ENVIRONMENTAL REGULATIONS AND THE CLEANUP OF MANUFACTURING: PLANT-LEVEL EVIDENCE

Nouri Najjar and Jevan Cherniwchan*

Abstract—For much of the industrialized world, pollution from manufacturing has been falling despite increased output. We examine how air quality standards—a common environmental regulation—have contributed to this cleanup of manufacturing. We develop a general equilibrium model to show how air quality standards can lead to a cleanup by causing reductions in plant emission intensity, relative changes in plant output, and plant entry and exit. We provide quasi-experimental evidence from Canada to highlight the magnitude of these responses. Our results suggest that air quality standards explain just under 40% of the cleanup of manufacturing.

I. Introduction

MANUFACTURING emission intensity—the amount of pollution emitted from the manufacturing sector per dollar of output shipped—has plummeted in several countries around the world in recent decades. For example, from 1990 to 2008, the emission intensity of U.S. manufacturing fell by up to 77% (Levinson, 2015). Similar trends have been observed in Europe (Brunel, 2017) and in several countries that emit sulphur dioxide (Grether, Mathys, & de Melo, 2009). This implies that for much of the industrialized world, manufacturing is becoming cleaner.

Previous research has concluded that this "cleanup" has been primarily driven by reductions in the emission intensity of individual industries, a phenomenon that has been labeled the "technique effect" (Levinson, 2009, 2015; Brunel, 2017). Understanding the potential causes of this technique effect is important, as pollution has considerable negative effects on human health, as well as productivity, cognition, and labor supply (Graff Zivin & Neidell, 2013).

In this paper, we ask how air quality standards, a common form of environmental regulation, have contributed to the technique effect. We first develop a simple general equilibrium model to show how air quality standards can generate a technique effect by causing changes in plant emission intensity (a "process" effect), output at regulated plants (a "reallocation" effect), and the number of plants operating in an industry (a "selection" effect). As we will show, together, these plant-level responses completely determine the technique effect. Second, we exploit quasi-experimental variation created by a major revision to Canadian air quality regulation, the implementation of the Canada Wide Standards for Particulate Matter and Ozone (CWS), to estimate the effects of this air quality standard on the production choices, emission intensity, and entry and exit decisions of affected manufacturing plants over the period 2004 to 2010. We use these coefficient estimates to construct estimates of the process, reallocation, and selection effects and determine the contribution of air quality standards to the cleanup of Canadian manufacturing.

While there are several possible explanations for the technique effect, our focus on environmental regulation is motivated by the recent work of Shapiro and Walker (2018), who develop and estimate a quantitative model to understand the sources of the cleanup of manufacturing in the United States.[1] The counterfactual exercises considered by Shapiro and Walker provide compelling evidence that the cleanup is primarily due to environmental policy. However, their single, model-derived measure of regulation remains silent on the efficacy of the various environmental regulations that have been used in practice. Determining how different environmental regulations have contributed to the cleanup of manufacturing is important because different policies may yield different process, reallocation, and selection effects due to their design. As a result, two alternative policies that have the same effect on aggregate emission intensity may have different implications for the broader economy.

Our choice to study air quality standards stems from the fact that these regulations have been a predominant form of environmental policy in the jurisdictions where manufacturing cleanups have been documented previously. For example, both the United States (via the Clean Air Act, CAA) and Europe (via the EU Clean Air Directives, CAD) have relied on air quality standards as a means to address air pollution. Given the prominence of these policies, a natural hypothesis is that they have played a large role in the cleanup of manufacturing.

Our motivations for studying Canada are twofold. The first is pragmatic. Estimating the process effect requires information on plant emission intensity, which means we must observe both pollution emissions and output. We obtain this information from a new confidential data set that covers the majority of polluters from the Canadian manufacturing sector.

Our second motivation for studying Canada is that the air-quality regulation that we study, the CWS, is very similar in design to the CAA and shares many features with the air quality standards used in Europe. This commonality suggests our

Received for publication June 5, 2019. Revision accepted for publication December 9, 2019. Editor: Rema N. Hanna.

*Najjar: Western University; Cherniwchan: Carleton University.

We thank Brian Copeland, Carol McAusland, Thomas Lemieux, David Green, Arik Levinson, Joseph Shapiro, Michael Willox, Danny Leung, members of Environment and Climate Change Canada's Legislative and Regulatory Affairs Division, and various seminar and conference participants for helpful discussions and suggestions. Funding from the Economics and Environmental Policy Research Network, the Alberta School of Business Nova Faculty Fellowship, the Centre for Applied Business Research in Energy and the Environment at the University of Alberta, and the Productivity Partnership is gratefully acknowledged. All results have been reviewed by Statistics Canada to ensure no confidential information is disclosed. The opinions expressed in this paper are the ours' alone.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_00904.

[1]Other possible explanations include the effects of international trade (Levinson, 2009; Cherniwchan, 2017) and changes in productivity (Shapiro & Walker, 2018).

The Review of Economics and Statistics, July 2021, 103(3): 476–491

https://doi.org/10.1162/rest_a_00904

estimates should be informative of how air quality standards have contributed to the cleanups of manufacturing that have been observed elsewhere.

In addition to having a similar policy design, it appears that Canada's manufacturing sector has experienced a similar cleanup. Given that it had not been examined previously, we begin our analysis by examining the cleanup of Canadian manufacturing using the industry decomposition developed by Levinson (2009, 2015). As we show in section II, emissions of most air pollutants from manufacturing have fallen substantially, primarily because of reductions in the emission intensity of individual industries. That is, in Canada, as in the United States and Europe, the cleanup of the manufacturing sector is driven by the technique effect.

Next, we develop a simple general equilibrium model in which plants face regulatory constraints similar to those imposed under a typical air quality standard. This model has two key features. First, it allows for plant productivity differences, which have been highlighted as a key determinant of the effects of environmental regulation in the existing theoretical literature (e.g., Anoulies, 2017). Second, it allows for endogenous technology adoption by plants to capture the fact that leading technologies were used as a benchmark in the design of the CWS. Under the CWS, regulated plants were required to either adopt technical changes to meet industry best practices or reduce activities generating the regulated pollutant. Similar technology standards were imposed under the CAA and as part of the CAD.[2] As we show below, the magnitude of the fixed costs associated with adopting technical changes to comply with these standards is key to understanding how air quality standards will generate a technique effect. If these costs are relatively high, only highly productive firms comply with regulation by adopting new technologies; low-productivity firms will instead comply with regulation by reducing output or exiting the market. In this case, the technique effect is primarily driven by reallocation and selection effects.

We then turn to estimate the effects of the CWS on affected manufacturing plants. To do so, we exploit variation in regulatory stringency created by the design of the CWS. The CWS was designed to ensure each region met a minimum level of air quality by establishing thresholds for the ambient concentration of $PM_{2.5}$. Regions in which ambient $PM_{2.5}$ levels exceeded the minimum threshold in a given year were subject to more stringent regulation relative to other regions. In addition, these regulations were 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 a CWS standard were subject to more stringent environmental regulation. 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 confounding factors such as localized recessions or industry demand shocks that would otherwise confound the effects of the CWS.

We find robust evidence that the CWS affected Canadian manufacturing plants.[3] We find that the CWS caused a 26.1% reduction in $PM_{2.5}$ emissions from the average surviving plant. This reduction in $PM_{2.5}$ was driven largely by reduced production: the CWS caused a 14.5% reduction in revenues, on average, but had no significant effect on plant emission intensity. We also find evidence that the CWS affected the extensive margin; it significantly reduced the number of plants from targeted industries in regulated regions.[4]

Finally, we develop an approach for decomposing the coefficient estimates into the selection, reallocation, and process effects. Our approach builds on the logic of the within-between decomposition commonly used in both labor economics and productivity studies. This approach allows us to perform a simple counterfactual exercise without imposing a structural model. In particular, we ask what fraction of the observed technique effect in Canada over our study period is attributable to each of the three channels induced by the CWS.

