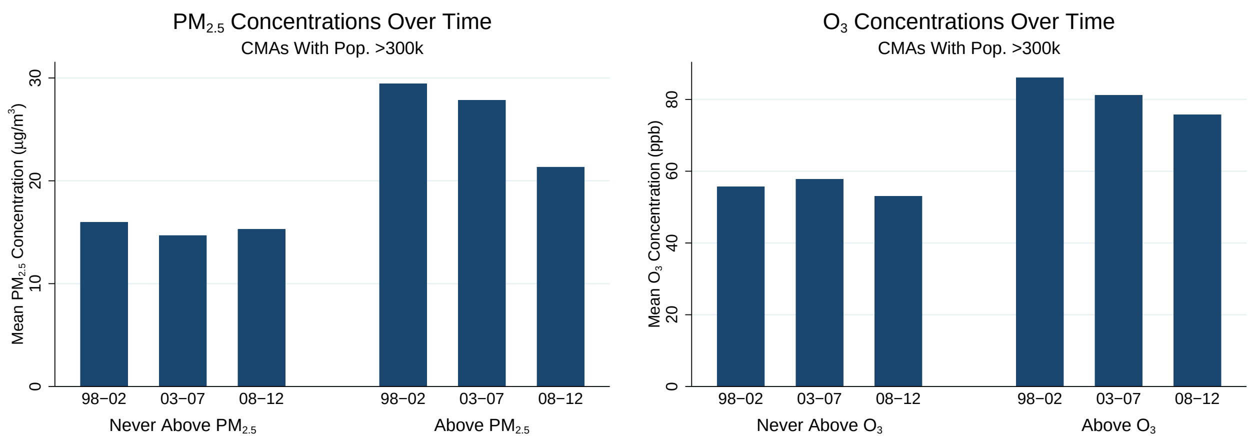

(a) $PM_{2.5}$ Concentrations

(b) O_3 Concentrations

Figure 2.6: Mean Pollution Concentrations by Year for Large CMAs

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

While there are potentially other factors that may explain the changes in air quality shown above, a proper treatment of the effect of the CWS on air quality is beyond the scope of this chapter. As our goal is to estimate the effects of the CWS on manufacturing plants, such a regional trend would only represent a potential identification problem if it were specific to the industries targeted by the CWS. This is because the cross-industry variation in regulatory stringency allows us to flexibly control for CMA trends.

2.4 Empirics

In Chapter 1, we presented a theoretical model of how a targeted environmental regulation, such as the CWS, may affect plants. This model provided a number of clear predictions as to how facilities would respond to the CWS. Taken together, these results imply that when the fixed costs of abatement are high, environmental regulations should primarily reduce industry emission intensity via reallocation and selection effects. In contrast, when the fixed costs of abatement are low, the industry clean-up should be driven by process effects. In this section, we explore those plant-level predictions empirically by estimating the CWS' effect on plant pollution intensity, production, and exit. We use the resulting estimates to determine how the process, reallocation, and selection effects created by the CWS have contributed to the clean-up of Canadian manufacturing.

2.4.1 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.[24]

Our design begins by comparing the average outcomes of plants in regulated CMAs while regulated (i.e. while violating one of the standards) to their average outcomes while unregulated. This allows us to control for any unobserved time-invariant industry, CMA or plant characteristics that would affect plant pollution emissions. Moreover, in the absence of any other shocks, this comparison would identify the average causal effect of the CWS on pollution emissions. Yet, such absence is unlikely; there is strong reason to believe that a simple before-and-after comparison of affected plants could also capture the effects of regional, industry, or aggregate economic shocks.[25] We discuss each in turn.

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

The simple before-and-after comparison could also be contaminated as a result of economic shocks that affect individual industries, which could arise due to the effects of increased foreign competition created by international trade, or by revisions to federal policies that target certain sectors. To address these issues we exploit cross-CMA variation in regulation, and

[24] It is worth mentioning that, while plants in dirty CMAs that were operating in a targeted industry were subject to more strict regulation and enforcement, it is possible that other plants in the country were regulated to some degree as a result of the CWS. If this is the case, then our research design produces estimates that give a lower bound on the CWS' effects on the manufacturing sector. We describe these potential policies in the appendix (see appendix A.2.5).

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

utilize the fact that in any particular industry, only plants in areas with poor air quality were subject to stringent environmental policy. This allows us to use the average outcomes from plants in a targeted industry in an unregulated CMA as a counterfactual for the average outcomes of plants from that industry that are located in a regulated CMA. This captures the effects of any industry specific shocks.

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

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

$$y_{pict} = \beta_{PM} T_{ict}^{PM} + \beta_{O3} T_{ict}^{O3} + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \tag{2.2}$$

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.[26] T_{ict}^{j} is an indicator of treatment for standard j, and takes a value of one for plants that are in industries targeted by the CWS for years in which their CMA exceeds threshold j.

Equation (2.2) also includes plant (ρ_p), CMA-year (ξ_{ct}), industry-year[27] (λ_{it}) fixed effects and an error term (ε_{pict}). The plant fixed effects account for any unobserved plant-specific heterogeneity, as well as time-invariant industry and CMA characteristics. The CMA-year fixed effects capture any region specific shocks. The industry-year fixed effects account for any industry-wide events. Finally, the error term captures idiosyncratic changes in outcomes across plants.

The coefficients of interest in Equation (2.2) 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

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

[27] The CWS defined the targeted industries at the 3- or 4-digit North American Industry Classification System (NAICS) level.[28] We create an industry indicator that corresponds to either the 3- or 4-digit NAICS level. All 3-digit industries that contain targeted industries defined at the 4-digit level are grouped at the 4-digit level. The remaining industries are grouped at the 3-digit level.

identified from within plant comparisons over time.[29,30]

Changes in plant regulatory status must be plausibly exogenous for this research design to credibly identify the effects of the CWS. There are two reasons to believe this is the case. First, as with the CAAAs in the US, regulations are determined by a nationally set air quality threshold that does not vary over time. As a result, these standards are unrelated to differences in local tastes, characteristics or economic conditions (Greenstone, 2002), an issue that could arise with any region-specific policy (Besley and Case, 2000). Second, $PM_{2.5}$ 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.[31] As variation in regional air quality determines assignment to treatment, this means it is unlikely a plant could manipulate their regulatory status. 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). The transboundary nature of pollution means it is unlikely a single plant can directly manipulate their treatment status.

2.4.2 Data and Measurement

Our analysis relies on a unique confidential micro-dataset that contains information on the $PM_{2.5}$ and NO_X emission intensity of Canadian manufacturing plants. This dataset was created by merging data from two existing sources: the National Pollutant Release Inventory (NPRI) and the Annual Survey of Manufactures (ASM).[32] The NPRI contains information on the emissions of various pollutants from Canadian manufacturing plants. By law, any facility that emits one of the covered pollutants above a minimum threshold must report to the NPRI. The ASM was used as Statistics Canada's manufacturing census until 2012, and it provides

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

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

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

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

longitudinal information on plant sales, production costs, employment, and other plant characteristics for the majority of manufacturing plants in Canada.[33] Plants in these two datasets were linked by Statistics Canada, allowing us to create a longitudinal dataset containing information on $PM_{2.5}$ and NO_X emission intensity as well as other plant characteristics over the period 2004-2010. Additional details on each data source and the construction of the dataset used in our analysis are given in Appendix A.2.1.

While there are several additional datasets that contain linked plant-level pollution and production information, the NPRI-ASM has a number of advantages relative to these other sources. A commonly used approach to compile this type of data is to match publicly available plant-level pollution data produced by the Environmental Protection Agency with the National Establishment Time Series (NETS), a proprietary dataset with plant employment and sales. For example, Cherniwchan (2017), Holladay (2016), and Cui et al. (2015) all used this type of data in their work on trade and pollution.[34] A clear advantage of our data relative to this proprietary dataset is that the NPRI-ASM is produced using Canada's official manufacturing census, which derives its core information from tax files and representative surveys. In contrast, concerns have been raised over how representative the NETS is of the actual universe of U.S. manufacturing facilities (see, e.g., Barnatchez et al. (2017); Haltiwanger et al. (2013)).[35]

Descriptive statistics for the key variables that we employ are reported in Table 2.2. Each column in Table 2.2 presents averages and standard deviations for a different sample corresponding to emitters of each pollutant. The first column corresponds to the set of plants that emit $PM_{2.5}$, the second column shows statistics for plants that emit NO_X, and the final column of the table reports summary statistics for the entire sample of plants in the ASM. The statistics in columns one and two are weighted to account for potential sample bias induced by the linking procedure used to match plants across datasets (see Appendix A.2.1 for further details). Each sample is an unbalanced panel; the sample for $PM_{2.5}$ contains 6501 plant-year observations, the sample for NO_X contains 3012 plant-year observations, and the full sample contains 309,541 plant-year observations.

The summary statistics reported in Table 2.2 suggests that there are systematic differences in plants that emit different types of pollutants. For example, on average, the NO_X sample

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

[34]In addition, (Tang et al., 2015) utilize a dataset produced by the National Bureau of Statistics of China that contains firm-, but not facility-, level pollution and production data, although this data is cross-sectional.

[35]Others have used datasets that contain plant-level production and fuel-use data (Barrows and Ollivier, 2018; Martin et al., 2014). While fuel-use is an important input into pollution, having this data alone does not allow one to fully capture the process effect.

Table 2.2: Summary Statistics

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

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

emitted more pollution, produced more output, had higher employment levels, and had higher labour-productivity levels than the $PM_{2.5}$ sample. This potentially reflects substantial differences in how pollution is produced and abated, given that pollutants are typically produced by a few industries (Greenstone, 2002), and there are substantial differences in the fixed costs of abatement across pollutants (Canadian Council of Ministers of the Environment, 1998b; Environment Canada, 2002).

Table 2.2 also shows that polluters represent the largest plants in the manufacturing sector. Relative to the full manufacturing sector, the sample of plants that emit either $PM_{2.5}$ or NO_X sell more goods (15 to 30 times on average), employ more workers (7 to 10 times), and have higher value added per worker (2 to 3 times) than the average manufacturing plant.[36] This is, in part, due to the reporting requirements for the NPRI; by law, plants only report if they emit at least one covered pollutant above a minimum threshold level and employ at least 10 individuals or operate an on-site generator (Environment and Climate Change Canada, 2016c). While this means we systematically exclude small facilities, our analysis covers plants that account for the majority of manufacturing pollution in Canada.[37]

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

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

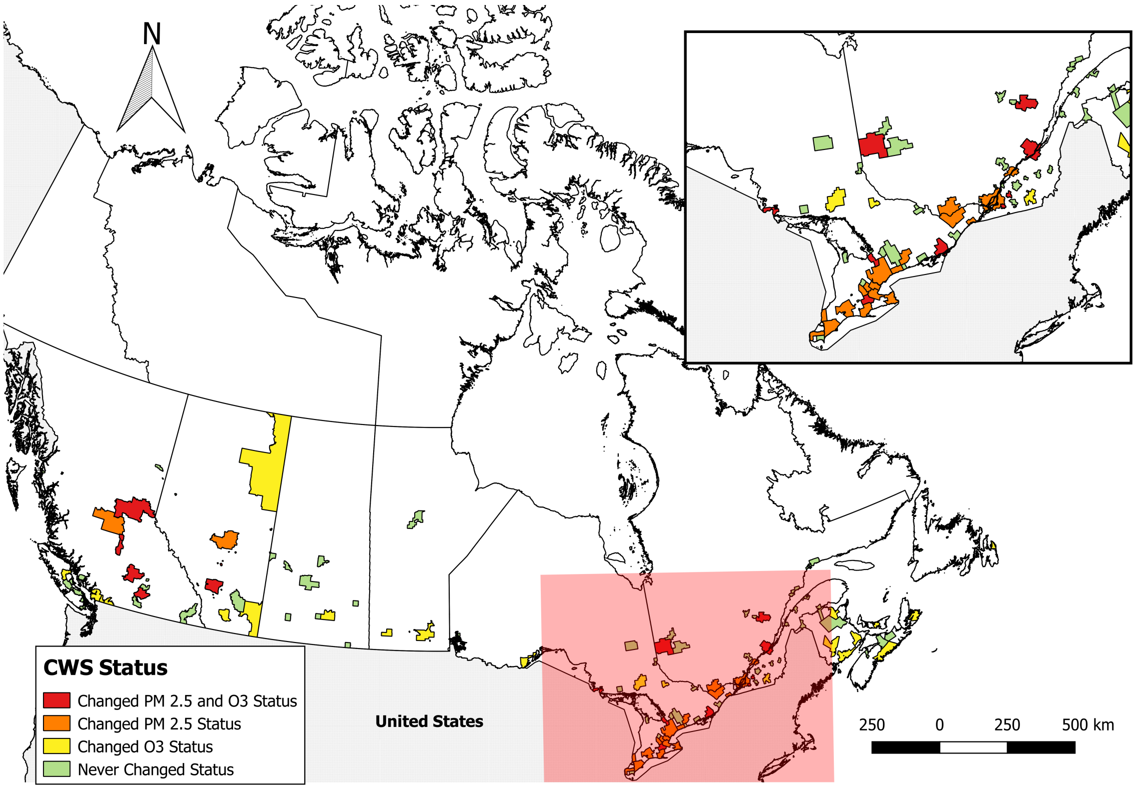

Figure 2.7: Regulatory Status Changes under the CWS

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

Determining Regulatory Status under the CWS

Our analysis also requires determining which CMAs were affected by the CWS. To do so, we use the local air quality information described in Section 2.3.1. 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.

The variation in regulatory status created by changes in ambient air quality is illustrated in Figure 2.7, which shows the CMAs that changed regulatory status for the $PM_{2.5}$ and O_3 standards. In Figure 2.7, the red CMAs changed status under both the $PM_{2.5}$ and O_3 standards, the orange CMAs only changed status for the $PM_{2.5}$ standard, the yellow CMAs only changed status for the O_3 standard, and the green CMAs didn't change status under either standard. As the figure shows, there was substantial variation in which CMAs changed their regulatory status over the 2000-2010 period. Of the 149 CMAs in our sample, 23% changed status under the $PM_{2.5}$ standard, 26% changed status under the O_3 standard, 11% changed status under both standards, and 60% never changed regulatory status.

2.4.3 Empirical Results

The CWS and Plant Pollution Emissions

We begin our analysis by estimating the effects of the CWS on the level of pollution emitted by affected Canadian manufacturing plants.[38] We start here for two reasons. First, it provides some indication as to the effectiveness of the CWS; if the regulations were responsible for the reduction in pollution levels documented in Section 2.2, then we should observe reductions in the emissions of targeted pollutants as a result of the CWS. Second, this also provides us with a means to assess the external validity of our results. As we discussed above, there is little evidence as to the effects of environmental regulation on the emission intensity of manufacturing plants. Focusing on pollution levels allows us to directly compare the effects of the CWS with the effects of other environmental policies.

Table 2.3 reports our estimates of the effects of the CWS on plant pollution emissions. We estimate Equation (2.2) 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.[39] The first column of each panel reports estimates from a version of Equation (2.2) that only includes the particulate matter standard. Similarly, the second column reports estimates from a specification that only includes the ozone standard. Finally, column (3) in each panel reports estimates from the specification given in Equation (2.2). The first row in each panel reports the effect of the $PM_{2.5}$ standard (β_{PM} in Equation (2.2)); the second row shows the effect of the O_3 standard (β_{O3} in Equation (2.2)). The dependent variable in each of these regressions is the natural log of plant pollution emissions for the relevant pollutant. Each regression is weighted to correct for potential sample bias introduced by the procedure used to match plants in the NPRI with plants in the ASM.[40] In all cases, standard errors clustered at the CMA-industry level are reported in parentheses.

[38] A related, but distinct question, is to ask what the CWS did to regional air quality. While this is beyond the scope of this chapter, in the appendix we provide descriptive evidence that air quality improved in Canada over this period (see Section 2.3.1).

[39] There are other pollutants that may also contribute to $PM_{2.5}$ and O_3 pollution, including volatile organic compounds and carbon monoxide. In the appendix we examine the CWS' effects on the emissions of a number of other pollutants (see Appendix A.2.3).

[40] In brief, the potential bias happens because the probability of a successful match is positively correlated with a plant's size. If the effects of the CWS vary by plant-size, then relying on the matched data would produce bias estimates. Details on the weighting procedure used to address this can be found in Appendix A.2.1.

Table 2.3: The Effects of the CWS on Plant Pollution Emissions

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

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

The estimates reported in Table 2.3 indicate that the CWS regulations led to statistically significant reductions in the emissions of both particulate matter and nitrogen oxide from affected plants. Our baseline estimates for $PM_{2.5}$, reported in column (3) of Panel A, indicate that the CWS particulate matter regulations are associated with a 15.1% reduction in emissions from affected plants. Our baseline estimates for NO_X, reported in column (6) of Panel B indicate that the ozone regulations are associated with 32.5% decrease in emissions from affected plants. The estimates reported in Panels A and B also show no statistically significant cross-effects of either standard. That is, O_3 regulation did not significantly affect particulate matter emissions and PM regulation did not significantly affect NO_X emissions.

We view the results in Table 2.3 as an exploratory analysis of the CWS' effects on plants. While the effect of O_3 on NO_X emitters is only marginally significant, we call attention to these estimates because, as we show later in this section, the average effects of the CWS mask considerable heterogeneity across plants (see Section 2.4.3). Taken together, our evidence suggests O_3 regulation had a meaningful effect on manufacturing plants. Moreover, the theory presented in Chapter 1 suggests an average treatment effect is not very illustrative of how plants respond to a policy such as the CWS.

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

2014.[41] This suggests that the CWS had similar effects on pollution levels as the environmental policies enacted elsewhere.

It is also worth noting that in Section 2.4.3 we present evidence that the estimates reported in Table 2.3 are robust to a number of potential identification concerns, including the potential for pre-trends and the effects of a negative relationship between a CMA's air quality and the production choices of the plants therein.

The CWS and the Clean-up of Manufacturing

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

Plant-Level Estimates

In Table 2.4 we report our estimates of the CWS' effect on the emission intensity of manufacturing plants. As in Table 2.3, panel A shows estimates of Equation (2.2) for the sample of plants that emit $PM_{2.5}$ and panel B shows estimates for the NO_X emitters. In each panel, we report estimates from two separate regressions each with a different measure of emission intensity, as well as reproducing our baseline estimates of the CWS' effects on plant pollution levels. The first column shows the CWS effect on pollution levels. In the second column, we show the CWS' effects on emission intensity, measured as the ratio of emissions to total plant shipments (sales), given this is the measure of output used previously in the literature documenting the manufacturing clean-up. In the third column, we measure emission intensity as the ratio of emissions to value-added. Value added may provide a more accurate reflection of the level of productive activity that occurs in each plant (Cherniwchan et al., 2017). However, we focus on the estimates in the second column of each panel, as our goal is to contribute to a literature that uses shipments as its measure of output.[42] In both cases, the dependent variable is the natural log of emission intensity. The first row in each panel reports the effect of the $PM_{2.5}$ regulation (β_{PM} in Equation (2.2)) and the second row reports the effect of the O_3 regulation (β_{O3} in Equation (2.2)). As before, each regression is weighted to correct for potential

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

[42] 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.

Table 2.4: The Effects of the CWS on Plant Emission Intensity

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

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

bias from matching the NPRI and ASM, while standard errors clustered by CMA-industry are reported in parentheses.

The estimates reported in Panel A of Table 2.4 indicate $PM_{2.5}$ regulation had little-to-no-effect on the emission intensity of plants that emitted $PM_{2.5}$, with an estimated coefficient in column (2) that is relatively small and statistically insignificant. In contrast, the CWS O_3 regulations appear to have caused a significant reduction in NO_X pollution intensity. The estimate reported in column (5) of Table 2.4 indicate that the CWS ozone regulations are associated with a 28.6% decrease in the level of NO_X emitted per unit of output.[43]

In addition, $PM_{2.5}$ regulation caused a significant increase in NO_X intensity measured in value added terms. These results are driven by a very small number of plants that are regulated by the $PM_{2.5}$ standard and emit NO_X, but not $PM_{2.5}$. For these plants, $PM_{2.5}$ regulation caused a large increase in NO_X emissions and decrease in value added. We do not probe these findings further, as they are driven by fewer than ten plants.[44]

$PM_{2.5}$ regulation caused a sizable reduction in plant $PM_{2.5}$ emissions, but had no significant effect on plant emission intensities. On the other hand, the O_3 standard caused a large reduction in NO_X emissions in both levels and pollution intensity. This implies the $PM_{2.5}$ standard must have led to large decreases in output from affected plants, whereas the ozone standard

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

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

Table 2.5: The Effects of the CWS on Plant Output

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

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

had relatively minor affects on output. We confirm these conclusions by directly estimating Equation (2.2) on both sales and value added.

Estimates of the effects of the CWS on plant output are given in Table 2.5 for both $PM_{2.5}$ (Panel A) and NO_X (Panel B) emitters, with each panel reporting estimates from two separate regressions. In the first, we measure output as the value of total plant shipments (sales), and in the second as value added. In both cases, the dependent variable is the natural log of output. The first row in each panel reports the effect of the $PM_{2.5}$ regulation (β_{PM} in Equation (2.2)) and the second row reports the effect of the O_3 regulation (β_{O3} in Equation (2.2)). As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure, and standard errors clustered by CMA-industry are reported in parentheses.