We find the effects of the CWS explain close to 38% of the reduction in the $PM_{2.5}$ intensity of Canadian manufacturing. Moreover, this was primarily due to selection and reallocation: the selection effect explains approximately 18% of the CWS's contribution to the cleanup, while the reallocation effect explains close to 65%.

Our findings make three contributions to a burgeoning literature examining the sources of the cleanup of the manufacturing sector. First, we show the aggregate trends that have been documented in the United States (Levinson, 2009, 2015) and Europe (Brunel, 2017) extend to Canada. Second, we develop a methodology for using quasi-experimental estimates of the effects of air quality standards to construct estimates of how these regulations have contributed to the cleanup of manufacturing. Our approach is quite general and could be used to examine other potential explanations for the cleanup of manufacturing. Third, we provide causal evidence of how air quality standards have contributed to the cleanup of manufacturing.

[2]The National Ambient Air Quality Standards (NAAQS) used by the U.S. CAA require regulated facilities to adopt state-of-the-art abatement technology and fines those that fail to do so (Greenstone, 2002). EU Directive 2010/75/EU requires member states to issue production permits based on the use of best available techniques and to levy penalties on producers that fail to comply (European Parliament, 2010).

[3]We show that these estimates are robust to a number of potential confounding factors, including preexisting differences in trends across treatment and control plants, a systematic relationship between regional air quality and the production choices of the plants therein, large-emitter effects, and production shifts across plants in multiplant firms. We also show our estimates are not simply capturing changes in emissions reporting behavior.

[4]Our model suggests these findings may be due to the fixed costs associated with complying with the technological standards imposed under the CWS. While the available evidence suggests these costs are high, we do not observe them directly. Hence, to show that our estimates are consistent with this mechanism, we test another prediction of our model, which suggests the effects of regulation will depend on plant productivity if these fixed costs are high. In the online appendix, we report evidence that supports this prediction.

This paper also contributes to a large empirical literature examining the effects of air quality standards. This literature has examined the effects of standards on various margins of plant activity, including entry and exit (e.g., Henderson, 1996; Becker & Henderson, 2000), output (e.g., Greenstone, 2002), and pollution emissions (Gibson, 2019). We contribute to this literature by providing the first unified set of estimates of the effects of air quality standards on plant entry and exit, plant output, and plant emission intensity. While previous studies have examined the first two outcomes independently, estimates of all three are needed to fully characterize the aggregate effects of air quality regulations.

Our paper also contributes to a recent theoretical literature studying environmental regulations in general equilibrium models featuring firm heterogeneity. We do so via our examination of a two-part regulatory structure in which firms must either make technological changes to meet industry best practices or are subject to a regulatory penalty. This is a common form of environmental regulation; it is used in the CWS, the CAA, and the CAD. Yet despite its prevalence, the existing literature has focused on pollution taxes (Andersen, 2018), permit trading (Anoulies, 2017), or intensity standards (Tombe & Winter, 2015).[5]

The remainder of this paper proceeds as follows. In section II, we document the cleanup of Canadian manufacturing. Section III provides an overview of the CWS. Section IV presents our model. Section V presents our data, research design, and empirical specification and our estimates of the CWS on individual manufacturing plants. Section VI presents the aggregate implications of our plant-level estimates. Section VII concludes.

II. The Cleanup of Canadian Manufacturing

Our goal in this paper is to determine how the responses of individual plants to air quality standards have contributed to the cleanup of manufacturing. While the cleanup has been documented in several countries, including the United States (Levinson, 2009, 2015) and the European Union (Brunel, 2017), it has not yet been 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.

For this exercise, we examine the pollution intensity of manufacturing for four common air pollutants: particulate matter ($PM_{2.5}$), nitrogen oxide (NO_X), volatile organic compounds (VOCs), and carbon monoxide (CO). To do so, we rely on industry-level pollution data from Environment and Climate Change Canada's Air Pollutant Emission Inventory, and industry-level output data, constructed using data from Statistics Canada.[6]

Changes in the emission intensity of Canadian manufacturing, relative to 1992 levels, are illustrated in figure 1. As the figure shows, the emission intensity of Canada's manufacturing sector has fallen by an average of 3.5% to 4.7% per year since 1992. This implies that the magnitude of Canada's cleanup was similar to those that occurred in the United States and Europe. Manufacturing emission intensity fell by 3.6% to 4.3% annually from 1990 to 2008 in the United States (Levinson, 2015) and by 3.4% to 5.5% per year from 1995 to 2008 in Europe (Brunel, 2017).

While this evidence suggests that they are alike, it reveals little as to whether these trends were driven by similar sources. As such, we follow Levinson (2009) and adopt the Grossman and Krueger (1993) decomposition to study the potential sources of the cleanup. 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 sectors to clean sectors with relatively low emission intensities or by a technique effect created by reductions in the emission intensity of individual industries.

To make this decomposition explicit, let Z, X, and $E = Z/X$ denote the pollution emissions, output, and emission intensity of the manufacturing sector, respectively. Let Z_i, X_i, and E_i denote the same for individual industries,[7] 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

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

The first term in equation (1) is the composition effect, and the second term is the technique effect.

We follow Levinson's (2015) approach and take equation (1) directly to the data.[8] 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 to 2015. We report these estimates in the online appendix.

The results of this exercise suggest that the cleanup of the Canadian manufacturing sector is primarily due to the

[5] One notable exception is Ryan (2012), who examines the dynamic consequences of the CAA on the Portland cement industry using a structural partial equilibrium model featuring Cournot competition. Our work is also related to that of Cherniwchan et al. (2017), who develop a similar model to the one developed here. However, they use their model to examine the effect of trade liberalization on emissions in the presence of a uniform pollution tax. In contrast, we focus on the effect of air quality standards on industry emission intensity.

[6] We constructed our output measures 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.

[7] Due to data limitations, we define industries as either three- or four-digit NAICS codes.

[8] We follow Levinson (2015), and calculate the technique effect as the percentage change in manufacturing pollution intensity implied by a Laspeyre's-type index, computed relative to 1992. Adopting the same definitions as above and letting t denote time, the index is given by $\left[\sum_i E_{it} \times X_{i,1992}\right] / \left[\sum_i E_{i,1992} \times X_{i,1992}\right]$. The composition effect is calculated as the residual.

FIGURE 1.—CANADIAN MANUFACTURING POLLUTION INTENSITY: 1992–2015

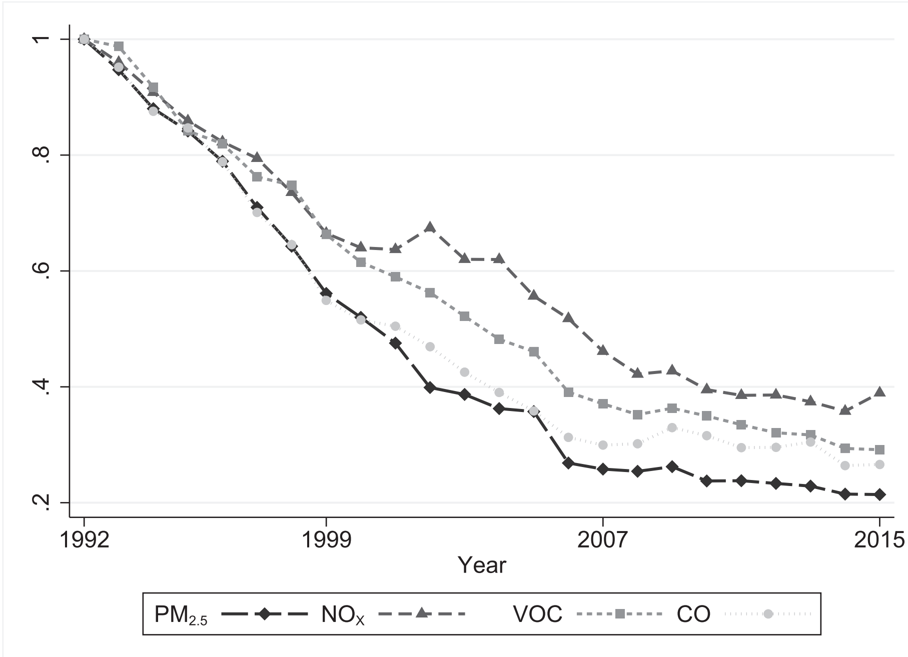

Figure depicts trends in manufacturing pollution per dollar of output (measured as real shipments) for $PM_{2.5}$, NO_X, VOCs, and CO relative to 1992.

technique effect. Nearly 99% of the reduction in $PM_{2.5}$ intensity, 93% of the reduction in NO_X intensity, 96% of the reduction in VOC intensity, and 100% of the reduction in CO intensity from Canadian manufacturing between 1992 and 2015 is due to the technique effect. This is further evidence that the Canadian cleanup is similar to those observed elsewhere; as shown by Levinson (2009, 2015) and Brunel (2017), the cleanups of U.S. and European manufacturing are also primarily due to the technique effect.

A. *How Do Industries Clean Up?*

The evidence presented indicates that the technique effect—changes in industry emission intensity—is the predominant source of the cleanup of manufacturing in Canada. The evidence suggests that environmental regulations may be an important contributor to this technique effect, as they cause some plants to exit and surviving plants to shrink and adopt new production processes (Henderson, 1996; Greenstone, 2002).

To see why these changes may reduce an industry's emission intensity, it is useful to further decompose the technique effect given in equation (1) into plant-level responses, as in Cherniwchan et al. (2017). Suppose, as above, that the pollution intensity of industry i is given by $E_i = Z_i/X_i$. In addition, suppose each industry is composed of a continuum of plants, and let $x_i(n)$, $z_i(n)$, and $e_i(n)$ denote output, pollution, and pollution intensity from plant n. Finally, 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.[9] In this case, the emission intensity of industry i can be expressed as a weighted average of plant emission intensities: $E_i = \int_0^{n_i} e_i(n)\lambda_i(n)dn$. Totally differentiating yields the following expression for the change in emission intensity of any industry i in response to a shock:

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

Equation (2) shows that the technique effect produced by a shock, such as a change in regulation, can be decomposed into three channels driven by plant-level changes. We call the first of these channels—the first term on the right-hand side of equation (2)—the process effect. This captures the change in industry emission intensity due to changes in plant emission intensity resulting from the adoption of new production processes.[10] The second term on the right-hand side of equation (2) captures the effects of the shock on the relative size of plants within an industry. This reallocation effect would arise if the shock does not affect plants uniformly. If the shock only affects a subset of plants in an industry, as is common with

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

[10] As discussed by Cherniwchan et al. (2017), what we term the process effect could be generated by either changes in the emission intensity or mix of individual activities done at each plant. For example, the set of activities undertaken could change as plants change their product mix in response to a shock (Barrows & Ollivier, 2018). Here, we abstract from such changes as we do not observe these activities directly.

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 industry emission intensity created by a change in the set of plants operating within the industry owing to plant entry and exit.

In what follows, we present direct estimates of the regulatory process, reallocation, and selection effects in the cleanup of Canadian manufacturing induced by the introduction of an air quality standard. These estimates allow us to quantify the degree to which this policy contributed to the technique effect and the cleanup of Canadian manufacturing.

III. Air Quality Regulation in Canada

In order to understand how air quality standards have contributed to the cleanup 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 to 2012.[11] Moreover, its design makes it an attractive setting for studying the effects of air quality regulation.

First signed in 2000, the CWS was an agreement between the federal government of Canada and the various provincial environment ministries.[12] 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, CMAs).[13] Much like the NAAQS in the United States, these standards created a common target level of air quality to be achieved by each CMA in Canada.[14] To that end, plants in CMAs with ambient $PM_{2.5}$ or O_3 concentrations in excess of the relevant standard were subject to more stringent environmental regulation than other plants. In what follows, we examine the effects the CWS on plants that emit $PM_{2.5}$.[15]

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.[16] These industries were chosen because they were viewed as major contributors to the air quality problems that motivated the CWS and were common across all CMAs.

Broadly, the CWS was a tiered regulatory regime in which the federal and provincial governments agreed on local air quality targets (the CWS standards), the federal government developed best practice and guidance documents for targeted industries to provide management tools to the provincial governments (Government of Canada, 2003), and the provincial governments regulated plants in targeted industries to meet these regional standards.

Provinces used their annual operation permit systems to regulate plants. This required plants to prove compliance with certain environmental regulations in order to operate in any year (see, e.g., Environment Canada & FPAC, 2004), and meant they could effectively follow one of two paths: either adopt technical changes recommended to their industry or reduce polluting activities. When local air quality was relatively clean (i.e., in compliance with the CWS), permitting constraints were laxer than when air quality was poor. Consequently, regulatory stringency varied over time according to regional air quality.

It is important to note that the costs associated with adopting production processes that emit less $PM_{2.5}$ appear to be substantial. Industrial $PM_{2.5}$ emissions are caused by a number of processes, several of which could potentially occur at the same facility. These include, for example, the combustion of fossil fuels, chemical reactions, and wear and tear on machinery. There are both low- and high-cost approaches to reducing $PM_{2.5}$ emissions. The low-cost approaches include fuel switching and the use of inertial separators or wet scrubbers. These methods, however, can have limited applicability in industrial uses and are not particularly effective for $PM_{2.5}$ (World Bank Group, 1998). Instead, for the typical manufacturing 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 (Environment Canada, 2002).

IV. Theoretical Framework

The CWS required plants to either adopt technical changes to meet industry best practices or reduce polluting activities. Mandating technical changes, and penalizing those that fail

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

[13] These included census agglomerations (CA) and CMAs. CMAs have populations of at least 100,000, while CAs have at least 10,000 people. We use the term CMA for both.

[14] The standard for $PM_{2.5}$ required each CMA's 24-hour $PM_{2.5}$ concentration lie below 30 $\mu g/m^3$. Achievement of the $PM_{2.5}$ standard was based on the 98th percentile of each region's annual 24-hour ambient concentration. The O_3 standard was applied as an 8-hour standard that required each CMA's O_3 concentration to lie below 65 parts per billion (ppb). Achievement of the O_3 standard was based on the 4th highest 8-hour concentration reported in a given year. In comparison, the NAAQS in the United States currently contain a 24-hour $PM_{2.5}$ standard set at 35 $\mu g/m^3$, and an 8-hour O_3 standard set at 70 ppb.

[15] Given that the source of the pollution data we use in our main analysis does not include information on ozone emissions, we also examined the effects of the CWS on plants that emit NO_X, a key ozone precursor. We have refrained from reporting these results, as we are unable to ensure they are capturing the effects of the CWS as opposed to other factors. We also account for the effects of the O_3 standard throughout our empirical analysis.