The estimates reported in Panel A of Table 2.5 confirm the $PM_{2.5}$ standard led to a large decrease in output from affected plants that emitted particulate matter. The estimate in column (1) of Panel A indicate the CWS particulate matter regulation is associated with a 10.8% decrease in sales from plants that emitted $PM_{2.5}$. Conversely, the estimates in panel B show the O_3 standard had no statistically significant effects on output.[45]

Lastly, we estimate a variant of our main specification (Equation (2.2)) in which we compare the number of plants operating in a treated industry-CMA-year cell to the number oper-

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

Table 2.6: The Effects of the CWS on Plant Exit

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

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

ating in an untreated industry-CMA-year cell. That is, we estimate the following regression

$$N_{ict} = \beta_{PM} T_{ict}^{PM} + \beta_{O3} T_{ict}^{O3} + \alpha I(\text{CWS})_{ic} + \xi_{ct} + \lambda_{it} + \varepsilon_{ict}, \tag{2.3}$$

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

As the dependent variable is a count variable, we estimate Equation (2.3) using both ordinary least squares and Poisson regression. As above, we report estimates for two groups of plants: those that emit $PM_{2.5}$ (Panel A) and those that emit NO_X (Panel B). These results are presented in Table 2.6, which includes standard errors clustered by CMA's in parentheses.

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

Aggregate Implications

The main implication of the results presented in Table 2.4, Table 2.5, and Table 2.6 is that the CWS contributed to the manufacturing clean-up through different channels for different pollutants. The particulate matter standard primarily caused a reduction in output at regulated plants and plants to exit. In contrast, the ozone standard caused regulated plants to adopt cleaner processes. To quantify the total contribution of the CWS to the manufacturing clean-up we present a simple counterfactual exercise in which we ask how much of the clean-up can be attributed to the process, reallocation and selection effects induced by the CWS. We do this by using our estimates, paired with a decomposition of an industry's emission intensity, to compute the implied change in manufacturing pollution intensity over our sample that occurred because of each of the CWS channels. We then compare these estimates to the observed change in manufacturing pollution intensity.[46]

For our decomposition, we follow an approach used in much of the labor literature and consider total changes in emission intensity over time (for a relevant review, see Foster et al. (2001)). This exercise is, in effect, an extension of the decomposition presented in Section 2.2 that documented how industry-level changes contribute to a change in a sector's pollution intensity. In contrast, this decomposition documents how plant-level changes contribute to a change in an industry's pollution intensity. This decomposition can also be thought of as an empirical analogue to the decomposition in Section 1.2 of Chapter 1.

To that end, let output and pollution from manufacturing industry i in year t be given by X_{it} and Z_{it}, respectively. We define an industry's pollution intensity as the amount of pollution emitted per unit of output produced, and let this be given by $E_{it} = Z_{it}/X_{it}$. In addition, suppose each industry is composed of a continuum of plants and let $x_{it}(n)$, $z_{it}(n)$, and $e_{it}(n)$ denote output, pollution, and pollution intensity from plant n. Lastly, let $\lambda_{it}(n) = x_{it}(n)/X_{it}$ be plant n's share of production in industry i and year t, and n_i denote the marginal plant that is endogenously determined by the industry's profitability.[47] In this case, the emission intensity of industry i at time t can be expressed as $E_{it} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn$. Assuming, for convenience, that plants only exit the industry over time and never enter, then the change in an industry's emission intensity from $t-1$ to t is given by

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

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

[47]As in Cherniwchan et al. (2017), we assume plants are ranked in reverse order of productivity. Consequently, selection removes the least productive plants.

In Appendix A.2.4, we show that the percentage change in an industry's emission intensity, $\dot{E}_{it} = \frac{E_{it} - E_{it-1}}{E_{it-1}}$, can then be expressed as

$$\begin{aligned} \dot{E}_{it} = & \int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{\lambda}_{it}(n)dn \\ & - \int_{n_{it}}^{n_{it-1}} s_{zit-1}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)\dot{\lambda}_{it}(n)dn, \end{aligned} \tag{2.4}$$

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 term on the right-hand side of equation (2.4) is the "process effect". This captures the change in industry emission intensity due to changes in plant emission intensity over time resulting from the adoption of new production processes. The second term on the right-hand side of equation (1.2), captures the effects changes in the relative size of plants within an industry over time, which we call the "reallocation effect". 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 over time owing to plant exit. The final term is an interaction effect created by the interaction between the process and reallocation effects.[48]

We use our estimates presented above in Section 2.4.3 to construct the four terms on the left-hand side of Equation (2.4). As such, let $\hat{\beta}_e$, $\hat{\beta}_x$, and $\hat{\beta}_n$ denote our estimates of the effects of the CWS on plant pollution intensity (from Table 2.4), plant output (from Table 2.5), and selection (from Table 2.6), respectively. Moreover, recall that, given our identification assumptions, $\hat{\beta}_e$ captures the average change in emission intensity due to the CWS, meaning that we can write

$$\dot{e}_{it}(n) = \begin{cases} \hat{\beta}_e, & \text{if } n \text{ is treated} \\ 0, & \text{otherwise.} \end{cases} \tag{2.5}$$

In addition, an estimate of $\dot{\lambda}_{it}(n)$ and $\int_{n_{it}}^{n_{it-1}} s^z_{it-1}(n)dn$ can be constructed from $\hat{\beta}_x$ and $\hat{\beta}_n$, respectively. In the appendix (see Appendix A.2.4), we show that

$$\dot{\lambda}_{it}(n) = \begin{cases} \frac{\hat{\beta}_x(1 - s^{Treat}_{xit-1}) + s^{Exit}_{xit-1}}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}}, & \text{if } n \text{ is treated} \\ \frac{s^{Exit}_{xit-1} - \hat{\beta}_x s^{Treat}_{xit-1}}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}}, & \text{otherwise,} \end{cases} \tag{2.6}$$

[48]Note that this can be thought of as the approximation error in the decomposition presented in Section 1.2 of Chapter 1 caused by focusing on small, rather than potentially large, changes.

where s^{Treat}_{xit-1} and s^{Exit}_{xit-1} are the fraction of output in time $t-1$ from treated and exiting plants, respectively. Substituting Equation (2.5) and Equation (2.6) into Equation (2.4) gives estimates of the process, reallocation, and interaction effects. Letting s^{Treat}_{zit-1} be the share of industry i's pollution in time $t-1$ from treated plants, then the process effect is

$$\widehat{PE} = \hat{\beta}_e s^{Treat}_{zit-1}. \tag{2.7}$$

Similarly, the reallocation effect is given by

$$\widehat{RE} = \frac{s^{Exit}_{xit-1} + \hat{\beta}_x (s^{Treat}_{zit-1} - s^{Treat}_{xit-1})}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}}, \tag{2.8}$$

and the interaction effect is given by

$$\widehat{IE} = \hat{\beta}_e s^{Treat}_{zit-1} \left[\frac{\hat{\beta}_x (1 - s^{Treat}_{xit-1}) + s^{Exit}_{xit-1}}{1 - s^{Exit}_{xit-1} + \hat{\beta}_x s^{Treat}_{xit-1}} \right] \tag{2.9}$$

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

$$\widehat{SE} = \hat{\beta}_n N^{Treat} \bar{s}^{Exit}_{zit-1}, \tag{2.10}$$

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

In Table 2.7 we present our estimates of each of the CWS channels relative to the observed change in manufacturing pollution intensity from 2004 to 2010. The first row shows the fraction of the $PM_{2.5}$ clean-up due to the CWS and the second shows the fraction of the NO_X clean-up due to the CWS. Our estimates of the process effect, reallocation effect, selection effect and interaction effect for each pollutant are reported in columns (1)-(4), respectively. Column (5) reports the implied change in manufacturing pollution intensity that can be explained by the CWS.

The results of this exercise show that both the $PM_{2.5}$ and O_3 standards enacted under the CWS played a considerable role in the clean-up of Canadian manufacturing. The estimates in column (5) show that, from 2004 to 2010, the O_3 standard is responsible for 61% of the reduction in manufacturing NO_X intensity and the $PM_{2.5}$ standard is responsible for 21% of the reduction in manufacturing $PM_{2.5}$ intensity. However, the channels responsible varied considerably across pollutants. The process effect, for example, associated with NO_X regula-

Table 2.7: Counterfactual Estimates

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

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

tion 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.

Explaining How Industries Clean Up

The results presented above show that the channels through which the CWS caused the manufacturing sector to clean up varied across pollutants. The theoretical model presented in Chapter 1 provides a potential explanation for this: differences in the fixed costs of adopting cleaner production processes across pollutants. Indeed, as we discussed in Section 2.3, engineering assessments of these pollutants argue abatement of NO_X can be accomplished at low-cost, while abatement of $PM_{2.5}$ pollution typically requires large fixed costs (Canadian Council of Ministers of the Environment, 1998b; Environment Canada, 2002; Environmental Protection Agency, 1999a). We now turn to assess this mechanism further and examine other potential explanations for our findings.

Differential Effects by Plant Productivity Level

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

To test this prediction, we use an approach similar in spirit to that of Bustos (2011) and allow the effects of the CWS to differ across plants on the basis of their initial productivity level. An obvious limitation with our data is that we do not observe a plant's capital stock information, thereby preventing us from being able to estimate plant total factor productivity. Instead, we use value added per worker (that is, labor productivity) as our productivity measure. As productivity is potentially affected by regulation, we use a plant's initial labor productivity in the first year they enter our dataset to avoid contamination with the CWS. To account for potential differences in average productivity levels across industries and time, we regress these plant initial productivity levels on entry-year and industry fixed effects, and use the residuals from this regression as our measure of plant productivity. For consistency, we use the same industry definitions we employed in constructing our industry fixed-effects.[49]

To avoid functional form assumptions, we use a non-parametric approach and assign plants into one of three productivity bins, according to where they lie on the productivity distribution. We then interact these bins with the CWS treatment indicators. The bins are constructed by first sorting plants in each sample into terciles based on their residual initial labor productivity. These bins are then used to create indicators Q_q, which indicate whether a plant is in the bottom- (Q_1), middle- (Q_2), or top- (Q_3) third of the productivity distribution. 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.

We accomplish this by estimating the following regression

$$Y_{pict} = \sum_{q=1}^{3} \beta_{PM}^{Q_q}[T_{ict}^{PM} \times Q_q] + \sum_{q=1}^{3} \beta_{O3}^{Q_q}[T_{ict}^{O3} \times Q_q] + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \tag{2.11}$$

where Q_q is the indicator that takes the value one if plants that are in productivity tercile q, T_{ict}^{j} takes a value of one for all plants in targeted industries for years in which their CMA violates standard j, $\beta_j^{Q_q}$ is the treatment effect of standard j on plants in productivity tercile q, and the remaining variables are as defined for Equation (2.2).

We use Equation (2.11) to examine the CWS' effects on plant pollution levels, emission intensity, and sales. Examining pollution levels allows us to assess whether the CWS affected emissions from plants of all productivity levels, whereas examining emission intensity and sales allows us to quantify the channels by which regulation affected each plant.

These results are shown in Table 2.8. Panel A reports our estimates for $PM_{2.5}$ emitters; Panel B for NO_X emitters. In each panel, we report estimates from three separate regressions. The first column in each panel shows our estimates from Equation (2.11) on plant emissions, the second column shows the effects on plant emissions per dollar of sales, and the third on plant sales. Natural logarithms are taken of all dependent variables. The first three rows in each panel report the effects of the $PM_{2.5}$ regulation (the $\beta_{PM}^{Q_q}$ coefficients in Equation (2.11)). The first row shows the effect on plants in the lowest productivity tercile, the second row the effect on plants in the middle tercile, and the third row the effects on plants in the highest

[49]See Section 2.4.1 for more details on the industry definitions employed.

Table 2.8: The Effects of the CWS by Plant Productivity Level

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

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

tercile. Similarly, the final three rows report the effects of the O_3 regulation ($\beta_{O3}^{Q_q}$ in Equation (2.11)). The fourth row shows the effect on plants in the lowest productivity tercile, the fifth row the effect on plants in the middle tercile, and the sixth row the effects on plants in the highest tercile. As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

The results for the $PM_{2.5}$ standard show stark differences across $PM_{2.5}$ plants of different productivity levels. $PM_{2.5}$ regulation caused a drop in emissions among the bottom two-thirds of the productivity distribution, with a reduction in emissions of 16.3% for low productivity plants and 27.9% for middle productivity $PM_{2.5}$ emitters. In contrast, $PM_{2.5}$ regulation had no significant effect on the most productive $PM_{2.5}$ plants, suggesting they were unaffected by regulation.

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

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

The results in Table 2.8 are consistent with our theory, and the hypothesis that the channels of the CWS clean-ups varied across pollutants because of differences in abatement costs.[50] As the abatement of $PM_{2.5}$ requires paying a relatively high fixed cost, only relatively productive plants should choose to do so. These highly productive plants, in turn, experience a reduction in pollution intensity with a relatively small change in output and production inputs. The less productive plants, on the other hand, experience an increase in production costs, leading to a reduction in input use, output, and productivity. In contrast, as NO_X can be abated at a relatively low cost, there are smaller differences across plants of different productivity levels. For both pollutants, the most productive plants in an industry use state of the art technology, and are thus unaffected by the CWS.

Other Margins of Plant Adjustment

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

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

Thus far, the hypothesis we have focused on is that $PM_{2.5}$ and NO_X have different abatement costs, which affects a plant's willingness to adopt cleaner production processes. An alternative hypothesis is that the opportunities for input substitution may vary across pollutants. For example, there could be readily available alternatives to the inputs that create NO_X pollution, but not for the inputs that create $PM_{2.5}$ pollution. If this were the case, then regulation would reduce NO_X intensity but not $PM_{2.5}$ intensity.

Examining the effect of the CWS on input use allows us to asses the above hypothesis. If this hypothesis were true, then the CWS should have caused an increase in spending on inputs for NO_X emitters.[51] 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.[52] If the average plant became less productive in response to $PM_{2.5}$ regulation, but more productive in response to NO_X regulation, then this could generate the findings reported in Section 2.4.3. Examining the effect of the CWS on plant productivity allows us to test this hypothesis.

We examine these alternative hypotheses using data on the total number of plant employees[53], 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 2.9. Panel A shows estimates of Equation (2.2) 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

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

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

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

Table 2.9: Other Margins of Plant Adjustment

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

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

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

an indicator for whether the plant is involved in research and development using a linear probability model. In each specification, the first row reports the effect of the $PM_{2.5}$ regulation and the second row reports the effect of the O_3 regulation. As before, each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

We also examine if the effects of the CWS on productivity and input use differ across the initial plant productivity distribution. These estimates are reported in Table 2.10.[54] Panel

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

A shows the results for $PM_{2.5}$ emitters and Panel B for NO_X emitters. Each column in each panel corresponds to a different dependent variable, each measured in natural logarithms. Each regression is weighted to correct for potential bias from the NPRI-ASM matching procedure. In all cases, standard errors clustered by CMA-industry are reported in parentheses.

As the estimates reported in Table 2.9 and Table 2.10 show, the main channels by which the average $PM_{2.5}$ emitting plant responded to $PM_{2.5}$ regulation appears to be through changes in intermediate input use and labor productivity. $PM_{2.5}$ regulation decreased spending on production materials by 11.9%, caused a drop in energy spending (although not significant at conventional levels), and reduced labor productivity (also not significant at conventional levels). $PM_{2.5}$ regulation also caused a significant reduction in labor productivity among 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 2.9 and Table 2.10 also suggest O_3 regulation did not have a significant effect on input use, employment, labor productivity, or R&D propensity at the average affected plant. The exception to this is an increase in energy spending among $PM_{2.5}$ emitters. Allowing the effects of the CWS to vary across plant productivity levels, we still find no significant effect on NO_X emitter employment, input spending, or labor productivity. These results are inconsistent with the two additional hypotheses described above, as neither productivity nor input spending rise in response to regulation, which further suggests our results are driven by the fixed costs of abatement.

Table 2.10: CWS Mechanisms by Plant Productivity Level

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

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

Robustness of the CWS

In this section we present results of a several exercises we perform to examine the robustness of our main findings. In the interest of space, we only provide the estimation results for the average effects of the CWS on emissions.[55]

We begin assessing the robustness of these results with a series of placebo tests. The intent of these tests are to ensure our research design does not produce significant differences in plant emissions when plants are "randomly assigned" into (potentially) false treated and control groups. As our research design exploits three margins of variation – across industries, regions, and time – we present two separate series of placebo tests. In the first series of tests, we keep the same variation in regulatory exposure across regions and time, and randomize the plants that are in targeted and non-targeted industries. This allows us to ensure that the effect of a plant being in a violating region only matters for plants that are in targeted industries. In the second series of tests, we keep the same variation in regulatory exposure across industries, and randomize the plant-years that are in violating and non-violating regions. This allows us to ensure that the effect of a plant being in a targeted industry only matters for plants in violating regions.

The results of these placebo tests are presented in Table 2.11 and Table 2.12, for $PM_{2.5}$ and NO_X emitters, respectively. For $PM_{2.5}$ emitters we only show the effect of (placebo) $PM_{2.5}$ regulation, and for NO_X emitters we only show the effect of (placebo) O_3 regulation. The dependent variable in each regression is the natural log of plant emissions. Panel A of each table reports estimates from our first series of placebo tests, randomizing the plants that are assigned to targeted industries. Panel B of each table reports estimates from our second series of placebo tests, randomizing plant-years that are assigned to violating regions. Within each panel, we display estimates of specifications with different fractions of plants assigned to targeted industries or violating regions. All regressions include plant, industry-year, and CMA-year fixed effects, and standard errors are clustered by CMA-industry.

The results in Table 2.11 and Table 2.12 show no significant effects of these placebo regulations on plant emissions. This suggests the estimates presented in Section 2.4.3 are not simply driven by the structure of the research design.

Next, we examine whether our results are simply capturing the effects of a non-linear relationship between CMA air quality and the production choices of plants therein.[56] We do this

[55]Each robustness check we also performed for the CWS' effects on output and by plant-productivity level.

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

Table 2.11: CWS Placebo Tests - $PM_{2.5}$ Emissions

	Panel A: Within-CMA Placebo			Panel B: Within-Industry Placebo		
	(1)	(2)	(3)	(4)	(5)	(6)
Placebo $PM_{2.5}$ Reg.	-0.015	0.010	-0.033	0.024	0.015	0.045
	(0.036)	(0.046)	(0.042)	(0.033)	(0.042)	(0.037)
R^2	0.938	0.938	0.938	0.938	0.938	0.938
N	7058	7058	7058	7058	7058	7058
Fraction of plants in targeted industry	0.50	0.80	0.20			
Fraction of plants in violating region				0.50	0.80	0.20

Notes: Table reports estimates of placebo tests of $PM_{2.5}$ regulation's effect on plant $PM_{2.5}$ emissions. Panel A reports estimates from placebo tests that randomize the plants that are assigned to targeted industries, but preserves the actual variation in the CMA-years that violate the $PM_{2.5}$ regulation. Each column within Panel A assigns a different fraction of plants into the targeted industries. The Actual fraction of plants in targeted industries is 0.49. Panel B reports estimates from placebo tests that randomize the plants that are assigned to violating regions, but preserves the actual variation in the industries that are targeted. Each column within Panel B assigns a different fraction of plants into the violating regions. The Actual fraction of plants in violating regions is 0.18. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

Table 2.12: CWS Placebo Tests - NO_X Emissions

	Panel A: Within-CMA Placebo			Panel B: Within-Industry Placebo		
	(1)	(2)	(3)	(4)	(5)	(6)
Placebo O_3 Reg.	-0.010	-0.006	-0.041	-0.023	-0.002	0.024
	(0.029)	(0.031)	(0.032)	(0.039)	(0.049)	(0.046)
R^2	0.978	0.978	0.978	0.978	0.978	0.978
N	2779	2779	2779	2779	2779	2779
Fraction of plants in targeted industry	0.50	0.80	0.20			
Fraction of plants in violating region				0.50	0.80	0.20

Notes: Table reports estimates of placebo tests of O_3 regulation's effect on plant NO_X emissions. Panel A reports estimates from placebo tests that randomize the plants that are assigned to targeted industries, but preserves the actual variation in the CMA-years that violate the O_3 regulation. Each column within Panel A assigns a different fraction of plants into the targeted industries. The Actual fraction of plants in targeted industries is 0.54. Panel B reports estimates from placebo tests that randomize the plants that are assigned to violating regions, but preserves the actual variation in the industries that are targeted. Each column within Panel B assigns a different fraction of plants into the violating regions. The Actual fraction of plants in violating regions is 0.42. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

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.