[16] These were pulp and paper, lumber and wood products, power generation, iron and steel, base metal smelting, and the concrete and asphalt industries (CCME, 2000).

to do so is a common regulatory approach. For example, under the NAAQS used as part of the U.S. CAA regulated facilities must adopt state-of-the-art abatement technology or face fines (Greenstone, 2002). The EU Industrial Emissions Directive also requires the imposition of industrial emissions regulations that mandate the adoption of Best Available Techniques and penalizes facilities that fail to do so (European Parliament, 2010).

It is not immediately clear how this type of policy would affect manufacturing facilities. Hence, before we turn to estimate the effects of the CWS on Canadian manufacturing plants, we develop a simple general equilibrium model based on the work of Melitz (2003) and Bustos (2011) that features this type of regulation. We use this model to derive several predictions as to the effects of air quality standards. Below, we outline the features of the model and highlight its key empirical predictions. For brevity, we relegate details of the model's solution and relevant derivations to the online appendix.

A. *Setup*

We consider a closed economy comprising 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 by-product, and this harms consumers, lowering their utility. For convenience, we let labor be the numeraire.

The demand side of the economy is represented by a consumer who derives utility from the consumption of goods and disutility from aggregate pollution Z 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 that 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 must pay a fixed cost f_ϵ. Entrants then draw a productivity φ from a common Pareto distribution $G(\varphi) = 1 - \varphi^{-k}$.

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

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

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

The technologies described embed two important assumptions: for a given technology, more productive firms are less pollution intensive, and there are leading technologies that, if adopted, lower a firm's pollution intensity, even given the same level of productivity. We adopt the first assumption because in our context more productive $PM_{2.5}$ emitters are also less pollution intensive, a fact we document in the online appendix. We adopt the second assumption because regulations of this nature typically require firms to meet industry best practices, meaning some firms use cleaner technologies even absent regulation.

B. *The No-Regulation Equilibrium*

Our interest is in understanding the effects of imposing an air quality standard. Hence, we begin by considering the no-regulation equilibrium (*no*) in which pollution is not regulated. In this case, labor is the only variable cost of production. Below, we describe firm behavior in this case; the industry equilibrium is presented in the online appendix.

Given consumer preferences, firms set prices at a constant markup over marginal costs.[17] Hence, absent regulation, firms that employ technology b or r charge the same price: $p_b^{no}(\varphi) = p_r^{no}(\varphi) = 1/[\rho\varphi]$. If, instead, a firm employs technology s, it charges $p_s^{no}(\varphi) = 1/[\rho\alpha\varphi]$.

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

The exit and technology choices made by firms are highlighted in panel a of figure 2, which depicts the profits associated with adopting each technology as a function of firm productivity. 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 $\varphi < \varphi_\epsilon^{no}$, it exits the market. If firms do not exit, they choose the technology that yields the highest profit. This means that a firm with productivity level

[17]It is worth noting that in principle, regulation could also affect firm markups, for example, by changing market concentration. Here, we abstract from this channel for expositional simplicity. However, this assumption could be relaxed by adopting an alternative demand system, such as the linear demand system in Melitz and Ottaviano (2008).

FIGURE 2.—TECHNOLOGY CHOICES WITH AND WITHOUT ENVIRONMENTAL REGULATION

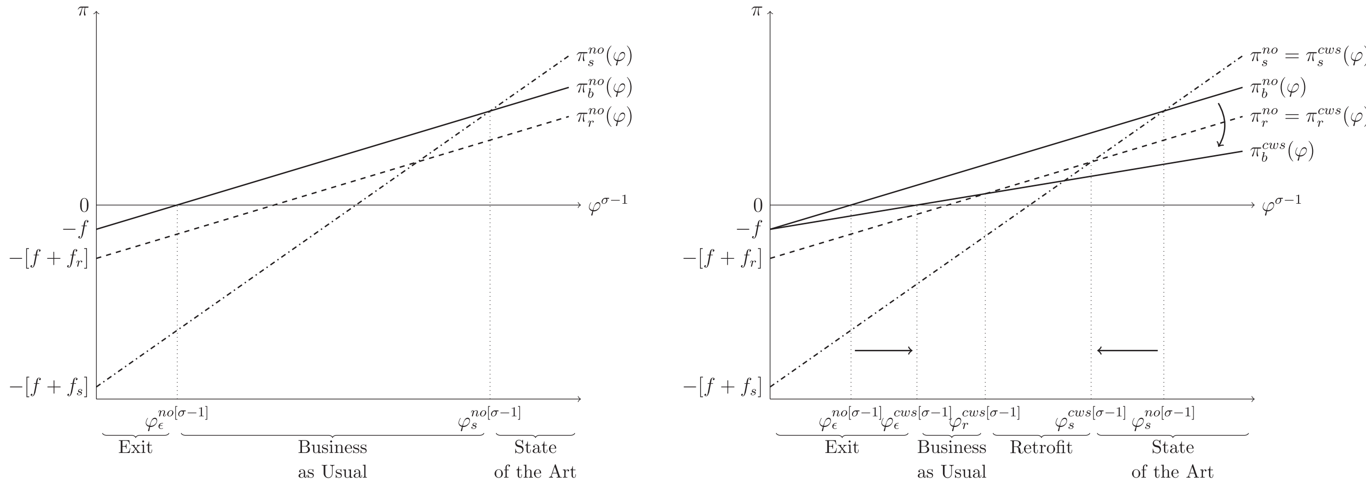

(a) Without Regulation

(b) With Regulation

$\varphi \in \{\varphi_\epsilon^{no}, \varphi_s^{no}\}$ will produce using technology b. Similarly, if a firm has productivity level $\varphi > \varphi_s^{no}$, then the reduction in variable cost created by adopting s is great enough to justify the fixed cost of adoption, meaning that it will adopt the technology s.

As panel a of figure 2 shows, firms never choose to retrofit absent regulation. Adopting technology r reduces the emission intensity of production but has no effect on variable production costs without regulation. As a result, retrofitting simply lowers firm profits below what can be obtained using technology b by increasing the average costs of production.

C. *The Effects of Air Quality Regulation*

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

In this regime (cws), the government regulates pollution using a two-part regulatory rule. If a firm uses a clean production process (either technology r or s), it is not subject to regulation because it is operating with the lowest emission intensity currently achievable for its productivity level. As a result, the marginal costs and profits from using these technologies are unaffected by regulation. In contrast, a firm that employs a dirty production process (technology b) is subject to a regulatory constraint in the form of a charge τ on each unit of pollution emitted.[18] Given that prices feature constant markups, this increase in marginal costs raises the price of output for firms producing with technology b. That is, $p_b^{no}(\varphi) < p_b^{cws}(\varphi) = [1+\kappa\tau]/[\rho\varphi]$, so profits are now $\pi_b^{cws}(\varphi) = [1/\sigma]I\,[P\rho]^{\sigma-1}\,\varphi^{\sigma-1}[\frac{1}{1+\kappa\tau}]^{\sigma-1} - f$. This means holding industry prices (P) fixed, the profit from using b falls for any φ.

This partial equilibrium outcome is depicted in panel b of figure 2, which displays the technological choices made by firms when faced with an air quality standard, holding industry prices fixed. As the figure shows, a reduction in variable profits under technology b increases the exit cutoff from φ_ϵ^{no} to φ_ϵ^{cws}. As such, regulation creates a selection effect: firms with $\varphi \in \{\varphi_\epsilon^{no}, \varphi_\epsilon^{cws}\}$ exit in response to regulation. Moreover, the increase in technology b's variable production cost makes technology upgrading a profitable alternative for some firms. As depicted, it is profit maximizing for firms with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^{no}\}$ to adopt technology r in response to regulation. For these firms, the benefit of avoided emission charges outweighs the increase in fixed costs. Similarly, firms with productivity $\varphi \in \{\varphi_c^{cws}, \varphi_s^{no}\}$ adopt technology s because it is now profit maximizing to do so.

While figure 2 clearly highlights how air quality standards create selection effects by causing firms to exit in response to regulation, the reallocation and process effects are not readily apparent from the figure. These effects are instead displayed in figure 3 for firms that survive regulation (those with $\varphi > \varphi_\epsilon^{cws}$). This figure depicts the effects of environmental regulation on firm revenues (panel a) and emission intensity (panel b) holding industry prices fixed. Both panels show that the most productive firms, with productivity $\varphi > \varphi_s^{no}$, are unaffected by regulation because they use the most advanced technology regardless of regulatory regime. In contrast, regulation causes the least productive firms, with productivity $\varphi < \varphi_r^{cws}$, to produce less; they use technology b under either regime, and regulation causes the variable cost of production to increase. However, as shown in panel a, this production is reallocated to firms with productivity $\varphi \in \{\varphi_s^{cws}, \varphi_r^{cws}\}$

[18]We use a pollution charge for analytical tractability. This charge could reflect actual penalties, as is the case with the U.S. CAA, or it could reflect the costs of dealing with increased regulatory oversight, such as measuring and reporting emissions. Alternatively, we could impose a production cap, which would be more in line with the CWS, without substantively affecting the results. Finally, to ensure consistency with the design of the CWS, we assume that the revenue from the charge is spent outside the model.

FIGURE 3.—SURVIVING FIRMS WITH AIR QUALITY REGULATION

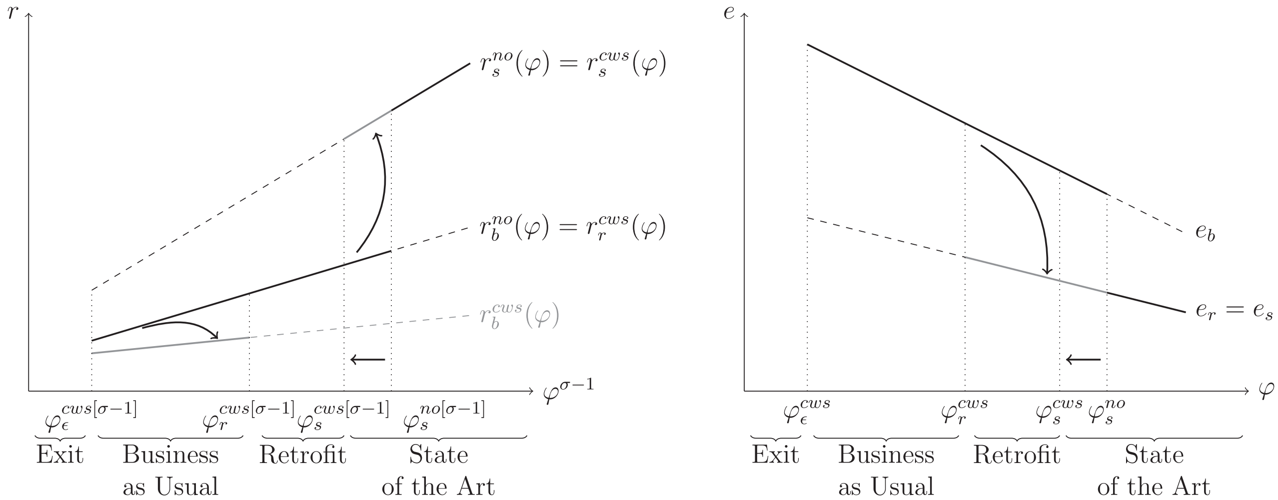

(a) Firm Revenues (b) Firm Pollution Intensity

because variable costs for these firms fall. Finally, regulation induces a process effect by causing the emission intensity of firms in the middle of the productivity distribution, those with productivity $\varphi \in \{\varphi_r^{cws}, \varphi_s^{no}\}$, to fall. This occurs because these firms either retrofit from technology *b* or adopt technology *s*.

While the discussion above is illustrative, in the online appendix we show that if retrofitting is costly, that is, if $f_r > 0$, then in equilibrium: (a) regulation causes some firms to exit, (b) on average, revenue and emission intensity fall in response to regulation, and (c) regulation's effects vary across the productivity distribution. Specifically, revenues fall for the least productive surviving firms using technology *b*, while emission intensity falls for firms in the middle of the productivity distribution that adopt *r* or *s*.

In addition, we show that the fixed cost of retrofitting (f_r) plays an important role in determining the channels through which regulation causes an industry to clean up. Specifically, decreasing f_r increases the measure of firms that adopt a clean production process in response to regulation and reduces the measure of firms that exit in response to regulation.[19] Thus, when f_r is large, regulation should primarily cause an industry to clean up through reallocation and selection effects. In what follows, we test these predictions empirically.

V. The Effects of the CWS on Manufacturing Plants

Our theoretical model provides a number of clear predictions as to how manufacturing plants should respond to a policy such as the CWS. Taken together, these results imply that commonly used air quality regulations should primarily reduce industry emission intensity via reallocation and selection effects if the fixed costs of process changes are large, as we argue is the case for $PM_{2.5}$. In this section, we explore these predictions empirically by estimating the effects of the CWS on plant emissions, emission intensity, production, and exit for a sample of plants that emit $PM_{2.5}$.

A. *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 regulation across time, industries, and regions created by the design of the CWS.

We begin by comparing the average outcomes from plants in regulated CMAs while regulated to their average outcomes while unregulated. This allows us to control for any unobserved time-invariant industry, CMA, or plant characteristics that would affect plant outcomes. Moreover, in the absence of any other shocks, this comparison would identify the causal effect of the CWS. Yet such absence is unlikely; there is strong reason to believe that this comparison may also capture the effects of regional, industry, or aggregate shocks.

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

Similarly, to control for industry shocks, we exploit cross-CMA variation in regulation and utilize the fact that in any particular industry, only plants in areas with poor air quality were subject to stringent environmental policy. This allows us to use the average outcomes from plants in a targeted

[19]This result requires restricting the size of f_s.

industry in an unregulated CMA as a counterfactual for the average outcomes of plants from that industry that are located in a regulated CMA. This comparison captures the effects of industry-specific shocks, such as foreign demand shocks or revisions to federal policies that target certain sectors, that would otherwise confound identification.

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

Given our interest in determining the CWS's contribution to the cleanup of manufacturing in Canada, we employ this triple-difference research design in two baseline regressions. First, we estimate the effects of regulation on continuing manufacturing plants:

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

where y_{pict} is the natural log of the dependent variable of interest (e.g., pollution, sales) at plant p, in industry i, located in CMA c, at time t. T_{ict}^{j} is a treatment indicator that takes the value of 1 for plants that are in industries targeted by the CWS for years in which their CMA exceeded standard j.[20] Equation (3) also includes plant (ρ_p), CMA-year (ξ_{ct}), industry-year[21] (λ_{it}) fixed effects, and an error term (ε_{pict}). The plant fixed effects account for any time-invariant unobserved plant, industry, or CMA heterogeneity. The CMA-year and industry-year fixed effects capture region-specific and industry-specific shocks, respectively. Finally, the error term captures idiosyncratic changes in outcomes across plants.