This specification is given by:

$$\begin{aligned} Y_{pict} = &\sum_b \beta^b_{PM}[K_i \times I(\underline{A^{PM}_b} \leq a^{PM}_{ct} < \overline{A^{PM}_b})] \\ &+ \sum_b \beta^b_{O3}[K_i \times I(\underline{A^{O3}_b} \leq a^{O3}_{ct} < \overline{A^{O3}_b})] + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \end{aligned} \tag{2.12}$$

where b indexes air quality bin numbers, K_i selects all industries targeted by the CWS, a^j_{ct} is the air quality measured in CMA c for pollutant j in year t, $\underline{A^j_b}$ is the air quality lower bound for bin b for pollutant j, $\overline{A^j_b}$ is the air quality upper bound for bin b for pollutant j, and $I(\underline{A^j_b} \leq a^j_c < \overline{A^j_b})$ is an indicator for all CMA-years with air quality that corresponds to bin b for pollutant j.[57] The coefficient β^b_j gives the effects of standard j in air quality bin b.

In estimating Equation (2.12), 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 $\mu g/m^3$. For the O_3 standard, we break the air quality distribution into six equal-sized bins from 57 to 77 ppb.[58]

The results of the estimating of Equation (2.12) using the full sample of polluters from the NPRI are displayed in Figure 2.8 and Figure 2.9 for $PM_{2.5}$ and NO_X emitters, respectively. Only the coefficients for the $PM_{2.5}$ standard are shown for $PM_{2.5}$ emissions, and the O_3 stan-

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

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

dard 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.

Figure 2.8: The Effect of $PM_{2.5}$ Regulation on $PM_{2.5}$ Emissions, by CMA Air Quality

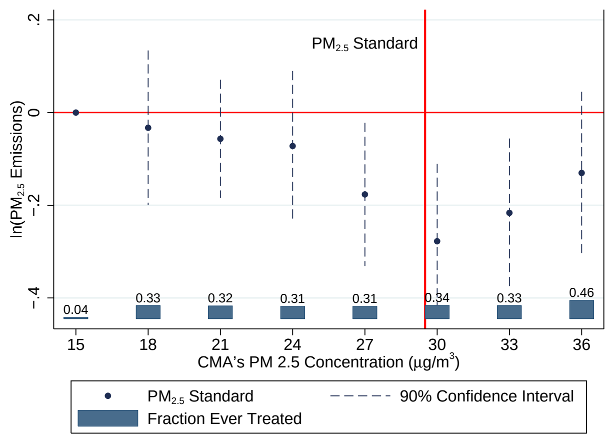

Notes: Figure displays estimates from a flexible DDD estimation of the $PM_{2.5}$ standard's effect on $PM_{2.5}$ emissions allowing the effects of regulation to vary by CMA air quality. Diamonds reflect the point estimates for each CMA air quality bin, while the dashed line displays the associated 90% confidence interval. These coefficients are measured relative to the excluded group (air quality below 18 $\mu g/m^3$ for $PM_{2.5}$ and below 57 ppb for O_3). Standard errors are clustered by industry-CMA. The histogram shows the fraction of observations in each bin treated by the respective standard at some point over the sample.

Figure 2.9: The Effect of O_3 Regulation on NO_X Emissions, by CMA Air Quality

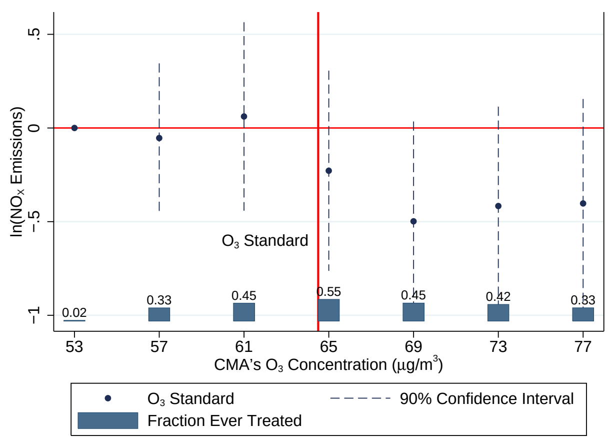

Notes: Figure displays estimates from a flexible DDD estimation of the O_3 standard's effects on NO_X emissions allowing the effects of regulation to vary by CMA air quality. Diamonds reflect the point estimates for each CMA air quality bin, while the dashed line displays the associated 90% confidence interval. These coefficients are measured relative to the excluded group (air quality below 18 $\mu g/m^3$ for $PM_{2.5}$ and below 57 ppb for O_3). Standard errors are clustered by industry-CMA. The histogram shows the fraction of observations in each bin treated by the respective standard at some point over the sample.

The results presented in Figure 2.8 Figure 2.9 show that a break that occurs just below the $PM_{2.5}$ standard's threshold for $PM_{2.5}$ 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 this chapter 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.

Finally, we adopt a common approach in the program evaluation literature 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 future.

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

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

$$Y_{pict} = \sum_{k=-3} \beta^k_{PM} T^{PM}_{ick} + \beta_{O3} T^{O3}_{ict} + \rho_p + \xi_{ct} + \lambda_{it} + \varepsilon_{pict}, \tag{2.13}$$

where T^{PM}_{ick} is an indicator for the years before ($k < 0$) or after ($k \geq 0$) a plant is treated for standard j, and T^{O3}_{ict} captures the average effect of the O_3 standard. We exclude the year prior to treatment for the $PM_{2.5}$ standard ($k = -1$), so the coefficients of interest (β^k_{PM}) report the semi-elasticity of treatment k years before or after treatment relative to the year before treatment. In other words, $\beta^k_{PM}$is the triple-difference coefficients relative to the year before a plant is first treated by the standard.[59]

[59]Note that in our main specification the triple-difference coefficient compares the average over all years

Figure 2.10: The Effect of $PM_{2.5}$ Regulation on $PM_{2.5}$ by Years Pre/Post Regulation

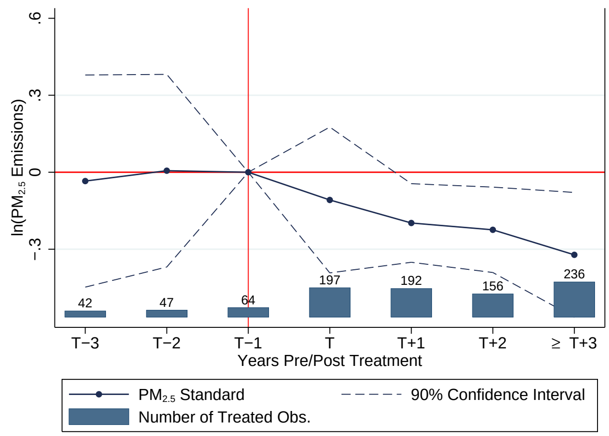

Notes: Figure shows the results of a flexible DDD estimation of the $PM_{2.5}$ standard for $PM_{2.5}$ emissions allowing the treatment effect to vary by years pre/post regulation. Diamonds show the triple-difference estimation coefficients by years before and after treatment, with a 90% confidence interval 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.

We estimate Equation (2.13) from three periods before a plant is treated onward. Separate coefficients are estimated up to three periods post treatment, and all periods greater than three years after treatment are pooled. We drop all observations that occur prior to three periods before a plant is treated. All plants in CMAs that began the sample period under treatment are dropped from the regression.

The results of the effects of the $PM_{2.5}$ standard on $PM_{2.5}$ emitters are shown in Figure 2.10. The dependent variable is the natural logarithm of $PM_{2.5}$ emissions and standard errors are clustered by CMA-industry.

Figure 2.10 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.

We can also show these 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 differences in firm ownership structure. For the sake of space, we relegate these results to the appendix (see Appendix A.2.2).

during which a plant is treated to the average over all years during which a plant is not treated.

2.5 Conclusion

In this chapter, we estimate the channels through which a change in environmental regulation 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). We then perform a decomposition exercise, first used in this literature by Levinson (2009), to show this clean-up was primarily driven by reductions in industry emission intensity. This result suggests the sources of the Canadian clean-up were similar to the sources of the clean-ups observed in the U.S. and Europe.

Next, we examine how Canadian manufacturing plants responded to a major revision to environmental policy, the Canada-Wide Standards for Particulate Matter and Ozone, and use the resulting empirical estimates to quantify the channels through which environmental regulations have contributed to the manufacturing clean-up. These estimates represent the first complete characterization of the regulatory channels driving the manufacturing clean-up. While these estimates are specific to Canada, given the similarity between the clean-ups and regulatory structures in Canada, the US, and Europe, we believe our results provide insights relevant for all three regions.

Our estimates imply that this policy explains approximately 60% of the reduction in nitrogen oxide emission intensity of the Canadian manufacturing sector from 2004 to 2010, and approximately 20% of the drop in particulate matter emission intensity. However, how this policy caused manufacturing to clean up varied considerably across pollutants. Over two-thirds of the nitrogen oxide clean-up caused by this policy was due to the adoption of cleaner production processes by surviving plants (the process effect). In contrast, over 80% of the particulate matter clean-up caused by this policy was due to plant exit (the selection effect) and the reallocation of output from regulated to unregulated plants (the reallocation effect).

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

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

Chapter 3

Environmental Regulation and the Pollution Haven Effect

3.1 Introduction

Debates over environmental regulations often centre on how these policies will impede international competitiveness by restricting the ability of domestic manufacturers to compete with foreign producers. The intuition underlying these concerns is simple: environmental policies raise production costs for domestic firms, making it more difficult to compete with foreign manufacturers who do not face similar policies, both at home and abroad. As a result, regulation may lead to what is known as a "pollution haven effect" (PHE), whereby domestic regulation reduces exports and increases imports (Copeland and Taylor, 2004). These outcomes are often seen as a problem by policy makers for two main reasons. First, there is a concern that regulations may overly disadvantage domestic producers, decreasing employment. Second, the change in relative costs may lead to increased foreign production. If the pollutant targeted by regulation is transboundary, this could increase foreign emissions and serve to undercut the effectiveness of domestic policy.[1]

While there is a growing empirical literature on the PHE, this literature has yet to directly test the PHE at the firm-level. Given recent innovations in the trade literature highlighting the heterogeneous nature of trade (Bernard et al., 2003; Bustos, 2011; Melitz, 2003), however, there is reason to believe the PHE may also be driven by heterogeneous effects across firms or facilities. For example, it is typically only a select group of firms within a given industry that

[1]This issue is particularly salient in the context of greenhouse gases, where it is termed "carbon leakage." For an overview of carbon leakage, see McAusland and Najjar (2015) or Fowlie et al. (2016).

are active in export markets. If an aspect of the PHE is the exit of firms from the export market, then this will clearly only be present for the firms that export absent regulation. Moreover, regulations themselves are often designed to differentially treat firms based on their characteristics.[2] Despite the real potential for heterogeneity in the PHE, the empirical work in this area has yet to push forward on this dimension. In this chapter, we assess the potential heterogeneity in the PHE by asking how environmental regulation affects the export behaviours of manufacturing facilities, both on average and across plants of different productivity levels.

Determining whether the PHE includes heterogeneous effects across producers may be important for understanding the discourse between industry and policy-makers. This is particularly evident in cases where policy makers believe that a particular policy will have a small effect on industry. While this may be true for the industry on average, if the effects of regulation on production costs vary across plants, this may not be true for the marginal exporter. As a consequence, even if a particular regulation were to produce a small PHE, measured in aggregate or average terms, it could still face substantial industry opposition.

In this chapter, we examine the effects of a major revision to Canadian environmental policy, the Canada Wide Standards for Particulate Matter and Ozone (CWS), on export volumes and the selection into (and out of) exporting of affected Canadian manufacturing plants. We estimate the effects of the CWS on plant export decisions using a unique longitudinal dataset that contains information on both the pollution emissions and production decisions of Canadian manufacturing plants over the period 2004-2010. This dataset allows us to clearly establish which plants in each industry were subject to regulation, making it possible for us to identify the effects of environmental policy on exports from individual plants.

Our empirical approach borrows heavily from our previous work examining how plant-level responses to the particulate matter and ozone regulations enacted under the CWS contributed to the clean-up of the Canadian manufacturing sector (see Chapter 2). The CWS implemented a pair of regional air quality standards across Canada, and targeted plants in select industries for regulation. The regulations used imposed a two-part regulatory constraint: requiring plants to either use clean production processes or face production constraints. To identify the effect of the CWS, we adopt a triple-difference research design, which exploits the variation in these regulations across industries, regions, and time. This approach allows us to flexibly control for factors such as regional demand shocks or industry-wide policy changes

[2]Regulations may explicitly do this, as is the case with intensity standards that do not affect the cleanest producers in an industry (e.g. former sulphur dioxide regulations in the US electricity sector (Lemoine, 2017)), and regulations that grandfather old firms based on past performance (e.g. the European Union's Emission Trading System (Knight, 2013)). Regulations may also do this implicitly, as can be the case of policies that penalize firms that fail to adopt clean technologies, such as the Clean Air Act (Greenstone, 2002).

that would otherwise confound the effects of environmental regulation.

Before discussing our empirical analysis, we present an overview of a simple theoretical framework of the heterogeneous nature of the PHE. As the nature of our contribution is primarily empirical, we do not explicitly solve the model. Instead, we use the framework to derive a set of empirical predictions that would hold under plausible circumstances. We use these predictions to provide intuition behind the PHE, and in particular highlight the potential heterogeneity of the PHE. Our model is similar in nature to that of Cherniwchan et al. (2017), in that we allow firm heterogeneity, assume firms compete via monopolistic competition, and allow firms to endogenously upgrade their technology. In contrast, however, we explicitly impose a small open economy assumption for the domestic market, and assume firms are regulated via a two-part regulatory rule, rather than a uniform pollution tax.

Our empirical predictions suggest regulation should cause some firms to exit the export market, and others to reduce their export volumes.[3] In addition, we show that under a two-part regulatory rule, the PHE should be most strongly felt by firms toward the middle of their industry's productivity distribution. The intuition for this particular form of heterogeneity is that, due to the fixed cost of exporting, the least productive firms in an industry will not participate in foreign markets. In contrast, while the most productive firms in an industry will export, they adopt technological improvements in response to regulation that allow them to avoid increases in their variable production costs. As a consequence, the firms that exit exporting and the surviving exporters that reduce export volumes should be moderately productive.

Next, we take our empirical approach to the data. We use a standard triple-difference research design to test for the average effects of the CWS. To test for the heterogeneity implied by our model, we follow the approach used by Bustos (2011) to study the differential effects of trade liberalization across plants and examine whether the effects of the CWS differ across plants on the basis of their initial productivity level.[4]

We find the CWS had a large negative effect on export volumes, causing a 22% reduction in export volumes for the average affected exporter. The policy, however, had no significant effect on the selection into or out of exporting for the average plant. In addition, as our model predicts, the effects of regulation varied considerably across plants of different productivity levels. In particular, arranging plants from least- to most-productive, the effects of regulation are most pronounced for plants in the second quintile of their industry's productivity distribution. For these plants, the CWS caused a 42% reduction in exports and reduced the probability

[3]These predictions require restricting the degree to which domestic wages and aggregate prices change in response to policy.

[4]As we do not observe capital information in our dataset, we define a firm's productivity as value added per worker in the first year they enter the dataset, after removing industry means.

of exporting by 10%.

In addition, we show that our results are robust to a number of potential identification problems. Specifically, we show that our baseline results are robust to controlling for the effects of foreign ownership, omitting the set of plants that do not sell domestically, and persistence in outcomes due to the presence of sunk costs.

This chapter relates to the large empirical literature on the PHE, for which there have been, essentially, three waves of literature. The earliest wave, prior to the late 1990s, tried to estimate the "competitiveness effects" of environmental regulations by studying changes in aggregate trade flows in response to policy. This early work typically found a paucity of evidence for the PHE (see, for example, Jaffe et al. (1995)). The lack of evidence for the PHE in these early papers was typically attributed to difficulties inherent in credibly identifying the effects of environmental policy on international trade (Cherniwchan et al., 2017). Levinson and Taylor (2008), for example, argue this difficulty is due in part to the potential endogeneity of environmental policy, as well as potential unobserved heterogeneity and aggregation issues.

The second wave of the empirical PHE literature attempted to address these methodological issues by either using an instrumental variable framework, or adopting a more nuanced view of the PHE. Work in this second wave has found evidence consistent with the PHE. As an example of the former, Levinson and Taylor (2008), examine the effects of environmental regulation on bilateral trade between the US and their main trading partners. They adopt a similar approach to much of the early literature, and rely on industry abatement spending as an indirect measure of the stringency of regulation facing an industry. Unlike the early literature, however, they develop an instrumental variable approach to address the potential endogeneity of this indirect measure. They find an increase in instrumented-abatement spending leads to an increase in net imports to the US from both Mexico and Canada. Other papers, such as Kellenberg (2009), Millimet and Roy (2016), and Broner et al. (2012), also use instrumental variable frameworks to address the endogeneity of environmental policy, finding evidence for the PHE. An example of the latter approach, Ederington et al. (2005) argue that work in the first wave failed to account for the geographic mobility of certain industries, as well as important characteristics of an industry's trade partners. Doing so, they find evidence for the PHE: industry abatement spending has a positive effect on imports from developing countries.

One potential issue with the second wave of the empirical PHE literature, as has been discussed by Cherniwchan et al. (2017), is that these approaches typically rely on model-driven arguments for the validity of their instruments. To address this limitation, the third wave of the empirical PHE literature has shifted from using an instrumental variable approach to quasi-experimental empirical strategies, such as difference-in-difference estimation, that exploit policy changes to identify the effect of environmental regulation. The advantage of the quasi-experimental approach is that the identification arguments typically require less structure, in terms of the underlying economic model, relative to the instrumental variable approach.

There are several examples of such work. Shi and Xu (2018), for example, exploit variation in regulatory stringency across industries and provinces in China stemming from China's eleventh five-year plan. They find that relatively pollution-intensive industries in highly regulated provinces experience a reduction in exports relative to counterfactual industry-provinces. In addition, Aichele and Felbermayr (2015) exploit variation in the stringency of environmental policy across countries stemming from the Kyoto Protocol, an international agreement on climate change policy. They find the Kyoto Protocol caused an increase in imports to relatively highly regulated countries. [5]

The work in these three waves of literature on the PHE all share at least one common feature: they all assess the PHE using industry or national-level data. In addition to missing the potential heterogeneity in the PHE, relying on industry- or national-level data is potentially problematic, as the presence of plant- or firm-level heterogeneity may have implications for the ability to identify the effect of regulation on international competitiveness. Identification is a potential issue, as it is harder to argue for the exogeneity of policy with respect to aggregate values than with respect to firms. For example, in the case of air quality standards, arguments that a single firm would be unable to influence the stringency of regulation rely on an assumption of that firm representing a relatively small amount of production in its region. Clearly, this same argument cannot be made for aggregate production in the region, particularly if regulators care about regional employment or firms are represented by regional trade associations.[6]

In addition to contributing to the literature on the PHE, this chapter also contributes to a large literature examining the effects of environmental policy on manufacturing facilities. Most of this work has focused on domestic outcomes, such as output (e.g. Greenstone (2002)), productivity (e.g. Berman and Bui (2001b), Greenstone et al. (2012)), employment (e.g. Berman and Bui (2001a), Walker (2013)), pollution (e.g. Greenstone (2003), Fowlie et al.

[5] In work related to the PHE, Hanna (2010) exploits variation in regulation across industries and regions stemming from the US Clean Air Act to ask whether regulation affects the foreign output of multinational firms. Hanna finds that multinationals with locations in the US responded to increased domestic regulation by increasing output in foreign markets. While not directly on the international flow of goods, this is suggestive evidence for the PHE.

[6] While Hanna's work avoids this pitfall, her focus on multinational firms means her results are potentially not representative of regulation's effect on small- and medium-sized firms, or large firms that only operate in a single market.

(2012)), and plant entry and exit (e.g. Becker and Henderson (2000), List et al. (2003)). Our research complements this body of work by showing how regulation affects a facility's participation in foreign markets.

Finally, our results contribute to a burgeoning empirical literature examining the microfoundations of the relationship between international trade and the environment. To date the majority of these studies have focused on how international trade affects environmental outcomes at individual manufacturing plants (e.g. Martin (2012), Cherniwchan (2017)). Our study contributes to this line of research by providing evidence of how environmental regulations affect manufacturing plants' participation in international trade.

The remainder of this chapter proceeds as follows. Section 2 presents our model. Section 3 discusses our data. Section 4 discusses our research design and presents our baseline empirical specification. Section 4 presents our results. Section 5 concludes.

3.2 The Pollution Haven Effect in a Small Open Economy

In this section, we present an overview of a simple model of how firms that differ on the basis of their productivity respond to environmental regulation in a small open economy. As the goal of this paper is primarily empirical in nature, we do not fully solve the model, nor do we dive into the myriad of important theoretical questions about the PHE in an economy with heterogeneous firms. Instead, we present a framework we believe adequately captures our empirical setting, and use this framework to present a series of empirical predictions. We use these empirical predictions to show outcomes that would hold under what we believe are realistic, although not necessary, industry-level responses (such as average domestic prices rising following regulation).

We begin by presenting the model's set-up and discussing how it would be solved. We then move onto our empirical predictions, and discuss what these predictions mean for our empirical setting.

3.2.1 Model Set-Up

We adopt a framework similar to that of Melitz (2003), in which heterogeneous firms compete via monopolistic competition in both domestic and foreign markets. Unlike Melitz (2003), however, we explicitly impose a small open economy assumption for the domestic market, following an approach used by Demidova and Rodríguez-Clare (2009). In addition, we assume firms emit pollution as a by-product of production and are regulated via a two-part regulatory constraint intended to capture a common feature of air pollutant regulation.[7] Our two-part regulatory constraint, which we describe further below, requires firms to adopt a clean production process, and levies a pollution tax against those that fail to do so.

Consider an economy with a single industry, which itself is comprised of a continuum of firms that compete via monopolistic competition. Firms produce a single good with a single input, labor (denoted l), the use of which creates pollution. As mentioned above, pollution emitted by domestic producers is regulated via a two-part regulatory constraint. Under this constraint, domestic firms that use a clean-production process, which we refer to as retrofitted technology (and label r), are unregulated. In contrast, domestic firms that use a dirty-production process must pay a penalty, τ, levied on each unit of pollution emitted. The penalty is meant to reflect both a "compliance cost" of dealing with increased regulatory oversight, as well as a direct penalty levied on dirty producers.[8] We adopt this type of regulation because it is a common form of regulation used to address air pollution.[9] In addition, while not a direct representation of our empirical setting, it is a good approximation of the policy we study.[10]

To enter the market, firms must employ f_ε worth of labor. This makes the fixed entry cost, wf_ε, where w is the wage rate. Entering the market allows a firm to draw a productivity level, φ, from some known distribution, denoted by $H(\varphi)$. The productivity draw dictates the firm's input requirement, such that a firm with a higher φ has a lower unit-labor requirement, given by $\frac{l}{\varphi}$.

As φ affects the unit-labor requirement, it also affects a firm's pollution. Firms that use the dirty-production technology have a pollution-labor ratio given by $\frac{z(\varphi)}{l(\varphi)} = \kappa$. Firms that use the retrofitted production technology have a pollution-labor ratio given by $\frac{z_r(\varphi)}{l_r(\varphi)} = \frac{\kappa}{\gamma}$. As a result, the pollution emitted by a firm that uses the dirty-technology and has productivity φ is $z(\varphi) = \frac{\kappa x(\varphi)}{\varphi}$, while pollution from a retrofitted producer is $z_r(\varphi) = \frac{\kappa x(\varphi)}{\gamma\varphi}$. Moreover, to adopt the retrofitted technology, a firm must pay a fixed cost of f_r units of labor. Notice that the assumptions on the retrofitted technology make it similar to the productivity-enhancing technology upgrading in Bustos (2011). Here, retrofitting affects a firm's pollution intensity, whereas in Bustos (2011), technology upgrading reduces labor intensity.

[7]For a Melitz-style model of heterogeneous firms that face a uniform pollution tax, see Cherniwchan et al. (2017).

[8]Compliance costs may arise as a result of regular meetings between managers and regulators, or information reporting, for example.

[9]For example, the National Ambient Air Quality Standards used as part of the US Clean Air Act feature this type of policy (Greenstone, 2002).

[10]The policy we study, the CWS, imposed production constraints on plants that failed to adopt clean processes, as well as required more stringent oversight of these plants.

In what follows, we use the superscripts *no* and *reg* to represent outcomes in a regime without regulation and with regulation, respectively. With this set-up, absent regulation, the firm's unit cost is given by $c^{no}(\varphi) = \frac{w}{\varphi}$, under either the dirty or retrofitted technology. When the domestic market is regulated, the firm's unit cost is $c^{reg}(\varphi) = \frac{w+\tau\kappa}{\varphi}$ and $c_r^{reg}(\varphi) = \frac{w}{\varphi}$, for dirty and retrofitted producers, respectively.

Domestic firms can sell in the domestic market or export goods to a foreign market. Serving the domestic market requires paying a fixed market access cost of wf, while serving the foreign market requires paying an additional fixed cost wf_x. Similarly, foreign producers can import goods into the domestic market as well as sell in their own market. Import fixed-costs are given by f_m, which is independent of the domestic wage.

Preferences of domestic consumers are given by

$$U = \left[\int_{v \in \Omega} q(v)^{\rho} dv + \int_{v' \in \Omega_m} q(v')^{\rho} dv' \right]^{1/\rho}, \tag{3.1}$$

where v and v' denote domestic and foreign varieties, and Ω and Ω_m denote the set of all available domestic and imported varieties, respectively. The elasticity of substitution is given by $\sigma = \frac{1}{1-\rho}$, where $0 < \rho < 1$.

The domestic market is comprised of L identical consumers, each endowed with a unit of labor that is inelastically supplied to the domestic market. By the small open economy assumption, changes in the domestic labor market do not affect wages in the foreign market. Hence, domestic wages are given by w and foreign wages are given by w_m. Domestic consumers exhaust all income, which means demand for domestic good, v, and imported good, v', are given by

$$q(v) = IP^{\sigma-1} p(v)^{-\sigma} \text{ and } q_m(v') = IP^{\sigma-1} p_m(v')^{-\sigma}, \tag{3.2}$$

where I denotes consumer income and the domestic price index is given by P, such that

$$P^{1-\sigma} = \int_{v \in \Omega} p(v)^{1-\sigma} dv + \int_{v' \in \Omega_m} p_m(v')^{1-\sigma} dv'.$$

Given this demand system, a domestic producer charges a mark-up over their unit (or marginal) costs. Absent regulation, a firm with productivity φ charges $p^{no}(\varphi) = \frac{w}{\rho\varphi}$ under either technology. As a result, domestic revenues absent regulation under either technology are given by

$$r^{no}(\varphi) = \frac{I\,[P\rho]^{\sigma-1}\,\varphi^{\sigma-1}}{w^{\sigma-1}}. \tag{3.3}$$

Domestic profits using the dirty technology are $\pi^{no}(\varphi) = \frac{r(\varphi)}{\sigma} - wf$, and domestic profits using retrofitted technology are $\pi_r^{no}(\varphi) = \frac{r(\varphi)}{\sigma} - w\,[f + f_r]$.

To make the small open economy assumption explicit, we adopt the approach of Demidova and Rodríguez-Clare (2009) and assume foreign demand for a domestic variety, v, is given by $Ap_x(v)^{-\sigma}$, where A is exogenous and $p_x(v)$ is the price charged by a domestic exporter.[11] Notice this means that even though the domestic market is small, domestic exporters have some price setting power. However, unlike in the domestic market, changes in their price only affect the demand of their goods, and do not affect the demand for other varieties in the foreign market. Using either technology, without regulation, an exporter with productivity φ charges $p_x(\varphi) = \frac{w}{\rho\varphi}$, and obtains revenues given by

$$r_x^{no}(\varphi) = \frac{A\,[\rho]^{\sigma-1}\,\varphi^{\sigma-1}}{w^{\sigma-1}}. \tag{3.4}$$

Export profits under either technology, absent regulation, are $\pi_x^{no}(\varphi) = \frac{r_x(\varphi)}{\sigma} - wf_x$.

Without regulation, adopting the retrofitted technology requires paying a fixed cost, but does not change the firm's revenues. As a result, no firms retrofit absent regulation.

Under regulation, domestic producers also set prices as a mark-up over their unit costs. Producers that use the dirty technology charge $p^{reg}(\varphi) = \frac{w+\tau\kappa}{\rho\varphi}$ for their output, whereas producers that use the retrofitted technology charge $p_r^{reg}(\varphi) = \frac{w}{\rho\varphi}$. As a result, domestic revenues for the dirty technology are given by

$$r^{reg}(\varphi) = \frac{I\,[P\rho]^{\sigma-1}\,\varphi^{\sigma-1}}{[w+\tau\kappa]^{\sigma-1}}, \tag{3.5}$$

and domestic revenues for the retrofitted technology are

$$r_r^{reg}(\varphi) = \frac{I\,[P\rho]^{\sigma-1}\,\varphi^{\sigma-1}}{w^{\sigma-1}}. \tag{3.6}$$

In addition, profits from production are $\pi^{reg}(\varphi) = \frac{r(\varphi)}{\sigma} - wf$ and $\pi^{reg}(\varphi) = \frac{r(\varphi)}{\sigma} - w\,[f + f_r]$

[11] As Demidova and Rodríguez-Clare (2013) show, this form of foreign demand is the limiting case of a two large economy model, where one economy becomes infinitesimally small.

for the dirty and retrofitted technology, respectively.

Similarly, it can be shown that when the domestic market is regulated, export revenue for a firm using the dirty technology is

$$r_x^{reg}(\varphi) = \frac{A\left[\rho\right]^{\sigma-1}\varphi^{\sigma-1}}{\left[w+\tau\kappa\right]^{\sigma-1}}. \tag{3.7}$$

In addition, export profits under the dirty technology are $\pi_x^{reg}(\varphi) = \frac{r_x^{reg}(\varphi)}{\sigma} - wf_x$. In contrast, export revenue for a firm using the retrofitted technology, when regulated, is

$$r_{x,r}^{reg}(\varphi) = \frac{A\left[\rho\right]^{\sigma-1}\varphi^{\sigma-1}}{w^{\sigma-1}}, \tag{3.8}$$

and profits are $\pi_{x,r}^{reg}(\varphi) = \frac{r_{x,r}^{reg}(\varphi)}{\sigma} - wf_x$

Notice that under regulation, a firm chooses to adopt the retrofitted technology if doing so increases revenues enough to cover the fixed cost of retrofitting. Assuming the fixed-cost of retrofitting is large enough relative to the fixed-cost of exporting[12], then a firm is indifferent between using the dirty and retrofitted technology if

$$\pi^{reg}(\varphi) + \pi_x^{reg}(\varphi) = \pi^{reg}(\varphi) + \pi_{x,r}^{reg}(\varphi).$$

Denoting the productivity level that would make a firm indifferent between producing using the dirty or retrofitted technology as φ_r, then this productivity cut-off is given by

$$\varphi_r^{\sigma-1} = \frac{\sigma w^{\sigma} f_r}{\rho^{\sigma-1}\left[wL\left[P^{reg}\right]^{\sigma-1} + A\right]\left[1 - G(w,\tau)\right]}, \tag{3.9}$$

where $G(w,\tau) = \left[\frac{w}{w+\tau\kappa}\right]^{\sigma-1}$ and we have assumed the tax revenue is destroyed, which means domestic income is $I = wL$. This means that, under regulation, any firm that draws a productivity level greater than φ_r would choose to adopt the retrofitted technology. Any firm that draws a productivity level below φ_r would use the dirty technology.

Three market clearing conditions are required to close the model: a labor market clearing condition, a free entry condition, and a trade balance condition. Labor market clearing requires total domestic labor supply to equal total domestic labor demand. Labor market clearing under

[12]The parameter restriction required for this to be a valid equilibrium is $\frac{f_r}{f_x} > \left[\frac{wL[P^{reg}]^{\sigma-1}+A}{A}\right]\left[\left[\frac{w^{reg}+\tau\kappa}{w^{reg}}\right]^{\sigma-1} - 1\right] > 1.$

regulatory regime j can be written as

$$L = M^j \left[f_\varepsilon + \bar{L}_d^j + \bar{L}_x^j + \left[1 - H(\varphi_\varepsilon^j)\right] f + \left[1 - H(\varphi_x^j)\right] f_x + \left[1 - H(\varphi_r^j)\right] f_r \right],$$

where M^j is the measure of entrants, φ_ε^j is the productivity cut-off below which firms choose to exit the market, $\bar{L}_d^j$ is average labor demand for production to serve the domestic market, and $\bar{L}_x^j$ is average labor demand for production to serve the foreign market. Free entry requires firms to earn zero expected profits from entering production. Letting δ be an exogenous exit probability, and $\bar{\pi}^j$ be average profits for domestic producers, then free entry gives

$$\bar{\pi}^j = \delta w f_\varepsilon.$$

Finally, the trade balance condition requires that the total value of exports equal the total value of imports. Trade balance requires

$$M^j \bar{R}_x^j = M_{FOR} \bar{R}_m^j,$$

where M_{FOR} is the measure of foreign firms that enter the domestic market, $\bar{R}_x^j$ are average export revenues for domestic producers, and $\bar{R}_m^j$ are average import revenues for foreign producers that sell in the domestic market.

3.2.2 Empirical Predictions of the Pollution Haven Effect

With this set-up, we now present a series of empirical predictions relevant for the Pollution Haven Effect. We present each prediction, and use our theoretical framework to explain why the prediction would hold under the specified conditions.

Recall that the PHE arises because domestic regulation increases the cost of producing in the domestic market relative to the cost of producing in the unregulated foreign market. Our first two predictions make this argument clear by examining the effect imposing environmental regulation has on a firm's total and export revenues, respectively.

The first of these predictions is intended to show the conditions under which a reduction in total revenues can be used to infer that variable production costs have increased for domestic producers. As we discuss in the empirical section, we cannot observe variable costs. This is problematic because changes in variable costs are the cornerstone of the PHE. However, as is made clear by Empirical Prediction 1, by observing changes in a plant's total revenue in response to regulation, we can conclude variable costs likely rose as a result of said policy.

To introduce this prediction, notice that we can express total revenues for a firm that uses production technology t in regulatory regime j as

$$tr_t^j(\varphi) = \begin{cases} \frac{\left[I^j\left[P^j\right]^{\sigma-1}\right]\rho^{\sigma-1}}{c_t^j(\varphi)} & \text{if } \varphi_\varepsilon^j \geq \varphi < \varphi_x^j, \\ \frac{\left[I^j\left[P^j\right]^{\sigma-1}+A\right]\rho^{\sigma-1}}{c_t^j(\varphi)} & \text{if } \varphi \geq \varphi_x^j, \end{cases} \tag{3.10}$$

where $c_t^j(\varphi)$ is the variable cost of producing with technology t in regime j, φ_ε^j is the productivity cut-off below which firms exit the domestic market, and φ_x^j is the productivity cut-off below which firms exit the foreign market. Clearly, firms that export, those with $\varphi \geq \varphi_x^j$, receive revenues from both the domestic and foreign markets, while firms that do no export only receive domestic revenues.

For a firm with a given productivity level, dividing total revenues under regulation by their total revenues without regulation gives the proportional change in revenues due to regulation. To make this explicit by way of example, taking this ratio for a non-exporting firm, using $I^j = w^j L$, and simplifying, gives[13]

$$\frac{tr_t^{reg}(\varphi)}{tr_t^{reg}(\varphi)} = \frac{\left[w^{reg}\left[P^{reg}\right]^{\sigma-1}\right]}{\left[w^{no}\left[P^{no}\right]^{\sigma-1}\right]} \frac{c^{no}(\varphi)}{c^{reg}(\varphi)}. \tag{3.11}$$

Regulation affects revenues because of its direct effect on production costs, and its indirect effects on wages and average industry prices. From Equation (3.11), it is straightforward to show that revenues fall for a firm that uses technology t and does not export under either regime if the increase in production costs satisfies

$$\frac{c_t^{reg}(\varphi)}{c_t^{no}(\varphi)} > \frac{w^{reg}\left[P^{reg}\right]^{\sigma-1}}{w^{no}\left[P^{no}\right]^{\sigma-1}}. \tag{3.12}$$

Similarly, restrictions on production cost changes that deliver a reduction in total revenues can be shown for firms that export in both regimes, or that drop out of exporting as a result of regulation.[14] This is all to say that total revenues will fall for any firm that experiences a large enough increase in production costs as a result of regulation (either directly, or indirectly through changes in wages).

[13]As discussed above, we assume tax revenues are destroyed.

[14]If regulation causes industry prices to rise relative to equilibrium wages, then it can be shown that Equation (3.12) is also a sufficient condition to ensure total revenues fall for any surviving firm that uses technology t. This claim follows because if $\left[\frac{P^{reg}}{P^{no}}\right]^{\sigma-1} > \frac{w^{reg}}{w^{no}}$, then $\frac{I^{reg}\left[P^{reg}\right]^{\sigma-1}}{I^{no}\left[P^{no}\right]^{\sigma-1}} > \frac{I^{reg}\left[P^{reg}\right]^{\sigma-1}+A}{I^{no}\left[P^{no}\right]^{\sigma-1}+A} > \frac{I^{reg}\left[P^{reg}\right]^{\sigma-1}}{I^{no}\left[P^{no}\right]^{\sigma-1}+A}$.

The discussion of the effects of regulation, so far, has ignored any potential heterogeneity involved in the effects of regulation on revenues. There are, however, two margins through which heterogeneity arises. First, recall that due to the nature of this policy instrument, variable production costs rise more for firms that do not retrofit than for the firms that retrofit. As we showed in the preceding section, these retrofitting firms will be relatively productive. Thus, the most productive firms in an industry will retrofit, and will experience a relatively small reduction in revenues as a result.[15]

Second, as we will discuss in greater detail for the second empirical prediction, export revenues may fall by more than domestic revenues following regulation. This would occur if average domestic prices rise following regulation, thereby insulating domestic revenues from the cost increase. As the least productive firms only serve the domestic market, and domestic sales are insulated due to domestic price adjustments, then regulation will cause a relatively small proportional reduction in total revenues for the least productive firms. The consequence of these two forms of heterogeneity is that the reduction in total revenues is largest, in percentage terms, in the middle of the productivity distribution.

We summarize these results in the following empirical prediction.

Empirical Prediction 1. *If the increase in variable production costs for a given firm is sufficiently large, then*

(a) imposing environmental regulation will cause a reduction in the firm's total revenues;

(b) in percentage terms, total revenues will fall the most for firms toward the middle of the productivity distribution if average domestic prices rise in the industry.

By increasing the variable production costs for domestic producers, regulation should decrease a firm's optimal scale, provided the increase in production costs is sufficiently large. Our second empirical prediction explicitly connects this increase in production costs to the PHE.

Empirical Prediction 2. *If the domestic price index rises, then for a given domestic firm, imposing environmental regulation will cause a larger proportional reduction in export revenues than domestic revenues.*

To see how Empirical Prediction 2 arises, it suffices to show how domestic and export revenues change in response to regulation for a firm with a given productivity level. Recall that absent regulation, firm revenues from domestic sales under either technology are given

[15] Note that if wages fall due to regulation, then revenues may even rise for these firms.

by Equation (3.3). In contrast, when regulation is imposed, a firm's domestic revenues when using the retrofitted technology is given by Equation (3.6), and their domestic revenues when using the dirty technology is given by Equation (3.5). In addition, as discussed above, only the most productive firms in the industry adopt the retrofitted technology.

Comparing the domestic revenues across the two regulatory regimes for a firm that uses the retrofitted technology under regulation gives

$$\frac{r_r^{reg}(\varphi)}{r^{no}(\varphi)} = \left[\frac{P^{reg}}{P^{no}}\right]^{\sigma-1} \left[\frac{w^{no}}{w^{reg}}\right]^{\sigma}. \tag{3.13}$$

That is, domestic revenues only change for firms that use the retrofitted technology because of the change in equilibrium industry prices and wages. In comparison, the change in domestic revenues for a firm that does not retrofit under regulation gives

$$\frac{r^{reg}(\varphi)}{r^{no}(\varphi)} = \left[\frac{P^{reg}}{P^{no}}\right]^{\sigma-1} \left[\frac{w^{no}}{w^{reg}}\right]^{\sigma} \left[\frac{1}{1+\frac{\tau\kappa}{w^{reg}}}\right]^{\sigma-1}. \tag{3.14}$$

This means domestic revenues change for the non-retrofitting producers because of increased production costs, as well as the change in equilibrium industry prices and wages.

A similar comparison gives the change in export revenues for the producers that choose to retrofit their technology as

$$\frac{r_{x,r}^{reg}(\varphi)}{r_x^{no}(\varphi)} = \left[\frac{w^{no}}{w^{reg}}\right]^{\sigma-1}, \tag{3.15}$$

and the change in export revenues for non-retrofitters as

$$\frac{r_x^{reg}(\varphi)}{r_x^{no}(\varphi)} = \left[\frac{w^{no}}{w^{reg}}\right]^{\sigma-1} \left[\frac{1}{1+\frac{\tau\kappa}{w^{reg}}}\right]^{\sigma-1}. \tag{3.16}$$

Comparing Equation (3.16) to Equation (3.14) shows that if domestic prices rise as a result of regulation, then the reduction in export revenues for non-retrofitting producers – those whose variable costs are directly affected by regulation – is larger than the reduction in domestic revenues. The comparison of Equation (3.15) to Equation (3.13) shows the same holds for producers that adopt the retrofitted technology when regulated.

Empirical Prediction 2 clarifies the intuition underlying the PHE in a small open economy. In equilibrium, domestic prices and wages adjust to domestic policy, which may insulate domestic producers who sell in the domestic market. However, because the domestic market is assumed to be small relative to the foreign market, this policy change does not affect foreign prices. This leads to a relatively large reduction in export revenues compared to domestic revenues.

Our third empirical prediction highlights the heterogeneous nature of the PHE under this regulatory environment. Under a two-part regulatory constraint, firms can avoid paying higher production costs by adopting clean (or retrofitted, as we have called it) technology. By avoiding increased production costs, these firms experience smaller reductions in both domestic and export revenues relative to the firms that use the dirty technology. This leads to our third prediction.