Second, we estimate the effects of the CWS on plant net exit:

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

In equation (4), N_{ict} is the number of active polluters[22] in industry i in CMA c and year t, T_{ict}^{j} is the treatment indicator for standard j (which takes a value of 1 for industries targeted by the CWS for years in which their CMA exceeds threshold j), $I(\mathrm{CWS})_{ic}$ 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. In equation (3), we cluster standard errors by industry-CMA, and in equation (4) we do so by CMA (Bertrand, Duflo & Mullainathan, 2004).

The coefficient of interest in both equations (3) and (4) is β_{PM}. However, it is important to note that these coefficients capture effects for different populations. In equation (3), β_{PM} measures the average percentage change in outcomes for continuing plants affected by the particulate matter standards relative to those that are not and is identified from within plant comparisons over time. In contrast, β_{PM} from equation (4) measures the net entry or exit of plants in an industry-CMA due to the CWS and are identified from within industry-CMA comparisons over time. This distinction is important because of how we use these coefficient estimates to determine the contribution of the CWS to the cleanup of Canadian manufacturing. As we discuss further in section VI, we use our estimate of β_{PM} from equation (3) to construct estimates of the reallocation and process effects, which are defined for the set of continuing plants, while we use the estimate of β_{PM} from equation (4) to construct an estimate of the selection effect, which is defined over the entire population.

For equations (3) and (4) to credibly identify the effects of the CWS, there must be no other factors driving differences in outcomes across treated and untreated plants. We have reason to believe this is the case, as variation in regional air quality determines assignment to treatment. As with the CAA in the United States, regulations are determined by a nationally set air quality threshold, meaning that they are unrelated to differences in local tastes, characteristics, or economic conditions (Greenstone, 2002). Moreover, $PM_{2.5}$ and O_3 are capable of being transported long distances by prevailing wind patterns, meaning that ambient pollution levels need not solely reflect local economic activity. Regardless, in what follows, we estimate two alternate specifications to test the credibility of our key identifying assumption.

B. *Data and Measurement*

Our analysis relies on a unique confidential data set that contains information on the pollution emissions and productive activities of Canadian manufacturing plants. This data set was created by merging data from two existing sources: the National Pollutant Release Inventory (NPRI) and the Annual Survey of Manufactures (ASM). The NPRI contains plant-level information on the emissions of various pollutants.[23] The ASM provides plant-level information on output,

[20] We have adopted this definition given the design of the CWS; plants are regulated on the basis of the air quality in their CMA in the current year. Note that unlike the NAAQS in the United States, the regulatory status of a CMA does not depend on the status of neighboring CMAs and did not require multiple periods of improvement to change status.

[21] The CWS defined targeted industries at the three- or four-digit NAICS level. Our industry definitions account for this. All targeted industries defined at the four-digit level are grouped at the four-digit level. The remaining industries are grouped at the three-digit level.

[22] We define an exit as a plant that does not report to the NPRI in a given year.

[23] A concern with the NPRI data is the fact that they are self reported, which means that one of our dependent variables is potentially measured with error. This would be problematic if the error varied systematically with ambient air quality. We do not believe this is the case, as regulatory status is unaffected by reported emissions. Furthermore, there are significant incentives for truthful reporting; facilities face fines of up to $12 million for

TABLE 1.—SUMMARY STATISTICS

	(1) Emissions (tonnes)	(2) Sales (\$1 mill.)	(3) Value Added (\$1 mill.)	(4) Employment	(5) VA/Worker (\$1,000)	(6) *N*
$PM_{2.5}$	20.13	259.87	79.25	317.07	211.39	3978
	(98.47)	(1259.37)	(345.96)	(838.82)	(271.67)	
Full ASM		11.12	4.29	35.69	84.78	309,541
		(123.56)	(34.34)	(125.27)	(166.11)	

Table reports means and standard deviations of variables examined in the analysis. The first row reports statistics for our estimation sample, while the second reports statistics for the entire ASM. All monetary values are reported in 2007 Canadian dollars.

production costs, employment, and other plant characteristics. Our matched data set contains longitudinal, plant-level information on $PM_{2.5}$ emissions, output, and other plant characteristics over the period 2004 to 2010. Further details on the data set and its construction are given in the online appendix.

Descriptive statistics for the key variables that we employ are reported in table 1. The first row reports summary statistics for the set of plants that emit $PM_{2.5}$, while the second row corresponds to all plants in the ASM.[24] Our estimation sample, row 1, is an unbalanced panel containing 3,978 plant-year observations.

The summary statistics reported in table 1 show that our sample of $PM_{2.5}$ polluters represents the largest plants in the manufacturing sector. These plants sell more goods, employ more workers, and have higher value added per worker than the average manufacturing plant. This is in part due to the reporting requirements of the NPRI; by law, plants report only if their emissions exceed a minimum threshold and have at least ten employees or operate an on-site generator (Environment & Climate Change Canada, 2016). Despite this, our analysis covers plants that account for the majority of manufacturing pollution in Canada.

CWS regulatory status. As the standards were enforced by the provincial governments, there is no exhaustive list that tracks compliance with the CWS. We construct this information using data on local air quality and the compliance guidelines contained in the CWS.[25] The air quality data come from Canada's National Air Pollution Surveillance Program, which provides data on hourly monitor-level $PM_{2.5}$ and O_3 concentrations. We use these data to construct CMA-level pollution concentration measures for each year in our sample, where the measures computed are those associated with each standard. We label a CMA as in violation of the CWS if its constructed pollution measure exceeds the relevant threshold.[26]

misreporting. Nonetheless, in the appendix, we provide evidence that our estimates are not simply capturing changes in reporting behavior.

[24]The statistics in row 1 are weighted to account for potential sample bias induced by matching plants across data sets (see the online appendix for further details).

[25]We discussed this approach with members of ECCC's Legislative and Regulatory Affairs division to ensure we adopted the correct methodology and data for this exercise.

[26]For details of this procedure, see the online appendix.

The variation in regulatory status created by changes in ambient air quality is illustrated in figure 4, which shows the CMAs that changed regulatory status for the $PM_{2.5}$ and O_3 standards. In figure 4, the black CMAs changed status under both the $PM_{2.5}$ and O_3 standards, the dark gray CMAs changed status only for the $PM_{2.5}$ standard, the light gray CMAs changed status only for the O_3 standard, and the white CMAs did not change status under either standard. As the figure shows, there was substantial variation in which CMAs changed their regulatory status over the 2000–2010 period. Of the 149 CMAs in our sample, 23% changed status only under the $PM_{2.5}$ standard, 26% changed status only under the O_3 standard, 11% changed status under both standards, and 60% never changed status.

C. Results

Flexible estimates. The key identifying assumption underlying equations (3) and (4) is that there are no other factors, aside from the CWS, causing differential trends in outcomes across plants. As we noted above, there is strong reason to believe that this is the case given the design of policy, but it is still possible that our main estimates will capture differences in trends across plants owing to other factors beyond the CWS. Hence, before we turn to our main empirical analysis, we conduct two exercises to ensure that this is not the case. For each exercise, we examine plant emissions and output and the number of emitters in a CMA-industry cell.

For our first exercise, we estimate versions of equations (3) and (4) that allow the effects of the CWS $PM_{2.5}$ standard to vary by year. Specifically, for our main plant-level outcomes, we implement this event-study design via the following regression based on equation (3):

$$Y_{pict} = \sum_{k=-3} \beta^k_{PM} T^{PM}_{ick} + \beta_{O3} T^{O3}_{ict} + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict}, \tag{5}$$

where T^{PM}_{ick} is an indicator for the years before ($k < 0$) or after ($k \geq 0$) a plant is treated 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.