Empirical Prediction 3. *In percentage terms, the reduction in export revenues as a result of environmental regulation will be largest for the firms that do not retrofit their production technology. These firms will be the least productive exporters in the industry.*

Empirical Prediction 3 follows immediately by comparing Equation (3.15) to Equation (3.16). Firms that do not retrofit while regulated, which are the least productive in the industry, experience an increase in production costs not felt by the firms that retrofit. As a result, the reduction in export revenues is largest for these low-productivity exporters. Empirical Prediction 3 says that the PHE, at least under a two-part regulatory constraint, will be most pronounced for the least productive exporters. It is also worth noting that, as exporting requires paying a fixed-cost, only relatively productive firms choose to export. This implies that, while the heavily affected exporters will not be the most productive firms in their industry, they will not be the least productive either, as the least productive firms will not export.

The fixed-cost of exporting produces one final empirical prediction worth discussing. Due to this export fixed-cost, not all domestic producing firms within a given industry serve the foreign market. Moreover, the marginal exporter will have a relatively low productivity level, and as a result will use the non-retrofitted production technology. This leads to the following empirical prediction.

Empirical Prediction 4. *If the reduction in wages following regulation is sufficiently small, then imposing environmental regulation causes some firms to leave the export market. These firms will be the least productive exporters in the industry.*

To see how Empirical Prediction 4 arises, recall that a firm's size is a monotone transformation of productivity. This means there is a productivity cut-off above which all firms export, and below which no firms export. Absent regulation, this export productivity cut-off, labeled

φ_x, is given by

$$\varphi_x^{no} = \left[\frac{\sigma\left[w^{no}\right]^{\sigma} f_x}{A\rho^{\sigma-1}}\right]^{\frac{1}{\sigma-1}}. \tag{3.17}$$

Under regulation, the export cut-off is

$$\varphi_x^{reg} = \left[\frac{\sigma\left[w^{reg}\right]^{\sigma} f_x}{A\rho^{\sigma-1}}\right]^{\frac{1}{\sigma-1}} \left[1 + \frac{\tau\kappa}{w^{reg}}\right]. \tag{3.18}$$

Comparing the two export productivity cut-offs gives

$$\left[\frac{\varphi_x^{reg}}{\varphi_x^{no}}\right]^{\sigma-1} = \left[\frac{w^{reg}}{w^{no}}\right]^{\sigma} \left[1 + \frac{\tau\kappa}{w^{reg}}\right]^{\sigma-1}, \tag{3.19}$$

which is greater than one if, and only if, $\left[\frac{w^{no}}{w^{reg}}\right]^{\sigma} < \left[1 + \frac{\tau\kappa}{w^{reg}}\right]^{\sigma-1}$. Moreover, as exit from exporting is caused by an increase in the exporting productivity cut-off, then by definition, the exiting firms will be the least productive exporters.

Notice that Empirical Prediction 4 also has implications for identifying the Pollution Haven Effect. In particular, it implies that an observed reduction in firm revenues need not be evidence of a Pollution Haven Effect. The intuition behind this is that, as only relatively large firms choose to export, observing a change in firm revenues may reflect changes at firms that are not present in foreign markets. This could, in principal, reflect general equilibrium effects of policy, rather than changes in the relative cost of production between domestic and foreign producers.[16]

To emphasize the intuition underlying the heterogeneous nature of the PHE we present a series of graphs showing how regulation affects a firm's export decision and technology choices. Figure 3.1 plots a firm's profit absent regulation as a non-exporter ($\pi^{no}(\varphi)$), as an exporter ($\pi_x^{no}(\varphi)$), and as an exporter that uses the retrofitted technology ($\pi_{x,r}^{no}(\varphi)$).[17] As the figure shows, firms with a productivity draw above φ_x^{no} choose to pay the fixed cost to export, and sell in both the domestic and foreign markets. The remaining producing firms, with $\varphi \in (\varphi_\varepsilon^{no}, \varphi_x^{no})$, only serve the domestic market. No firm chooses to retrofit, as doing so raises the fixed costs of production, and does not affect revenues.

Figure 3.2 shows how regulation affects these decisions.[18] In this figure, we plot firms'

[16]Note that it is straightforward to show that if $A < IP^{\sigma-1}$, then some firms only serve the domestic market.
[17]To simplify the graph, we do not plot profits for firms that do not export but use the retrofitted technology.
[18]For simplicity, we hold industry prices and wages fixed.

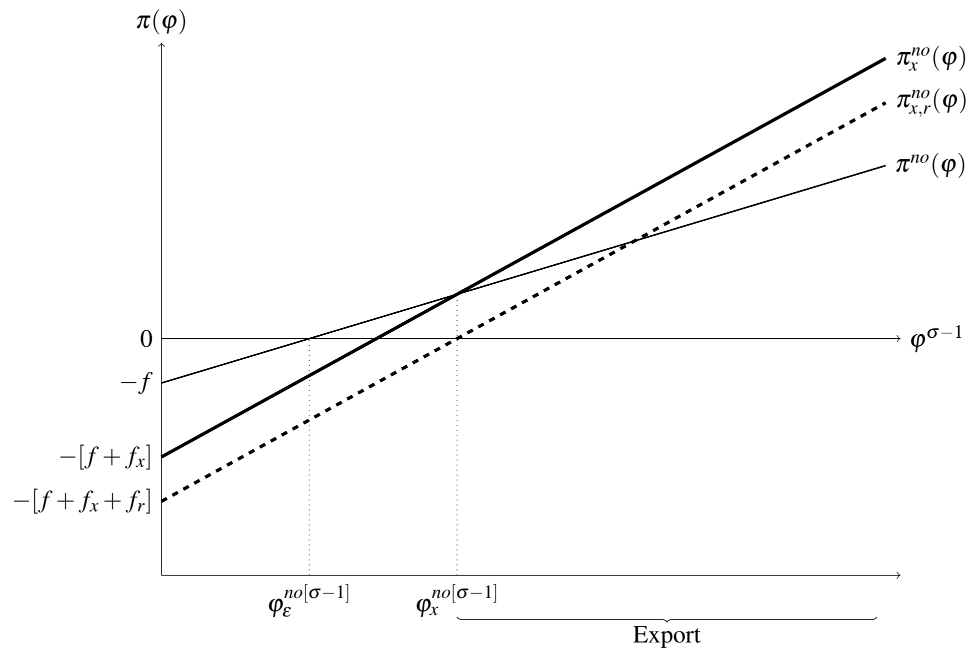

Figure 3.1: Export and Technology Choices without Environmental Regulation

profits both with and without regulation for non-exporters, exporters, and exporters that use the retrofitted technology. Profits under regulation are labeled with superscript *reg*, while the profits without regulation are labeled *no*. Regulation reduces the profitability of operating using the dirty technology for both exporters and non-exporters alike. This causes the least productive exporters, those with $\varphi < \varphi_x^{reg}$, to exit the export market, as we discussed in Empirical Prediction 4. In addition, the most productive exporters, those with $\varphi \geq \varphi_r^{reg}$, now find it worthwhile to retrofit. The remaining firms, however, remain in the export market, but do not retrofit. As we showed in Empirical Prediction 3, these low-productivity exporters face a large reduction in export sales because their production costs increase, and foreign prices cannot adjust due to the small open economy assumption.

In the empirical analysis that follows, we test these predictions by examining the effect of a change in Canadian environmental policy on the export behaviours of Canadian manufacturing plants. Our tests of Empirical Prediction 1 and Empirical Prediction 2 come from examining regulation's effect on a plant's total sales and comparing this to regulation's effect on sales that occur in foreign markets. We test Empirical Prediction 3 by examining how regulation's effect on export sales varies across plants of different productivity levels. Finally, we test Empirical Prediction 4 by estimating regulation's effect on a plant's decision to select into and out of exporting, in particular, across plants of different productivity levels.

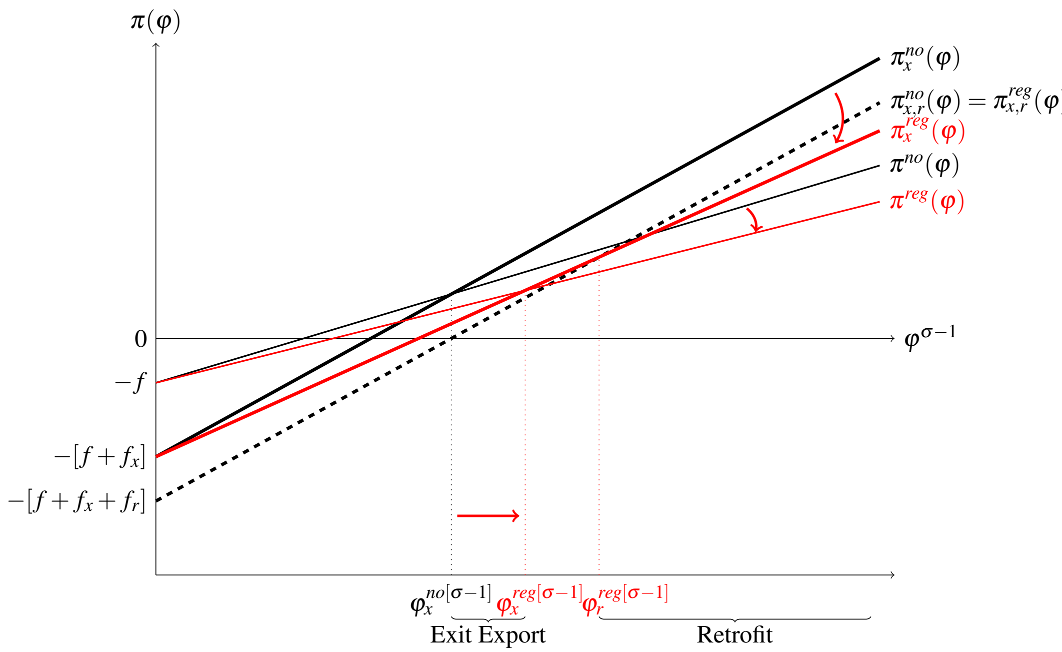

Figure 3.2: Export and Technology Choices with Environmental Regulation

3.3 Data and Measurement

Our goal in this chapter is to determine the effect of environmental regulation on the export decisions of Canadian manufacturing plants. To do so, we utilize a unique micro dataset that contains information on both the pollution emissions and export decisions of Canadian manufacturing plants over the period 2004-2010. This dataset was created by linking the data from the National Pollution Release Inventory (NPRI), a publicly available dataset containing information on the pollution emissions of Canadian manufacturing facilities, with the confidential data on plant characteristics from the Annual Survey of Manufacturers (ASM).[19] Together, these data sources allow us to create a longitudinal dataset containing information on the export decisions of plants that emit fine-scale particulate matter ($PM_{2.5}$), a pollutant regulated in Canada over our period of study as part of the suite of of environmental regulations called the Canada-Wide Standards.

3.3.1 Descriptive Statistics

In order to understand how environmental regulations affect a plant's participation in international markets, we examine three outcomes: the likelihood of exporting, total sales, and

[19]These data were linked by Statistics Canada. For further details on the data and its construction, see Chapter 2.

Table 3.1: Summary Statistics

	$PM_{2.5}$ (1)	Full ASM (2)
Sales ($1 mill.)	194.62	11.12
	(890.55)	(123.56)
Exports ($1 mill.)	97.88	6.661
	(709.67)	(89.74)
Pr(Export)	0.76	0.36
	(0.43)	(0.48)
N	6501	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 plants that emit $PM_{2.5}$ and column (2) reports plant characteristics for the entire manufacturing sector. Statistics in column (1) 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.

export sales. Together, these variables allow us to determine if environmental regulations affect international competitiveness along the intensive (total sales, and exports) or extensive (likelihood of exporting) margins.

Summary statistics for total sales, total exports and the probability of exporting are reported in Table 3.1. Column (1) of the table reports statistics for our main dataset, which comprises an unbalanced panel of manufacturing plants that emit $PM_{2.5}$ pollution. The summary statistics in column (1) are weighted to account for any possible sample bias created by the procedure used to link the NPRI and ASM.[20] For comparison, column (2) of the table reports summary statistics for the entire sample of plants in the ASM. Although we do not use the full ASM dataset in our analysis, we present these statistics to highlight the difference between our sample of polluters and the average Canadian manufacturing facility.

The descriptive statistics reported in Table 3.1 suggest that, on average, the manufacturing plants that emit $PM_{2.5}$ are substantially larger, are much more likely to export, and export more than the average plant in the Canadian manufacturing sector. This is driven both by reporting requirements for the NPRI, as plants typically only report to the NPRI database if they have at least ten employees, and structural differences between polluters and non-polluters.[21]

[20]For details, see Chapter 2.

[21]It is worth noting that the NPRI requires any plant that operates a boiler or generator on-site to report their $PM_{2.5}$ emissions, regardless of their number of employees. As many industrial $PM_{2.5}$ emitters use an on-site boiler or generator, it is unlikely that the employment threshold is the main cause of the differences reported in Table 3.1.

3.3.2 Canadian Environmental Regulations

We supplement the data from the NPRI-ASM dataset with data from Chapter 2 on whether each plant faced regulation under the Canada-Wide Standards for Particulate Matter and Ozone (CWS).[22] The CWS was a major revision to Canadian environmental policy that occurred in the year 2000 as a result of an agreement between the federal government of Canada and the provinces. The policy was intended to improve air quality across the country by creating air quality standards for fine-scale particulate matter ($PM_{2.5}$) and ground-level ozone (O_3) that applied to each major town or city in Canada.[23] These standards required each CMA to meet an air quality target; those cities with poor ambient air quality were required to adopt stringent environmental regulations, while the remaining CMAs had to ensure that their air quality did not deteriorate. In addition, the CWS designated a set of "targeted" industries that were to be the focus of regulation given that they were viewed as key determinants of poor air quality.[24] Provincial authorities regulated plants that were in targeted industries and violating regions using two-part regulatory constraints. As part of the annual provincial operation permitting system, regulated plants had to either show they were operating using clean production processes, or face a production constraint.

While the CWS regulated emissions of both $PM_{2.5}$ and various O_3 pre-cursors, in this chapter we only focus on its effect on plants that emit $PM_{2.5}$. We make this choice because, due to differences in technical constraints facing emitters of $PM_{2.5}$ and O_3 pre-cursors, only the $PM_{2.5}$ standard appears to have had a meaningful impact on variable production costs. As we show in Chapter 2, emitters of O_3 pre-cursors appear to have responded to the CWS by adopting process changes that produced pollution reductions without increasing variable production costs. Moreover, the existing evidence suggests the fixed costs associated with these process changes were relatively small.[25] As the PHE is based on regulation's effect on production costs, either variable or fixed, we focus on the pollutant where this effect appears to be meaningful. We leave an examination of the O_3 standard's effect on the export decisions

[22]The Canada-Wide Standards for Particulate Matter and Ozone were two of the many environmental standards enacted under the Canada-Wide Standard system. Canada-Wide Standards were created for benzene, mercury, and dioxins and furans, among others.

[23]Under the terms of the agreement, an urban area's status as a major town or city was determined using Statistics Canada's definitions of Census Agglomeration (CA) or Census Metropolitan Area (CMA). For convenience, we use the terms CMA and city to refer to both CAs and CMAs.

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

[25]In Chapter 2, we present a detailed discussion on the technical differences in process changes available to emitters of these two pollutants.

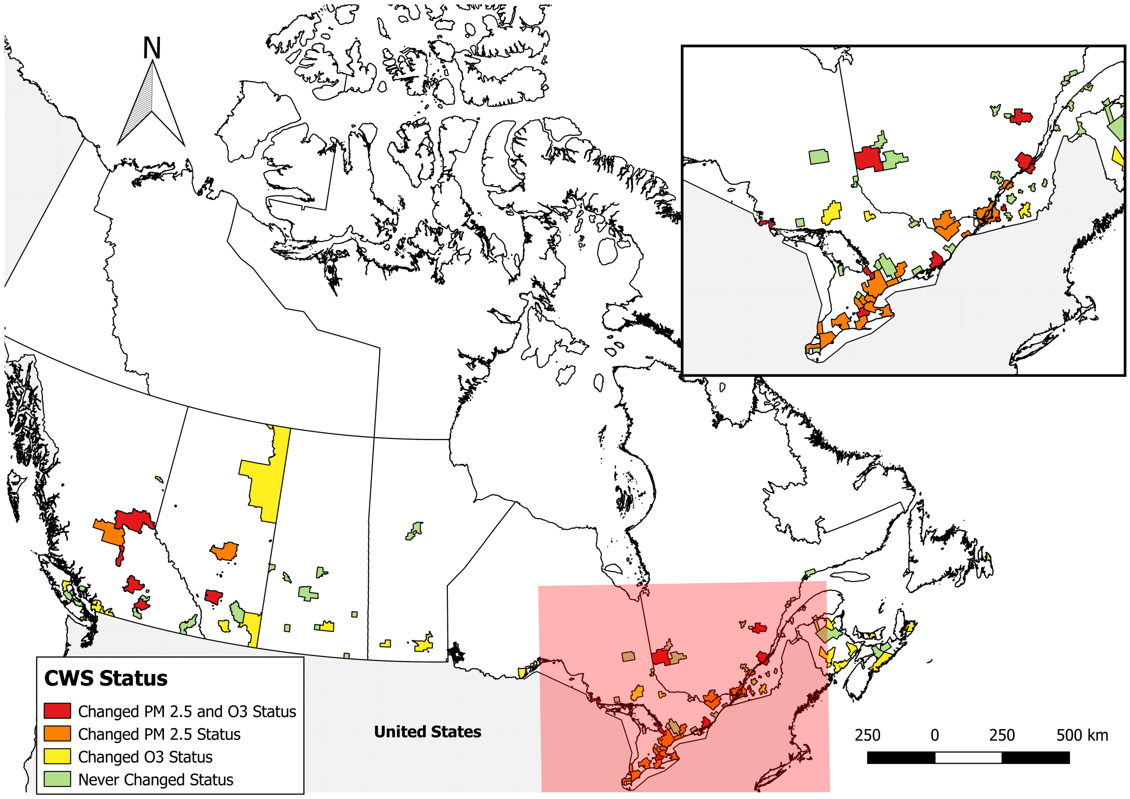

Figure 3.3: Regulatory Status Changes under the CWS

Notes: Figure depicts $PM_{2.5}$ and O_3 standard status changes for each CMA from 2000 to 2010. Red CMAs changed status under both the $PM_{2.5}$ and O_3 standards. Orange CMAs only changed status for the $PM_{2.5}$ standard. Yellow CMAs only changed status for the O_3 standard. Green CMAs didn't change status under either standard. The mainland United States is shown in light gray. Part of the northern Canadian Territories are trimmed for scale. The inset shows detail on the most densely populated area of Canada, colored in light red on the main map. *Source:* Chapter 2.

of the emitters of O_3 pre-cursors for future work.[26]

The variation in environmental regulation created by the CWS is shown in Figure 3.3, which depicts which CMAs were forced to adopt more stringent environmental regulations to address ambient $PM_{2.5}$ and O_3 problems at least once over the period 2000-2010. In the figure, CMAs that adopted more stringent environmental regulations due to ambient pollution concentrations exceeding the relevant air quality standard are depicted in red, orange and

[26]It is worth noting that in our model, only changes in variable costs as a result of regulation directly affect export revenues and a firm's choice to enter the foreign market. However, in principle, general equilibrium changes in wages resulting from regulation could also affect a firm's exports. This means even if regulation only directly affected the fixed-cost of producing, and not variable production costs, changes in equilibrium wages resulting from general equilibrium effects would change exports.

yellow. The red CMAs were required to adopt more stringent policy under both the $PM_{2.5}$ and O_3 standards, while the orange and yellow CMAs were only required to adopt more stringent policy under the $PM_{2.5}$ standard or the O_3 standard, respectively. CMAs depicted in green were not required to adopt more stringent policy under either standard.

As the figure shows, the CWS created substantial variation in environmental regulations across CMAs. Of the 149 CMAs in our sample, 23% adopted new regulations under only the $PM_{2.5}$ standard, 26% adopted new regulations under only the O_3 standard, and 11% adopted new regulations under both standards. We exploit this variation, and the fact that the CWS targeted a subset of industries, to identify the effects of environmental regulation on the export decisions of Canadian manufacturing plants.

3.4 Research Design

To determine the effects of environmental regulation on the export decisions of Canadian manufacturing plants, we adopt the research design we developed in our previous work to study the effects of the CWS on the clean-up of the Canadian manufacturing sector (see Chapter 2). As such, we exploit the variation in regulation created by the design and implementation of the CWS. As shown in Figure 3.3, ambient air quality changes led to variation in the environmental regulations faced by different CMAs over time. Moreover, these regulations targeted a subset of industries in the manufacturing sector, meaning that the regulations varied across industries as well. We use these three sources of variation – across industries, regions, and time – in a triple difference research design to isolate the causal effects of the $PM_{2.5}$ standard on the export decisions of affected plants.[27]

Our approach starts by exploiting the variation in $PM_{2.5}$ regulation over time. To that end, our research design compares the average outcomes from plants in targeted industries located in regulated CMAs (the plants "treated" by $PM_{2.5}$ regulation) while regulated to their outcomes while unregulated. This comparison allows us to control for any unobserved time-invariant plant, industry or CMA-specific heterogeneity that would otherwise confound the effects of regulation. We then exploit the variation in regulation across industries, by comparing the average outcomes from plants in targeted industries to the average outcomes from plants in non-targeted industries located in the same CMAs in the same year. This allows us to control for any unobserved time-varying CMA specific heterogeneity, such as localized recessions, that might affect the decision to export. We then exploit the variation in regulation across regions by taking plants in the same industry in the same year, and comparing the av-

[27] As we noted in Section 3.3, while the CWS regulated both $PM_{2.5}$ and O_3 pre-cursors, we focus on the $PM_{2.5}$ standard for our analysis.

erage outcomes from plants in regions that violate the $PM_{2.5}$ standard, at some point in time, to the average outcomes from plants in non-violating regions. This allows us to control for any time-varying industry heterogeneity, such as foreign demand shocks, that would otherwise confound identification. Finally, our approach compares the average outcomes from treated plants to the average outcomes from plants from non-targeted industries located in CMAs that did not experience a change in $PM_{2.5}$ regulation. These plants serve as a counterfactual that allow us to capture the effects of any unobserved aggregate shocks, such as changes in technology or exchange rate fluctuations, common across all facilities in the country.

3.4.1 Empirical Specification

We implement this research design by estimating several variants of the following equation:

$$y_{pijt} = \beta T_{ijt}^{PM_{2.5}} + \rho_p + \mu_{jt} + \lambda_{it} + \varepsilon_{pijt} \tag{3.20}$$

where y_{pijt} is the outcome of interest (either an indicator of export status, total shipments, or export sales) at plant p, in industry i, located in CMA j at time t. $T_{ijt}^{PM_{2.5}}$ is a treatment indicator for the particulate matter (PM) standards implemented under the CWS. This indicator takes the value one for plants that are in industries targeted by the CWS for years in which their CMA exceeds the relevant pollution threshold. The ρ_p are plant fixed effects that capture any time-invariant plant specific heterogeneity. The μ_{jt} are CMA×year fixed effects that capture any time-varying region specific heterogeneity, such as localized recessions. The λ_{it} are industry×year fixed effects that capture time-varying industry heterogeneity, such as demand shocks. Finally, ε_{pijt} captures idiosyncratic changes in outcomes across plants.

The coefficient of interest in (3.20) is β. This coefficient capture the average percentage difference in outcomes across plants that were affected by the $PM_{2.5}$ standard relative to those that were not, and is identified from within-plant comparisons over time. These comparisons will identify the causal effect of environmental regulations if there are no other factors aside from the CWS particulate matter standard driving differences in export behaviours across plants over time. As we discuss in Chapter 2, there are two reasons to believe this is the case. First, the CWS air quality standards were set federally, meaning that they are unrelated to local tastes, characteristics and economic conditions. Second, ambient pollution levels in a CMA do not necessarily reflect local economic activity due to the fact that particulate matter can be transported long distances via wind patterns. These facts suggest that treatment is exogenous.

While equation (3.20) will produce estimates of the average effect of the CWS particulate matter standard, our model predicts that the effects of regulation will differ across plants on the basis of their productivity level. To investigate this heterogeneity, we adopt the approach first used by Bustos (2011) to study the differential effects of trade liberalization across plants on the basis of their initial productivity. Specifically, we estimate:

$$y_{pijt} = \sum_{q=1}^{5} \beta_{B_q} \left[T_{ijt}^{PM_{2.5}} \times B_q \right] + \rho_p + \mu_{jt} + \lambda_{it} + \varepsilon_{pijt} \tag{3.21}$$

where B_q is an indicator variable equal to one if plant p is in productivity bin q, β_{B_q} is the effect of the CWS particulate matter standard on plants in productivity bin q, and all other variables are defined as in equation (3.20). Given that the ASM does not include information on plant capital stocks, making it impossible to calculate TFP measures using standard methods, we proxy for initial plant productivity using value added per worker in the first year a plant enters our sample, and construct B_q by dividing plants into productivity bins.[28]

We assign plants into productivity bins according to where they lie on their industry's productivity distribution, using the entire set of plants that emit $PM_{2.5}$. Consequently, B_1, for example, corresponds to the first quintile of the productivity distribution of $PM_{2.5}$ emitters. We refer to these as bins, rather than quintiles, because in some regressions we restrict our sample to plants that are continuing exporters. For these regressions, we do not redefine the productivity bins. As a result, a bin may have less than 20% of the observations for the sample used in these regressions.

3.5 Results

3.5.1 Plant Revenue

We begin our empirical analysis by examining the effects of the CWS particulate matter regulations on total plant revenue (Empirical Prediction 1). While our ultimate interest is in understanding how environmental regulation affected participation in export markets and export revenue, examining total revenue is a useful first step because it provides indirect evidence as to the effects on production costs of complying with the CWS particulate matter regulations. Although the available evidence suggests that these costs are large, we do not observe them directly in our data. However, as stated in Empirical Prediction 1, if the compliance costs are sufficiently high, then the total revenues of affected plants should fall in response to regulation. If the CWS does not affect the revenue of affected plants, then it is unlikely that the costs

[28]To address the possibility that the plant productivity measures capture differences inherent across industries or time, we first demean all productivity measures by regressing initial productivity levels on entry-year and industry fixed effects. We then use the residuals from this regression as our measure of plant productivity.

of the regulations are substantial enough to have affected export decisions.

These results are presented in Table 3.2. The table reports estimates from seven separate regressions; in all cases, the dependent variable is the natural log of total revenue, and each regression is weighted to correct for potential sample bias induced by the procedure used to match the NPRI to the ASM.[29] Throughout, standard errors clustered by industry-CMA are reported in parentheses. The first six specifications reported in columns (1) to (6) are based on equation (3.20); as such, the estimated coefficient reports the average effect of the particulate matter standard on affected plants. Our baseline estimate, reported in column (1), includes plant, industry-year and CMA-year fixed effects. Columns (2) through (6) present evidence of the robustness of PM regulation's effect on total revenues. Column (2) adds an indicator of whether industry i in CMA j was also regulated under the CWS O_3 standard in year t. Column (3) adds an indicator of whether plant p was owned by a foreign parent at time t. Column (4) restricts the sample to exclude the set of plants that only export and do not sell domestically. Finally, columns (5) and (6) include a lagged dependent variable. Given that estimating these specifications with a fixed effects estimator would yield inconsistent estimates, we follow the approach of Arellano and Bond (1991) and adopt a GMM procedure with either one (column (5)) or two (column (6)) lags as instruments.

The last specification, reported in column (7), is based on equation (3.21) and reports the estimated effects of the particulate matter standard by bins corresponding to a plant's initial productivity level. The estimates reported in column (7) allow us to test the heterogeneity in the effects of environmental regulation implied in Empirical Prediction 1.

The estimates reported in the first six columns of Table 3.2 suggest that the CWS particulate matter regulations led to a significant reduction in the total revenue of affected plants. For example, the estimate reported in column (1) indicates that the CWS particulate matter standard is associated with a 10.8% reduction in total revenue at the average affected plant. In addition to being statistically significant, this effect is also economically meaningful; given that the average plant in our sample had revenues of close to \$195 million CAD, this estimate implies that the CWS particulate matter reduced the revenues at the average plant by just over \$21 million CAD. Our preferred estimate, reported in column (2), shows that this effect is robust to controlling for the effects of the other regulation imposed under the CWS. Indeed, adding an indicator of treatment status under the CWS O_3 regulations appears to have no effect on either the point estimate or standard error, suggesting that our estimates of the effect of the $PM_{2.5}$ standard are not capturing the effects of other CWS regulations.

Columns (3) and (4) provide further evidence that the CWS particulate matter standard

[29] Details of the match procedure are available in Chapter 2.

Table 3.2: Environmental Regulations and Plant Revenue

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$PM_{2.5}$ Std.	-0.108^{b}	-0.108^{b}	-0.108^{b}	-0.110^{b}	-0.076^{b}	-0.077^{b}	
	(0.050)	(0.050)	(0.050)	(0.051)	(0.038)	(0.039)	
$PM_{2.5}$ Std.$\times$B1							-0.097^{c}
							(0.058)
$PM_{2.5}$ Std.$\times$B2							-0.286^{a}
							(0.092)
$PM_{2.5}$ Std.$\times$B3							-0.043
							(0.078)
$PM_{2.5}$ Std.$\times$B4							0.048
							(0.059)
$PM_{2.5}$ Std.$\times$B5							-0.045
							(0.065)
O_3 Std.		X	X	X	X	X	X
Foreign Owner			X				
Rest. Sample				X			
R^2	0.224	0.224	0.225	0.225			0.227
AR1					-3.358	-3.402	
AR2					-0.802	0.828	
N	6501	6501	6501	6149	3694	3694	6501

Notes: Table reports estimates of the effects of the CWS on the natural log of manufacturing plant revenue. 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. Column (3) includes an indicator of whether the plant is owned by a foreign company. Column (4) restricts the sample to exclude plants that only sell abroad. Columns (5) and (6) report Arellano-Bond estimates with one and two lags, respectively. In all cases, standard errors are clustered by CMA-industry. ^{c}, ^{b}, and ^{a} denote significance at the 10%, 5%, and 1% level, respectively.

negatively affected manufacturing plant revenues. One concern with our preferred estimate is that it is not just capturing the effects of CWS regulation, but also the effects of foreign ownership, which is time-varying. For example, foreign owners may be able to help offset the effects of a negative shock such as the CWS, in ways not possible for domestically owned plants, by exploiting unique knowledge of their home markets. This type of activity would lead to a downward bias in our estimates. The estimate reported in column (3) shows that our preferred estimate is robust to accounting for this explanation; including an indicator of whether a plant is owned by a foreign entity has no discernible effect.

A second concern with our preferred estimate is that it may be driven by plants that do not sell domestically. While our theoretical framework considers the case where plants either only sell domestically or sell both domestically and export, over 5% of our sample is comprised of plants that do not sell in the Canadian market. In principle, the CWS could have a larger effect on the total revenue of these plants because they are potentially at the largest disadvantage when they are regulated. Unlike plants that sell domestically and potentially compete with other plants that are regulated under the CWS, plants that do not sell in the domestic market only compete with foreign plants that are unaffected by the CWS. As a result, our baseline estimates could be simply capturing the effects of the CWS on this set of plants. However, as the estimate reported in column (4) shows, restricting our sample to exclude the set of plants that only sell abroad has little effect on our results.

A final concern with our preferred estimate is the possibility that the estimating equation is mis-specified because we have failed to account for persistence in revenues due to sunk costs. This is a particularly important consideration in our context because previous work studying export behaviour has emphasized the role of sunk costs in determining the export decisions of plants (i.e. Roberts and Tybout (1997) or Bernard and Jensen (2004)). As such, we adapt the approach taken by Bernard and Jensen (2004) to add a lagged dependent variable to our estimating equation and estimate the resulting specification following the approach of Arellano and Bond (1991). As columns (5) and (6) show, these estimates are similar in magnitude and significance to our preferred estimate.

While the estimates reported in columns (1)-(6) provide robust evidence that the CWS particulate matter caused revenues of affected plants to fall, it is important to note that the estimated coefficients report the average effect of regulation on affected plants. Hence, while these results are broadly supportive of Empirical Prediction 1, they do not reveal if the effects of the CWS regulations differed across plants with different productivity levels.

To address this issue, in column (7) we report estimates from a version of our preferred specification where we allow the effects of the CWS to differ across plants according to their initial productivity bin. These estimates suggest that the the effects of the particulate matter standard are concentrated at plants in the two lowest productivity quintiles. Moreover, the effects are larger for plants closer to the middle of the initial productivity distribution as predicted in Empirical Prediction 1; the estimates reported in rows two and three of column (7) indicate that the CWS caused total revenues to fall by 9.7% at plants in the lowest productivity bin but by 28.6% for plants in the second productivity bin. Recall that the least productive plants in an industry are less likely to export. As a result, as Empirical Prediction 1 describes, they should experience a relatively small reduction in total revenues, in percentage terms, compared to the least productive exporters, as exports fall proportionally more than domestic sales.

3.5.2 Export Revenue

Next, we turn to examine the effects of the CWS on the revenue of exporting plants (Empirical Predictions 2 and 3). As such, we restrict our attention to the set of plants that are continuing exporters and export in all years.

Our estimates of the effects of the CWS on the revenue of exporting plants are reported in Table 3.3 and Table 3.4. For the sake of comparison, we start by first re-estimating the effects of the CWS on the natural log of total revenues for the set of plants that are continuing exporters. These estimates are reported in Table 3.3. Next, we estimate the effects of the CWS on the natural log of export revenues. These estimates are reported in Table 3.4. The regressions reported in each column of both tables correspond to the same column in Table 3.2. As such, column (1)-(6) report estimates based on equation (3.20). Column (1) only includes plant, industry-year and CMA-year fixed effects, while column (2) adds an indicator of whether the plant was regulated under the CWS O_3 standard in a given year, column (3) adds an indicator of foreign ownership and column (4) again restricts the sample to exclude the set of plants that only export and do not sell domestically. Columns (5) and (6) both include a lagged dependent variable and are again estimated using the approach of Arellano and Bond (1991), using either one (column (5)) or two (column (6)) lags as instruments.

Finally, column (7) reports estimates from a specification based on equation (3.21) and reports the estimated effects of the particulate matter standard by initial productivity bin. To make our results comparable across all tables, we maintain the same productivity bins used in Table 3.2. We maintain the same bins, as the goal behind allowing the effects to vary across productivity levels is to trace-out the heterogeneity in plant-responses across the industry's productivity distribution. This means that $B1$ in the following two tables, for example, contains all continuing exporting plants that have initial productivity levels low enough to put them in the first quintile of the productivity distribution defined by the full set of all $PM_{2.5}$ emitting plants. Notice that because the least-productive plants in an industry are less-likely to export, there are relatively few observations in the first productivity bin in the following tables. As a result, there is relatively little power in column (7) of both Table 3.3 and Table 3.4. [30]

The estimates reported in the first six columns of Table 3.3 are very similar to those reported in Table 3.2. For example, our preferred estimate indicates that the CWS is associated with a 10.7% reduction in total revenue from affected continuing exporter plants. This result also appears to be very robust; as the estimates reported in columns (3)-(6) show, controlling for foreign ownership, restricting the sample to exclude plants that do not sell domestically

[30] All regressions are weighted to correct for potential sample bias induced by the procedure used to match the NPRI to the ASM. Standard errors clustered by industry-CMA are reported in parentheses.

Table 3.3: Environmental Regulations and Revenue from Continuing Exporters

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$PM_{2.5}$ Std.	-0.107^c	-0.107^c	-0.107^c	-0.106^c	-0.082^c	-0.083^c	
	(0.055)	(0.056)	(0.056)	(0.056)	(0.047)	(0.047)	
$PM_{2.5}$ Std.×B1							-0.019
							(0.120)
$PM_{2.5}$ Std.×B2							-0.331^b
							(0.130)
$PM_{2.5}$ Std.×B3							-0.122
							(0.081)
$PM_{2.5}$ Std.×B4							-0.014
							(0.065)
$PM_{2.5}$ Std.×B5							-0.094
							(0.074)
O_3 Std.		X	X	X	X	X	X
Foreign Owner			X				
Rest. Sample				X			
R^2	0.322	0.322	0.323	0.316			0.324
AR1					-2.563	-2.571	
AR2					1.274	1.289	
Obs.	4093	4093	4093	3807	2367	2367	4093

Notes: Table reports estimates of the effects of the CWS on the natural log of plant revenues for plants that are continuing exporters. 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. Column (3) includes an indicator of whether the plant is owned by a foreign company. Column (4) restricts the sample to exclude plants that only sell abroad. Columns (5) and (6) report Arellano-Bond estimates with one and two lags, respectively. In all cases, standard errors are clustered by CMA-industry. ^{c}, ^{b}, and ^{a} denote significance at the 10%, 5%, and 1% level, respectively.

and allowing for lagged revenues to account for the possibility of sunk costs has little effect on our point estimates. The main difference between the results in these two tables is that the point estimates in Table 3.3 are less precise, owing to the fact that by focusing on continuing exporters we have fewer observations.

Despite the fact that the restricting our sample to continuing exporters has little effect on our estimates of the average effect of the CWS particulate matter regulations, it appears that the effect is driven by a different set of plants. This can be seen from the estimates reported in column (7) of Table 3.3. The particulate matter standard now appears to have little effect on plants with the lowest productivity levels. Instead the effect appears to be driven by plants in the second and third productivity quantiles of the initial productivity distribution. The estimate reported in row three indicates the CWS particulate matter standard is associated with a

Table 3.4: Environmental Regulations and Export Revenue from Continuing Exporters

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$PM_{2.5}$ Std.	-0.219^b	-0.219^b	-0.219^b	-0.242^b	-0.158	-0.157	
	(0.055)	(0.056)	(0.056)	(0.056)	(0.108)	(0.108)	
$PM_{2.5}$ Std.$\times$B1							-0.128
							(0.131)
$PM_{2.5}$ Std.$\times$B2							-0.423^a
							(0.156)
$PM_{2.5}$ Std.$\times$B3							-0.243
							(0.164)
$PM_{2.5}$ Std.$\times$B4							-0.154
							(0.144)
$PM_{2.5}$ Std.$\times$B5							-0.182
							(0.160)
O_3 Std.		X	X	X	X	X	X
Foreign Owner			X				
Rest. Sample				X			
R^2	0.288	0.289	0.289	0.297			0.289
AR1					-3.952	-3.952	
AR2					1.474	1.461	
Obs.	4093	4093	4093	3807	2367	2367	4093

Notes: Table reports estimates of the effects of the CWS on the natural log of export revenues for plants that are continuing exporters. 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. Column (3) includes an indicator of whether the plant is owned by a foreign company. Column (4) restricts the sample to exclude plants that only sell abroad. Columns (5) and (6) report Arellano-Bond estimates with one and two lags, respectively. In all cases, standard errors are clustered by CMA-industry. ^{c}, ^{b}, and ^{a} denote significance at the 10%, 5%, and 1% level, respectively.

33.1% reduction in revenues from plants in the second productivity quintile. Furthermore, the estimate reported in row four suggests the standard reduced revenues at plants in the third productivity quintile by 12.2%, although this effect is imprecisely estimated and not statistically significant at conventional levels. These are intuitive findings, as our model suggests the least productive plants in an industry should not export. As a result, there should be relatively few observations in the first productivity bin when we restrict the sample to continuing exporters.[31]

The estimates reported in columns (1)-(6) of Table 3.4 suggest that the CWS led to a significant reduction in export revenue from affected plants. For example, our preferred estimate, reported in column (2), indicates that the CWS particulate matter standard is associated with a 21.9% reduction in export revenue from affected manufacturing plants. This estimate also

[31] Recall that we maintain the same productivity bin groupings, regardless of sample restriction.

appears to be quite robust; controlling for foreign ownership (column (3)) or excluding plants that only export (column (4)) has little effect on the estimated effect of regulation. The estimated coefficients reported in columns (5) and (6) also suggest that the CWS had a negative effect on exporting, however, these estimates are imprecisely estimated and not statistically significant at conventional levels. The loss of precision is likely owing to the considerable reduction in sample size from this restriction.

The estimates reported in column (7) suggest that the reduction in export revenue is driven by the responses of plants in the second and third productivity quintiles. The estimate reported in row three indicates that the CWS particulate matter standard reduced the export revenue of plants in the second productivity quintile by 42.3%. The estimate reported in row four suggests that the standard reduced export revenues at plants in the third productivity quintile by 24.3%, but as is the case with total revenues, this effect is imprecisely estimated and not statistically significant at conventional levels.

Taken together, the estimates reported in Table 3.3 and Table 3.4 are supportive of Empirical Predictions 2 and 3. The estimates indicate that the CWS had a much larger effect on the export revenue than the total revenue of affected plants, which is consistent with environmental regulation causing a larger reduction in export revenues than domestic revenues, as predicted by our model. In addition, the effects of regulation are concentrated on plants in the second and third productivity quintiles, which is consistent with the model's prediction that the effects of regulation will be largest at the least productive exporters.

3.5.3 Export Status

Finally, we examine the effects of the CWS on plant export status (Empirical Prediction 4). If the costs associated with environmental regulation are large enough to reduce the total revenues of affected plants, as we have shown above, then our model suggests that some plants should exit the export market. To test this prediction, we again turn to examine our full sample of plants, that includes both exporting and non-exporting plants.