FIGURE 4.—REGULATORY STATUS CHANGES UNDER THE CWS

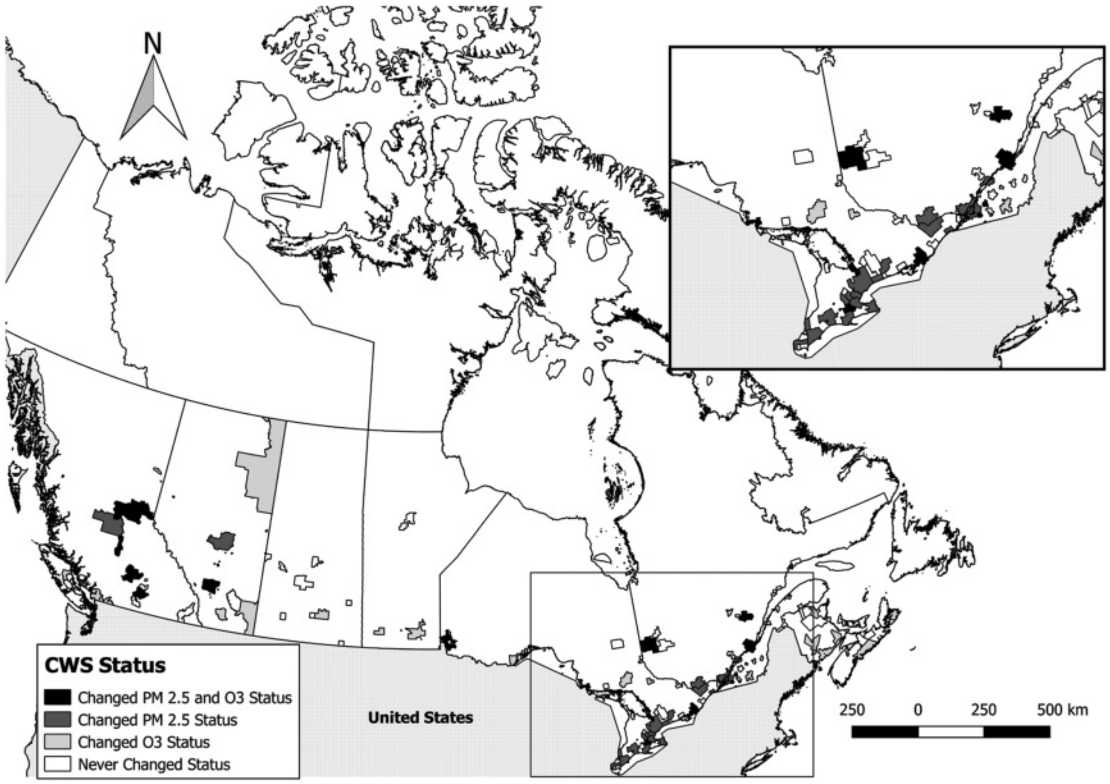

Figure depicts $PM_{2.5}$ and O_3 standard status changes for each CMA from 2000 to 2010. Part of the northern Canadian Territories are trimmed for scale. The inset shows detail on the most densely populated area of Canada and is outlined on the main map.

We also estimate an analogous event-study regression based on equation (4):

$$N_{ict} = \sum_{k=-3} \beta_{PM}^{k} T_{ick}^{PM} + \beta_{O3} T_{ict}^{O3} + \alpha I(\text{CWS})_{ic} + \xi_{ct} + \lambda_{it} + \epsilon_{ict}, \tag{6}$$

where all coefficients are as defined previously in equations (4) and (5).

We estimate both equations (5) and (6) from three periods before treatment onward. Separate coefficients are estimated up to three periods after treatment, and all periods greater than three years after treatment are pooled. We drop all observations from three periods prior to treatment. We weight the regressions based on equation (5) to correct for potential sample bias introduced by the matching procedure used to link the NPRI and the ASM.[27] Standard errors are clustered by CMA-industry in equation (5) and CMA in equation (6).

The results of equation (5) are shown in panel a of figure 5. This panel shows the results of three separate regressions, each with a different dependent variable: the natural logarithm of $PM_{2.5}$ emissions (circles), plant sales (squares), and plant value added (triangles). Panel b of figure 5 displays the net-exit results (diamonds) from equation (6).[28]

Panel a of figure 5 shows there were no significant differences in outcomes for our treatment and control plants prior to regulation. In addition, there was a clear break starting in the year of regulation and persisting following treatment, although the effect on output may dissipate by three years after regulation. Similarly, panel b of figure 5 shows no significant difference in manufacturing net exit between treatment and control CMA industries prior to regulation but a clear break following regulation. This suggests our main results will not simply capture preexisting differences in trends across plants or CMA industries.

For our second exercise, to ensure that our estimates are capturing the effects of the CWS, we estimate versions of equations (3) and (4) in which we allow the effect of the CWS's $PM_{2.5}$ standard to vary on the basis of a CMA's air quality. Given the design of the CWS, there should be no significant difference between targeted and nontargeted

[27]The potential bias arises because match probability is positively correlated with plant size. If the effects of the CWS vary by plant size, then the matched data would produce biased estimates. Details on the weighting procedure can be found in the online appendix.

[28]For brevity, tables containing the point estimates and associated standard errors from the regressions presented in figure 5 are reported in the online appendix.