The results of estimating a linear probability model of the effects of the CWS particulate matter regulations on an indicator of plant export status are presented in Table 3.5. We again report estimates from seven specifications. The first six, reported in columns (1)-(6) respectively, are based on equation (3.20), and report the average effect of particulate matter regulation on export status at affected plants. The seventh specification, reported in column (7), is based on equation (3.21), and reports the effects of the regulation by initial productivity quintile. Again, column (1) includes plant, industry-year and CMA-year fixed effects only, while column (2) adds an indicator of whether the plant was regulated under the CWS O_3

Table 3.5: Environmental Regulations and Export Status

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$PM_{2.5}$ Standard	-0.018	-0.018	-0.019	-0.013	-0.055	-0.053	
	(0.026)	(0.026)	(0.026)	(0.026)	(0.038)	(0.038)	
$PM_{2.5}$ Std.$\times$B1							-0.008
							(0.037)
$PM_{2.5}$ Std.$\times$B2							-0.102[b]
							(0.045)
$PM_{2.5}$ Std.$\times$B3							0.020
							(0.033)
$PM_{2.5}$ Std.$\times$B4							0.051
							(0.060)
$PM_{2.5}$ Std.$\times$B5							-0.003
							(0.058)
O_3 Standard		X	X	X	X	X	X
Foreign Ownership			X				
Restricted Sample				X			
R^2	0.129	0.129	0.130	0.140			0.131
AR1					-2.545	-2.702	
AR2					1.225	1.325	
Observations	6501	6501	6501	6149	3694	3694	6501

Notes: Table reports estimates of the effects of the CWS on an indicator of plant export status. 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. Column (3) includes an indicator of whether the plant is owned by a foreign company. Column (4) restricts the sample to exclude plants that only sell abroad. Columns (5) and (6) report Arellano-Bond estimates with one and two lags, respectively. In all cases, standard errors are clustered by CMA-industry. [c], [b], and [a] denote significance at the 10%, 5%, and 1% level, respectively.

standard in a given year, column (3) adds an indicator of foreign ownership and column (4) restricts the sample to exclude the set of plants that only export and do not sell domestically. Columns (5) and (6) both include a lagged dependent variable and are again estimated using the approach of Arellano and Bond (1991), using either one (column (5)) or two (column (6)) lags as instruments.[32]

The estimates reported in the first six columns of Table 3.5 suggest that the average effect of the CWS on affected plants is small and insignificant. For example, our preferred estimate, reported in column (2), shows that the CWS only reduced the likelihood of a plant exporting by less than 2%. Not only is this effect economically small, it is statistically insignificant. This is still true when we allow for differences in foreign ownership, restrict our sample to exclude

[32]All regressions are weighted to correct for potential sample bias induced by the procedure used to match the NPRI to the ASM. Standard errors clustered by industry-CMA are reported in parentheses.

foreign exporters, and allow for the possibility of sunk costs in exporting using a dynamic panel specification.

While our estimates of the average effect of the CWS on affected plants are small and insignificant, the estimates reported in column (7) suggest that they mask substantial heterogeneity in how plants respond to regulation. Specifically, the estimates reported in column (7) show a 10% reduction in the likelihood of exporting for plants in the second quintile of the plant productivity distribution as a result of the CWS. This result is consistent with our model's prediction that the least productive exporting firms will exit the export market in response to regulation. The intuition behind this heterogeneity, as our model highlights, is that the marginal exporter should lie somewhere in the middle of an industry's productivity distribution. As a result, exit from exporting should be restricted to this group of firms.

3.6 Conclusion

In this chapter we present plant-level evidence of the heterogeneous effects of environmental regulation on the international competitiveness of domestic industry. The concern that stringent environmental policy may impede the competitiveness of domestic producers, referred to as the Pollution Haven Effect (PHE), has been present in debates on environmental regulation for decades (Jaffe et al., 1995). The PHE stems from the observation that policy imposed unilaterally by one country should increase production costs for domestic producers relative to their unregulated foreign counterparts.

Thus far, the empirical literature on the PHE has focused exclusively on the industry- or region-level effects of regulation. We argue, however, that if producers have differing productivity-levels, then the PHE should vary across firms or plants on the basis of their productivity. This heterogeneity has implications for both identifying the PHE, and understanding the mechanisms through which regulation disadvantages domestic industry.[33]

We use a simple model to clarify the logic underpinning the heterogeneous nature of the PHE. In this model, firms that differ on the basis of their productivity levels compete via monopolistic competition and face a two-part regulatory constraint on their pollution emissions. The model shows that, for an exporting firm of a given productivity level, regulation should cause a larger reduction in exports relative to domestic sales because domestic prices adjust to alleviate the effects of regulation. Moreover, under the type of policy we study, there should

[33]Note that this heterogeneity is particularly pronounced under the type of policy we study, which is a common form of air pollutant regulation (for example, the US Clean Air Act features this type of policy). This policy imposes a two-part regulatory constraint, that requires firms to adopt cleaner production processes, and penalizes those that fail to do so.

be a U-shaped relationship between the effect of regulation on firm revenues and the firm's productivity. This occurs because only relatively productive firms select into exporting, which exposes them to larger loses from regulation, but the most productive of these firms adopt new technology to avoid increased production costs from regulation. Finally, we show that regulation causes some firms to leave the export market, and these firms should be in the middle of the industry's productivity distribution.

Using a unique dataset that contains plant-level production and pollution information, we test for the PHE by examining the effects of a major Canadian environmental policy on the export participation, export volumes, and total sales of Canadian manufacturing plants over the period 2004-2010. This policy, called the Canada Wide Standards for Particulate Matter and Ozone (CWS), implemented regional air quality standards in every major town or city of Canada. In addition, the CWS explicitly targeted plants in a select group of industries. As a result, the CWS created variation in regulatory stringency across industries, regions, and time, which we exploit using a triple-difference research design.

We find the CWS had a large negative effect on both total sales and export volumes, causing an 11% reduction in total sales for the average affected plant and a 22% reduction in export volumes for the average affected exporter. The policy, however, had no significant effect on the selection into or out of exporting for the average plant. As our model predicts, the effects of regulation varied considerably across plants of different productivity levels. Our results show a U-shaped relationship between a plant's productivity and the effect of regulation on both total sales and export volumes. Finally, in line with our model's prediction, the CWS reduced the probability of exporting by 10% for moderately productive plants.

Taken together, these results suggest that environmental regulations that increase variable production costs may reduce the international competitiveness of affected plants. This finding is consistent with recent empirical work on the PHE effect, such as Levinson and Taylor (2008), Shi and Xu (2018), and Broner et al. (2012). However, unlike the previous literature, we show evidence that regulation affects trade flows using plant-level data, rather than aggregate data. Importantly, our results also show that, at least in the context of a CWS-type regulation, there seems to be considerable heterogeneity in the PHE; exporters that feel the competitiveness effects of regulation are in the middle of their industry's productivity distribution.

Conclusion

This thesis provides new theoretical and empirical evidence of how environmental regulation affects manufacturing facilities. This new evidence is important for informing debates on environmental policy, both on the effects these policies have on the economy, and on the design of regulations targeting firms. In addition, this new evidence provides insights into the workings of firms, and in particular on how firms respond to regulatory constraints.

In the first chapter, we present a theoretical model to show the channels through which environmental regulation causes a reduction in an industry's pollution intensity (measured as the amount of pollution emitted per dollar of output). We study regulation that imposes a two-part regulatory constraint on firms: they must either adopt clean production processes, or face a penalty. While a common form of environmental policy, two-part regulatory constraints have not been studied in the theoretical literature on environmental regulation.

Our model shows that this type of regulation causes the least productive firms in an industry to exit, which we call the selection channel, and low-productivity surviving firms to produce less (a reallocation channel). If these affected firms are relatively pollution intensive, then these channels will serve to reduce an industry's pollution intensity. In addition, this type of regulation causes moderately-productive firms to adopt cleaner production processes, which reduces an industry's pollution intensity through what we call the process channel.

In the second chapter, we estimate the regulatory channels of the manufacturing clean-up. While there is much indirect evidence that suggests regulation may affect an industry's pollution intensity through several channels, work has yet to directly estimate the magnitude of these channels. In this chapter, we use a novel confidential dataset that contains plant-level pollution and production information for major manufacturing polluters in Canada to estimate the three plant-level channels through which regulation contributes to a clean-up. With this data, we estimate the effect of a major revision to environmental policy in Canada, called the Canada Wide Standards for Particulate Matter and Ozone, on plant pollution intensity, output, and entry and exit decisions. This policy created variation in regulatory stringency across industries, regions, and time, which we use to identify the effect of environmental regulation.

We find the Canada Wide Standards played a sizeable role in the Canadian manufacturing clean-up. From 2004 to 2010, this policy explains approximately 60% of the observed reduction in sector-level nitrogen oxide emissions and 20% of the reduction in particulate matter emissions. Moreover, the channels involved in the CWS clean-up varied starkly across pollutants. The clean up of nitrogen oxide was primarily caused by the process channel, while the channels driving the particulate matter clean-up were primarily selection and reallocation. We argue these differences arise because of differences across pollutants in the fixed costs plants need to pay to adopt cleaner production processes, and show additional empirical evidence consistent with this hypothesis.

In the third chapter, we ask whether environmental regulation affects the international competitiveness of manufacturing plants. By raising the costs of production in domestic markets relative to unregulated foreign markets, environmental regulation may disadvantage domestic producers relative to their foreign counterparts. This hypothesis, referred to as the Pollution Haven Effect (PHE), is typically tested by examining regulation's effect on the international flow of goods. Thus far, the literature on this topic has assessed the PHE at the industry or regional level. We contribute to this literature by asking how regulation affects the export decisions of individual plants.

We provide a theoretical framework to show how regulation affects a plant's decision of whether to export, as well as their export volumes. Our theoretical framework shows that a regulation that imposes a two-part regulatory constraint on firms should have differential effects on firms of different productivity levels. In particular, the PHE should be strongest for the least-productive exporters.

Finally, we estimate the effect of the same policy-change studied in the second chapter, the Canada Wide Standards, on the export behaviours of Canadian manufacturing plants. In line with our theory's predictions, we find the PHE appears strongest for the least productive exporting plants. Those low-productivity exporters experience a 10% reduction in the probability of exporting, as well as a 40% reduction in export sales.

Bibliography

Aichele, R. and G. Felbermayr (2015). Kyoto and carbon leakage: an empirical analysis of the carbon content of bilateral trade. *Review of Economics and Statistics 97*(1), 104–115.

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.

Andersen, D. C. (2018). Accounting for loss of variety and factor reallocations in the welfare cost of regulations. *Journal of Environmental Economics and Management 88*, 69–94.

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

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

Antweiler, W., B. R. Copeland, and M. S. Taylor (2001). Is free trade good for the environment? *American Economic Review 91*(4), 877–908.

Applied Technologies of New York Inc. (2018). Case Study: East River Housing. http://www.atiofny.com/case-study-east-river-housing/. [Online; accessed 5-May-2018].

Arellano, M. and S. Bond (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies 58*(2), 277–297.

Barnatchez, K., L. D. Crane, and R. A. Decker (2017). An Assessment of the National Establishment Time Series (NETS) Database.

Barrows, G. and H. Ollivier (2018). Cleaner firms or cleaner products? how product mix shapes emission intensity from manufacturing. *Journal of Environmental Economics and Management 88*, 134–158.

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

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

Berman, E. and L. Bui (2001a, feb). Environmental regulation and labor demand: evidence from the South Coast Air Basin. *Journal of Public Economics 79*(2), 265–295.

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

Bernard, A. B., J. Eaton, J. B. Jensen, and S. Kortum (2003). Plants and productivity in international trade. *American Economic Review 93*(4), 1268–1290.

Bernard, A. B. and J. B. Jensen (2004). Why some firms export. *Review of Economics and Statistics 86*(2), 561–569.

Besley, T. and A. Case (2000). Unnatural Experiments? Estimating the Incidence of Endogenous Policies. *The Economic Journal 110*(467), 672–694.

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

Broner, F., P. Bustos, and V. Carvalho (2012). Sources of comparative advantage in polluting industries. NBER Working Paper 18337.

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

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

Canadian Council of Ministers of the Environment (1998a). National Emission Guideline for Cement Kilns. https://www.ccme.ca/files/Resources/air/emissions/pn_1284_e.pdf.

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

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

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

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

Canadian Council of Ministers of the Environment (2006). Canada-wide Standards for Particulate Matter and Ozone: Five Year Report: 2000-2005. https://www.ccme.ca/files/Resources/air/pm_ozone/pm_oz_2000_2005_rpt_e.pdf.

Cao, J., L. D. Qiu, and M. Zhou (2016). Who invests more in advanced abatement technology? theory and evidence. *Canadian Journal of Economics 49*(2), 637–662.

Chang, T., J. G. Zivin, T. Gross, and M. Neidell (2016). Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy 8*(3), 141–169.

Chay, K. Y. and M. Greenstone (2003). The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession. *Quarterly Journal of Economics 118*, 1121–1167.

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

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

Copeland, B. and M. Taylor (2004). Trade, growth, and the environment. *Journal of Economic Literature XLII*, 7–71.

Cui, J., H. Lapan, and G. Moschini (2012). Are Exporters More Environmentally Friendly than Non-Exporters? Theory and Evidence.

Cui, J., H. Lapan, and G. Moschini (2015). Productivity, Export, and Environmental Performance: Air Pollutants in the United States. *American Journal of Agricultural Economics 98*.

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

Demidova, S. and A. Rodríguez-Clare (2009). Trade policy under firm-level heterogeneity in a small economy. *Journal of International Economics 78*(1), 100–112.

Demidova, S. and A. Rodríguez-Clare (2013). The simple analytics of the Melitz model in a small economy. *Journal of International Economics 90*(2), 266–272.

Ederington, J., A. Levinson, and J. Minier (2005). Footloose and pollution-free. *Review of Economics and Statistics 87*(February), 92–99.

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

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

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

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

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

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

Environment Canada and Forest Products Association of Canada (2004). Towards More Innovative Air Quality Management: Proposal for a Pulp and Paper Air Quality Forum. http://publications.gc.ca/site/eng/274270/publication.html.

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

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

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

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

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

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

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

Foster, L., J. Haltiwanger, and C. J. Krizan (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In 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*, 965–993.

Fowlie, M. L., M. Reguant, and S. P. Ryan (2016). Measuring Leakage Risk. Mimeo.

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

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

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

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 & 1977 Clean Air Act amendments and the census of manufacturers. *Journal of Political Economy 110*(6), 1175–1219.

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

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

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

Haltiwanger, J. C., R. S. Jarmin, and J. Miranda (2013). Who Creates Jobs? Small vs. Large vs. Young. *Review of Economics and Statistics 95*(2).

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

Hanna, R. and P. Oliva (2015). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics 122*, 68–79.

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

Holladay, J. S. (2016). Exporters and the Environment. *Canadian Journal of Economics 49*.

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.

Jaffe, A. B., S. R. Peterson, P. R. Portney, and R. N. Stavins (1995). Environmental Regulation and the Competitiveness of U.S. Manufacturing: What does the evidence tell us? *Journal of Economic Literature 33*(1), 132–163.

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

Kellenberg, D. K. (2009). An empirical investigation of the pollution haven effect with strategic environment and trade policy. *Journal of International Economics 78*(2), 242–255.

Knight, C. (2013). What is grandfathering? *Environmental Politics 22*(3), 410–427.

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

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

Lemoine, D. (2017). Escape from Third-Best: Rating Emissions for Intensity Standards. *Environmental and Resource Economics 67*(4), 789–821.

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

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

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

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

Li, Z. and S. Shi (2017). Emission taxes and standards in a general equilibrium with productivity dispersion and abatement. *Macroeconomic Dynamics 21*(8), 1857–1886.

Li, Z. and J. Sun (2015). Emission taxes and standards in general equilibrium with entry and exit. *Journal of Economic Dynamics and Control 61*, 34–60.

List, J. A., D. L. Millimet, P. G. Fredriksson, and W. W. McHone (2003). Effects of environmental regulations on manufacturing plant births: Evidence from a propensity score matching estimator. *Review of Economics and Statistics 85*(4), 944–952.

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.

McAusland, C. and N. Najjar (2015). Carbon Footprint Taxes. *61*(1), 37–70.

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

Millimet, D. and J. Roy (2016). Empirical Tests of the Pollution Haven Hypothesis When Environmental Regulation is Endogenous. *Journal of Applied Econometrics 31*, 652–677.

Roberts, M. J. and J. R. Tybout (1997). The decision to export in colombia: An empirical model of entry with sunk costs. *The American Economic Review 87*(4), 545–564.

Schlenker, W. and W. R. Walker (2016). Airports, air pollution, and contemporaneous health. *Review of Economic Studies 83*(2), 768–809.

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

Shi, X. and Z. Xu (2018). Environmental regulation and firm exports: Evidence from the eleventh Five-Year Plan in China. *Journal of Environmental Economics and Management 89*, 187–200.

Tang, E., J. Zhang, and Z. Haider (2015). Firm productivity, pollution, and output: theory and empirical evidence from China. *Environmental Science and Pollution Research 22*, 18040–18046.

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.

World Bank Group (1998a). Airborne Particulate Matter: Pollution Prevention and Control. https://www.ifc.org/wps/wcm/connect/ab848080488557a8bd44ff6a6515bb18/HandbookAirborneParticulateMatterPollutionPreventionAndControl.pdf?MOD=AJPERES.

World Bank Group (1998b). Nitrogen Oxides: Pollution Prevention and Control. https://www.ifc.org/wps/wcm/connect/60add78048855368aeecfe6a6515bb18/HandbookNitrogenOxidesPollutionPreventionAndControl.pdf?MOD=AJPERES.

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

Appendix A

Supporting Materials

A.1 Chapter 1 Appendix

A.1.1 Proof of Proposition 3

Consider firms that retrofit from technology b to technology r. For these firms, the effect of regulation can be determined by comparing equations (1.28) with (1.31):

$$\frac{r_r^{tar}(\varphi)}{r_b^n(\varphi)} = \left[\frac{\varphi_\varepsilon^{no}\left[1+\tau\kappa\right]}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1} \tag{A.1}$$

Note that

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

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

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

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 (1.29) and (1.32). This yields

$$\frac{r_s^{tar}(\varphi)}{r_s^{no}(\varphi)} = \left[\frac{\varphi_\varepsilon^{no}[1+\tau\kappa]}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1}, \tag{A.2}$$

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

$$\frac{r_s^{tar}(\varphi)}{r_b^{no}(\varphi)} = \left[\frac{\varphi_\varepsilon^{no}[1+\tau\kappa]\,\alpha}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1}. \tag{A.3}$$

This is greater than one if and only if

$$f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{\alpha^{\frac{1}{k}}[1+\tau\kappa]^{\frac{1}{k}}-1}{\alpha^{\frac{1}{k}}[1+\tau\kappa]^{\frac{1}{k}}}\right]\left[\left[\frac{\left[[1+\tau\kappa]^{\sigma-1}-1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right]\left[\frac{1}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}\frac{1}{\alpha^{\frac{1}{k}}} + \left[[1+\tau\kappa]^{\frac{k^2-1}{k}}\left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}}-\alpha^{\frac{1}{k}}\right]\frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}}\frac{1}{\alpha^{\frac{1}{k}}}\right]^{-1}, \tag{A.4}$$

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

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

$$\frac{r_r^{tar}(\varphi)}{r_s^n(\varphi)} = \left[\frac{\varphi_\varepsilon^n [1+\tau\kappa]}{\varphi_\varepsilon^{tar}\alpha}\right]^{\sigma-1}, \tag{A.5}$$

which is less than one if and only if

$$\begin{aligned} f^{\frac{k-\sigma+1}{\sigma-1}} > \left[\frac{[1+\tau\kappa]^{\frac{1}{k}} - \alpha^{\frac{1}{k}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right] \Bigg[\left[\frac{\left[[1+\tau\kappa]^{\sigma-1}-1\right]^{\frac{k}{\sigma-1}}}{[1+\tau\kappa]^{\frac{1}{k}}}\right] \left[\frac{1}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} \alpha^{\frac{1}{k}} \\ + \left[[1+\tau\kappa]^{\frac{k^2-1}{k}} \left[\frac{f_s}{f_s-f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} - \frac{1}{\alpha^{\frac{1}{k}}}\right] \frac{\Delta_1^{\frac{k}{\sigma-1}}}{f_s^{\frac{k-\sigma+1}{\sigma-1}}} \alpha^{\frac{1}{k}} \Bigg]^{-1}. \end{aligned} \tag{A.6}$$

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 $r_r^{tar}(\varphi)/r_s^n(\varphi) < 1$.

A.1.2 Proof of Corollary 1

Notice that Equation (1.33) implies $\left[\frac{\varphi_\varepsilon^{no}}{\varphi_\varepsilon^{tar}}\right]^{\sigma-1} = \frac{r_b^{tar}(\varphi)}{r_b^{no}(\varphi)}$. Substituting this into Equation (A.1) and Equation (A.2) gives

$$\frac{r_r^{tar}(\varphi)}{r_b^n(\varphi)} = \frac{r_s^{tar}(\varphi)}{r_s^{no}(\varphi)} = \left[\frac{r_b^{tar}(\varphi)}{r_b^{no}(\varphi)}\right][1+\tau\kappa]^{\sigma-1},$$

which means the relative reduction in revenues for firms that always use business-as-usual technology is larger than that of the firms that retrofit or always use state-of-the-art technology.

If $\varphi_s^{tar} < \varphi_s^{no}$, then some firms switch from business-as-usual technology to the state-of-the-art technology. Equation (1.33) shows that revenues rise for these firms.

If $\varphi_s^{tar} > \varphi_s^{no}$, then regulation causes some facilities to downgrade from state-of-the-art to the retrofitted technology. Substituting Equation (1.33) into Equation (A.5) gives

$$\frac{r_r^{tar}(\varphi)}{r_s^{n}(\varphi)} = \left[\frac{r_b^{tar}(\varphi)}{r_b^{no}(\varphi)}\right]\left[\frac{1+\tau\kappa}{\alpha}\right]^{\sigma-1}, \tag{A.7}$$

which is less than one by assumption. Notice that if we allowed α to be greater than $1+\tau\kappa$, then firms with $\varphi \in [\varphi_s^{no}, \varphi_s^{tar}]$ would experience the largest reduction in revenues.

A.1.3 Proof of Proposition 5

Differentiating the retrofitting cut-off with respect to f_r gives

$$\begin{aligned}\frac{\partial \varphi_r^{tar}}{\partial f_r} &= \left[\frac{1}{1+\tau\kappa}\right]\left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}\left[\frac{\partial \varphi_\varepsilon^{tar}}{\partial f_r} + \left[\frac{1}{\sigma-1}\right]\frac{\varphi_\varepsilon^{tar}}{f_r}\right] \\ &= \left[\frac{1}{1+\tau\kappa}\right]\left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}\left[\frac{1}{k[\varphi_\varepsilon^{tar}]^k}\frac{\partial[\varphi_\varepsilon^{tar}]^k}{\partial f_r} + \left[\frac{1}{\sigma-1}\right]\frac{\varphi_\varepsilon^{tar}}{f_r}\right].\end{aligned} \tag{A.8}$$

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

$$\frac{\partial\,[\varphi_\varepsilon^{tar}]^k}{\partial f_r} > -\left[\frac{k}{\sigma-1}\right]\frac{[\varphi_\varepsilon^{tar}]^k}{f_r}, \tag{A.9}$$

where $\frac{\partial[\varphi_\varepsilon^{tar}]^k}{\partial f_r} = [\frac{f}{\delta f}][1+\tau\kappa]^k f^{\frac{k-\sigma+1}{\sigma-1}}\left[[\frac{\Delta_1}{f_s-f_r}]^{\frac{k}{\sigma-1}} - [\frac{\Delta_2}{f_r}]^{\frac{k}{\sigma-1}}\right]$. One can show that Equation (A.9) reduces to

$$\begin{aligned}\left[\frac{k}{k-\sigma-1}\right]\frac{1}{f_r}\left[\frac{1}{f}\right]^{\frac{k-\sigma+1}{\sigma-1}}\left[\frac{1}{1+\tau\kappa}\right]^{\sigma-1} + \left[1+\left[\frac{k}{k-\sigma-1}\right]\frac{1}{f_r}\frac{1}{f_s-f_r}\right]\left[\frac{\Delta_1}{f_s-f_r}\right]^{\frac{k}{\sigma-1}} > \\ -\left[\frac{\sigma-1}{k-\sigma+1}\right]\left[\frac{\Delta_2}{f_r}\right]^{\frac{k}{\sigma-1}},\end{aligned} \tag{A.10}$$

which is always satisfied.

In addition, lowering f_r lowers the exit cut-off under the regulation regime if f_s isn't too large. Differentiating $\varphi_\varepsilon^{tar}$ with respect to f_r gives

$$\frac{\partial \varphi_\varepsilon^{tar}}{\partial f_r} = \left[\frac{\sigma-1}{k-\sigma+1}\right]^{\frac{1}{k}} \left[\frac{f}{\delta f_\varepsilon}\right]^{\frac{1}{k}} \left[\Lambda^{tar}\right]^{\frac{1-k}{k}} [1+\tau\kappa]^k \left[\frac{k-\sigma+1}{\sigma-1}\right] f^{\frac{k-\sigma+1}{\sigma-1}} \left[\Delta_2^{\frac{k-\sigma+1}{\sigma-1}} \left[\frac{1}{f_s-f_r}\right]^{\frac{k-2[\sigma-1]}{\sigma-1}} - \Delta_1^{\frac{k-\sigma+1}{\sigma-1}} \left[\frac{1}{f_r}\right]^{\frac{k-2(\sigma-1)}{\sigma-1}}\right], \tag{A.11}$$

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 $\varphi_r^{tar} < \varphi_s^{tar}$. If $k < 2[\sigma-1]$, then imposing $f_s > [\frac{\Delta_1+\Delta_2}{\Delta_2}] f_r$ ensures both results.

A.1.4 Proof of Corollary 2

Note that the ratio of φ_r^{tar} to $\varphi_\varepsilon^{tar}$ reflects the measure of firms that use business-as-usual technology in the regulated equilibrium. This ratio is given by

$$\frac{\varphi_r^{tar}}{\varphi_\varepsilon^{tar}} = \left[\frac{1}{1+\tau\kappa}\right] \left[\frac{f_r}{\Delta_2 f}\right]^{\frac{1}{\sigma-1}}. \tag{A.12}$$

The derivative of Equation (A.12) with respect to f_r is positive, which means lowering f_r lowers the measure of surviving firms using business-as-usual technology.

A.1.5 Technology Upgrading

In addition to the results shown in Proposition 4, 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 regulation and no regulation regimes can be written as

$$\frac{\varphi_h^{tar}}{\varphi_h^n} = \frac{1+[1+\tau\kappa]^k \left[\Delta_2^{\frac{k}{\sigma-1}} \left[\frac{f}{f_r}\right]^{\frac{k-\sigma+1}{\sigma-1}} + \Delta_1^{\frac{k}{\sigma-1}} \left(\frac{f}{f_s-f_r}\right)^{\frac{k-\sigma+1}{\sigma-1}}\right]}{[1+\tau\kappa]^{\frac{1}{k}} \left[\frac{f_s-f_r}{f_s}\right]^{\frac{1}{k[\sigma-1]}} \left[1+\Delta_1^{\frac{k}{\sigma-1}} \left[\frac{f}{f_s}\right]^{\frac{k-\sigma+1}{\sigma-1}}\right]}. \tag{A.13}$$

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

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

Hence, the effects of regulation 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.2 Chapter 2 Appendix

A.2.1 Data Appendix

Micro Data

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

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

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

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

For each pollutant, plants are required to report their releases by medium (to air, water, and land), quantities sent for disposal and recycling, methods used to compute releases, and abatement activities[3]. Detailed guidelines on how to compute emissions for each pollutant are provided for each sector and production activity (for a detailed list by sector, see: Environment and Climate Change Canada (2016a)). Each plant is also required to report a number of characteristics, including plant name, business number, industry, and location.

The ASM was used as Statistics Canada's manufacturing census until 2012, and provides longitudinal information for the majority of manufacturing plants in Canada.[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 to do with the representativeness of the matched sample. If the probability of a successful match is non-random, then the matched sample will not be representative of the universe of polluters. This means descriptive statistics from the matched sample will not be reflective of polluters in general. Rather they will be informative about the subset of polluters that were successfully matched.

The second issue is more problematic, as it could lead to biased estimates of the CWS' effects. This issue arises if the match probability is correlated with the CWS' treatment effect.

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

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

Note that if the effect of the CWS is homogenous, then the match probability cannot be correlated with treatment, and the estimated effect of the CWS from the matched data will be an unbiased estimate of the true effect of the CWS. That is, this issue only arises when the effect of treatment varies across plants.

In the case of the CWS, there is substantial heterogeneity in the treatment effects. As we show in the main body of this chapter, the treatment effects vary by plant productivity. Moreover, plant productivity is correlated with plant size, and the probability of a successful match also appears to be correlated with plant size. As a result, the match probability is potentially correlated with the treatment effect. This sample bias induced by the imperfect match should be addressed so as to obtain unbiased estimates of the CWS' effects. We correct for this bias using a simple weighting strategy.

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

$$\beta = Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2}, \tag{A.14}$$

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

The treatment effect in the matched sample is given by

$$\begin{aligned}\beta^{match} &= Pr(g_1|match)\beta^{g_1} + Pr(g_2|match)\beta^{g_2} \\ &= \frac{Pr(match|g_1)Pr(g_1)}{Pr(match)}\beta^{g_1} + \frac{Pr(match|g_2)Pr(g_2)}{Pr(match)}\beta^{g_2},\end{aligned} \tag{A.15}$$

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

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

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

$$\begin{aligned}\beta^{match,weighted} &= \omega_{g1} Pr(g_1|match)\beta^{g_1} + \omega_{g2} Pr(g_2|match)\beta^{g_2} \\ &= Pr(g_1)\beta^{g_1} + Pr(g_2)\beta^{g_2},\end{aligned} \tag{A.16}$$

which is the true treatment effect, β.

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

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

To show the effect of our weighting procedure, Table A.1 compares the average plant emissions of each of the CWS pollutants from the full NPRI, the unweighted matched sample, and the weighted matched sample. The first column shows the mean emissions for the universe of polluters, and the second the percentage differences between the mean emissions in the matched sample using our weighting procedure and the universe of polluters. The third column shows the percentage differences between the mean emissions in the matched sample without weighting and the universe of polluters.

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

A.2.2 Robustness

This section presents a series of additional robustness exercises described in Section 2.4.3 of the main text. We first examine when plants are regulated by the CWS. If the majority of

Table A.1: 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.

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.

Table A.2 shows the fraction of treated plants that are treated early in the policy. Panel A shows the plants treated in the first year of the sample, and Panel B shows the plants treated by the middle of the CWS phase-in. For each standard and pollutant, over half of the treated plants start the sample treated. That fraction increases to between 80% and 90% by 2005 for all standard-pollutant pairs with the exception of the $PM_{2.5}$ emitters treated by the O_3 standard, for which two-thirds are treated by 2005.

Restricting treatment to plants that start the sample treated (dropping all plants treated later from the sample) leaves the results qualitatively unchanged, and actually increases the magnitude of the main effects (though not significantly). The results for the average effect of the CWS on emissions of each pollutant are shown in Table A.3. For this group, the $PM_{2.5}$ standard reduced emissions of $PM_{2.5}$ by 17%, and the O_3 standard reduced emissions of NO_X by 56%.[5] 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 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

[5] Note 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.

Table A.2: Regulation Cohorts

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

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

Table A.3: CWS Effect on Emissions for Initial Treatment Cohort

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

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

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

Table A.4: CWS Effect on Emissions Dropping Large Emitters

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

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

until we reach a 1% threshold.[6] We report the results for emissions in Table A.4. 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 A.4 and the effect in the full sample. The effects of the O_3 standard are also consistent with the main results in this chapter, 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

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

productivity-levels for emissions and output are qualitatively unchanged in each of the size thresholds employed in Table A.4. 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 A.4, 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, 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 A.5. 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 O_3 regulation significantly reduced NO_X emissions from affected plants. Consequently, we conclude our results are unlikely to be reflective of differential trends across large and small emitters.

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

Table A.5: CWS Effect on Emissions with Large Emitter Trends

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

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

mutli- and single-plant firms.[7,8]

We use the parent company name information reported in the NPRI to identify multi-plant firms. This information is entered as a text string, which is 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.[9] We classify firms in two ways. In our first approach we classify firms as multi-plant if they own more than one plant that emit the same pollutant (either $PM_{2.5}$ or NO_X). In our second approach we classify firms as multi-plant if they own more than one plant in our dataset (that is, that emit any of 300 pollutants tracked in the NPRI). 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.

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

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

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

We estimate a version of our main specification in which we include a time-varying indicator that selects all plant-years owned by a multi-plant firm, and an interaction between the multi-plant indicator and our treatment indicators. For $PM_{2.5}$ emitters we estimate the $PM_{2.5}$ standard's effect on plant emissions, and for NO_X emitters we estimate the O_3 standard's effect on plant emissions. These results are reported in Table A.6. 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.

Additional Robustness

Recall that an identification problem exists if there is an unobservable characteristic that varies by CMA-industry-year and is correlated with treatment under the CWS. We believe there are two potential identification problems of concern that our robustness checks may not have addressed. The first is to do with differential trade shocks. If plants in targeted industries are more likely to export, and the CMAs that eventually exceed the CWS are more connected with Canada's major trading partners (i.e. the US), then exchange rate fluctuations would have a larger effect on the treated plants than the untreated plants. This is a potential issue because over the CWS phase-in period the Canadian dollar appreciated significantly with respect to the US dollar (in 2000, one Canadian dollar was worth 67 cents US, but by 2010 one Canadian dollar was worth 97 cents US). This appreciation in the Canadian dollar made Canadian goods more expensive, which could have depressed relatively export-intensive manufacturing plants. Note, however, that this is only an identification problem if the treated plants (those in the targeted industries in dirty regions) are more trade-exposed than the untreated plants. As we have plant-level data on exports to the US, we can test whether this is true. Testing for differences in trade exposure between our treated and untreated groups, we find plants that are eventually treated by the CWS are less export-intensive than those that are untreated.[10] As a result, differential trade shocks should be less costly to the treated plants, and would bias our results upwards, if at all.

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

[10]These results are available on request

A.2.3 Additional Results

The CWS' Effect on Other Pollutants

In this section we present our estimates of the effect of the CWS on plant-level emissions of pollutants not directly regulated by the CWS. These results are useful for two reasons. Firstly, ambient $PM_{2.5}$ and O_3 pollution may be formed through chemical reactions in the atmosphere between other pollutants besides $PM_{2.5}$ and NO_X, in particular other criteria air contaminants (CACs). Secondly, this allows us to assess whether there were positive or negative spillovers in response to the CWS. A positive spillover would occur if plants substitute toward unregulated pollutants, whereas a negative spillover would occur if emissions were were correlated across pollutants. The former is typically referred to as regulation-induced substitution, and the latter as co-pollutant effects.

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

GHG information is not available in the NPRI, however, Environment and Climate Change Canada collect GHG emission data for the largest plants in the country, and is publicly available through the Greenhouse Gas Reporting Program (GHGRP).[11] The GHGRP reports emissions for several GHGs (including carbon dioxide, methane, and nitrogen dioxide), total facility GHGs, and provides a crosswalk file to match plants in the GHGRP with plants in the NPRI. We use this crosswalk file to merge the facility GHG data to the NPRI. Virtually all manufacturing plants in the GHGRP over our sample were successfully matched to the NPRI. We report the effect of the CWS on total GHG emissions, carbon dioxide emissions, and nitrogen dioxide emissions.

The results of these regression are shown in Table A.7. Each column reports our estimates of the effect of the CWS on a different pollutant. The dependent variable in each of these regressions is the natural log of plant pollution emissions for the relevant pollutant, and standard errors clustered at the CMA-industry level are reported in parentheses.

The results in Table A.7 show the $PM_{2.5}$ standard caused a significant drop in PM_{10} emissions. This, to some extent, is a mechanical result: by definition, reported PM_{10} emissions *include* emissions of $PM_{2.5}$. Nonetheless, these results provide added confidence to our main results, and indicate there was no significant substitution from fine to large scale particu-

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

late matter emissions in response to the CWS. $PM_{2.5}$ regulation had only minor effects on emissions of the other air pollutants and greenhouse gases. $PM_{2.5}$ regulation caused a small (insignificant) increase in SO_2, CO, and lead, a small (insignificant) drop in VOCs, and had virtually no effect on GHGs. The only pollutant showing a sizeable response to PM regulation is zinc, which fell by 23%, although this is not significant at conventional levels.

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

Table A.6: CWS Effect on Emissions - Multi-Plant Firms

	Panel A: $PM_{2.5}$			
	Same Pollutant		Any Pollutant	
	Coarse Matching	Fine Matching	Coarse Matching	Fine 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
	Panel B: NO_X			
	Same Pollutant		Any Pollutant	
	Coarse Matching	Fine Matching	Coarse Matching	Fine 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 O_3 standard. The effect of the $PM_{2.5}$ standard is shown for PM emitters, and the O_3 standard is shown for O_3 NO_X emitters. All regressions include plant, industry-year and CMA-year fixed effects. Standard errors are clustered by CMA-industry. Asterisks denote significance at the 1% (***), 5% (**), and 10% (*) levels, respectively.

Table A.7: The Effects of the CWS on Plant Emissions of Unregulated Pollutants

	Panel A: CACs				Panel B: Metals		Panel C: GHGs		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	PM_{10}	VOCs	SO_2	CO	Lead	Zinc	GHGs	CO_2	N_20
$PM_{2.5}$ Std.	-0.17***	-0.07	0.05	0.07	0.05	-0.23	-0.02	0.02	0.04
	(0.06)	(0.08)	(0.21)	(0.09)	(0.32)	(0.18)	(0.19)	(0.19)	(0.15)
O_3 Std.	0.01	-0.22*	-0.51**	-0.44*	-0.09	0.01	-0.37**	-0.32**	-0.66***
	(0.09)	(0.12)	(0.21)	(0.24)	(0.56)	(0.79)	(0.17)	(0.15)	(0.19)
R^2	0.22	0.22	0.45	0.38	0.37	0.37	0.61	0.62	0.68
N	8,003	7,045	2243	3352	1411	1496	701	701	613

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

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

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

where u indexes an unregulated pollutant included in Table A.7, r indexes either $PM_{2.5}$ or NO_X, $\lambda_{u,i}$ is a pollutant-plant fixed-effect, and $\alpha_{u,r}$ is our estimate of the cross-pollutant elasticity between unregulated pollutant u and regulated pollutant r.

We restrict our sample to years before 2006 to try to limit the potential interference of the CWS in changing these cross pollutant elasticities (recall, most of the CWS regulations were implemented between 2005 and 2007). The results from these regressions are shown in Table A.8.

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

A.2.4 CWS Counterfactuals

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

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

Table A.8: Cross-Pollutant Elasticities

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

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

This can be written as

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

With some algebra, this reduces to

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

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

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

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

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

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

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

A.2.5 Policy Details

Nova Scotia

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

New Brunswick

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

Ontario

In 2005, Ontario adopted site-specific air quality standards (regulation O Reg 419/05). These standards targeted many different pollutants, including ozone, ozone pre-cursors (including nitrogen oxides and various volatile organic compounds), and particulate matter[12]. The regulation contained more stringent standards for a number of industries, including several of the industries targeted by the CWS.[13] In addition to more stringent standards, plants in these industries must submit annual emissions reports to the Ontario environment ministry.

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

Quebec

Quebec developed the Clean Air Regulations – which included local air quality standards and site-specific emissions standards – during the phase-in period (regulation QLR Q-2, r 4.1). Air quality standards were developed for a large number of pollutants, including ozone and particulate matter (the $PM_{2.5}$ standard was set at the level of the CWS and the ozone standard was set slightly more stringent than the CWS at 62.5 ppb). Emissions standards were developed for many different industries and industrial processes, including particle emissions from a variety of sources (chapter II), VOCs from a variety of sources (chapter IV), pollutants from combustion plants (chapter VI), and pollutants from incinerators (chapter VII). Although the regulations were first published in 2005, it took six years before they were officially made law.

Prince Edward Island

Prince Edward Island amended their Air Quality Regulations in 2004 to add particulate matter emissions fees for fuel-burning equipment (for details, see schedule D of http://www.gov.pe.ca/law/regulations/pdf/E&09-02.pdf and http://www.gov.pe.ca/photos/original/leg_table_regs.pdf).[14]

Newfoundland

In 2004, Newfoundland amended its Air Pollution Control Regulations, which had been in place since 1996 (NLR 39/04). The original regulations contained air quality standards for $PM_{2.5}$ that were more stringent than the CWS and a one-hour ozone standard. The amend-

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

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

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

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

Manitoba

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

Saskatchewan

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

Alberta

Alberta primarily manages air quality using ambient air quality objectives and guidelines, that are enforced through the provincial permitting and licensing process. Industrial facilities must be designed and operate so as to ensure the provinces ambient air quality objectives are met; however, they are given relative freedom in deciding how to manage their pollution. More stringent permitting regulations were passed in 2003 (Alberta Regulation 276/2003), under the Environmental Protection and Enhancement Act. In 2007, the CWS' $PM_{2.5}$ standard was adopted as an objective (for details, see: http://environment.gov.ab.ca/info/library/5726.pdf). An ozone objective has been in place since 1975, and was reviewed in 2007 but left unchanged. Firms can be fined for violating the conditions of an operation permit, such as failure to comply with air pollutant-related constraints. For example, in 2012 a refinery was fined for failing to install proper air pollution control equipment (for details, see https://www.alberta.ca/release.cfm?xID=32232CC295887-C17E-3ABE-EC7823B5948337D0).

British Columbia

British Columbia manages air quality using a combination of air quality objectives, local airshed management plans, and industrial codes of practice. Air quality objectives are non-binding standards that set the air quality levels to which regulators should aim. Over the phase-in period, the province adopted the $PM_{2.5}$ and O_3 Canada-Wide Standards as air quality objectives (for details, see: http://www.bcairquality.ca/regulatory/air-objectives-standards.html). Towards the end of the phase-in, the province adopted additional, more stringent, $PM_{2.5}$ objectives (see: http://www.bcairquality.ca/regulatory/pm25-objective.html). Provincial regulators achieve these objectives using mandatory codes of practice or other regulations (for details, see paragraph 4.3.2 of http://www.bcairquality.ca/reports/pdfs/pm25-implement-guide.pdf). These provincial regulations can target specific industries, regions, or facilities[15]. In addition, local regulators develop local airshed management plans to meet the air quality objectives. Over the CWS phase-in period, airshed management plans were developed for thirteen regions in the province (see http://www.bcairquality.ca/airsheds/bc-airsheds.html).

Subsidies

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

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