FIGURE 5.—THE EFFECT OF $PM_{2.5}$ REGULATION BY YEARS PRE-/POST-REGULATION

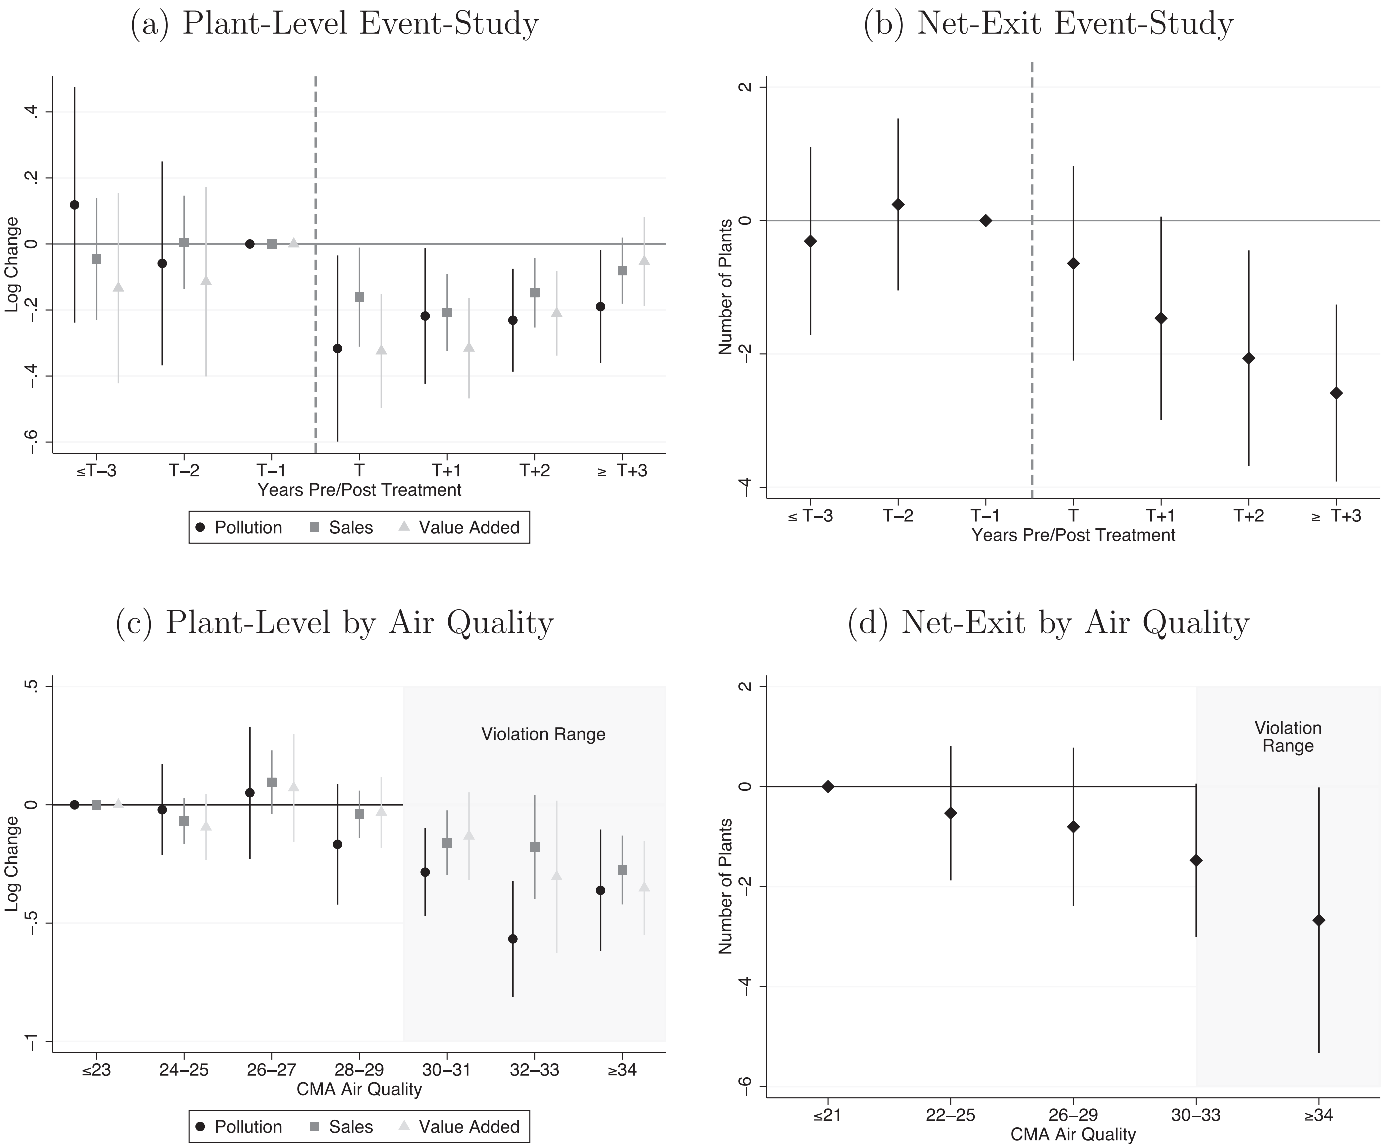

Figure shows estimates of an event study estimation of the effects of the $PM_{2.5}$ standard (panels a and b) and from a DDD estimation allowing the effects of regulation to vary by CMA air quality (panels c and d). Panels a and c show the effects on plants. Each plot results from three different regressions, each with a different dependent variable: the natural logarithm of $PM_{2.5}$ emissions, sales, and value added. Panels b and d show the effects on the number of plants in a CMA-industry. In all panels, estimated coefficients are displayed with a 95% confidence interval. All regressions in panels a and c 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, with standard errors clustered by industry-CMA. The regressions in panels b and d include industry-year and CMA-year fixed effects and an indicator for whether the industry-CMA was ever regulated, with standard errors clustered by CMA.

industries when their CMA's air quality is below the standard's threshold.

To accomplish this, we assign each observation to 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 nontargeted 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.

For our core plant-level outcomes, we implement this exercise using the following regression based on equation (3)

$$Y_{pict} = \sum_{b} \beta^{b}_{PM}[K_i \times I(\underline{A^{PM}_b} \leq a^{PM}_{ct} < \overline{A^{PM}_b})] + \beta_{O3} T^{O3}_{ict} + \rho_p + \xi_{ct} + \lambda_{it} + \epsilon_{pict}, \tag{7}$$

where b indexes air quality bin numbers, K_i selects all industries targeted by the CWS, a^{PM}_{ct} is the $PM_{2.5}$ air quality measured in CMA c in year t, $\underline{A^{PM}_b}$ and $\overline{A^{PM}_b}$ are the air quality lower and upper bounds for bin b, respectively, and $I(\underline{A^{PM}_b} \leq a^{PM}_{c} < \overline{A^{PM}_b})$ is an indicator for all CMA-years with air quality that corresponds to bin b. The coefficient β_j^b gives the effects of the $PM_{2.5}$ standard in air quality bin b. In estimating equation (7), we omit the cleanest $PM_{2.5}$ air quality bin and break the air quality distribution into seven equal-sized bins from 23 to 34 $\mu g/m^3$. We also weight each regression to account for potential sample bias created by matching the NPRI and ASM, and cluster standard errors by CMA-industry.

We also estimate an analogous specification for net exit based on equation (4):

$$N_{ict} = \sum_b \beta_{PM}^b [K_i \times I(\underline{A_b^{PM}} \leq a_{ct}^{PM} < \overline{A_b^{PM}})] + \beta_{O3} T_{ict}^{O3} + \alpha I(\text{CWS})_{ic} + \xi_{ct} + \lambda_{it} + \epsilon_{ict}, \quad (8)$$

where N_{ict} is the number of plants operating in a CMA-industry-year and all other variables are defined as in equations (4) and (7). When estimating equation (8), we break the air quality distribution into four equal-sized bins from 22 to 34 $\mu g/m^3$ and omit the "cleanest" bin. We split the air quality distribution into fewer bins than in equation (7) to ensure that we have a large enough sample in each. Standard errors are clustered by CMA.[29]

The results of estimating equations (7) and (8) are displayed in panels c and d of figure 5. Panel c displays the coefficient estimates and confidence intervals for three separate regressions, each corresponding to a different dependent variable: the natural log of $PM_{2.5}$ emissions (circles), the natural log of plant sales (squares), and the natural log of plant value added (triangles). Panel d displays the results of equation (8), which shows the effect of the $PM_{2.5}$ standard on net exit (diamonds). The dependent variable in this regression is the number of plants operating in a CMA-industry cell.

The results presented in panel c of figure 5 show no significant difference between plants in targeted and nontargeted industries in regions with air quality below the standard's threshold (a $PM_{2.5}$ concentration of 30 $\mu g/m^3$). A break, significant at the 10% level, occurs at the threshold level for pollution, sales, and value added. While not as precisely estimated, the results in panel d tell a similar story. There is no significant difference between targeted and nontargeted industries in regions with air quality below the standard's threshold, but a break, significant at the 10% level, occurs at the threshold level. These results provide further evidence in support of our identification strategy.

Main estimates. We now turn to our primary estimates of the effects of the CWS. We first examine plant pollution emissions, emission intensity, and output. These estimates are reported in table 2. The first column reports estimates of the CWS's effects on plant $PM_{2.5}$ emissions. The second and third columns report estimates of the CWS's effects on plant emission intensity, measured as the emissions-sales ratio (column 2), or the emissions-value added ratio (column 3). Finally, columns 4 and 5 report estimates of the CWS's effects on plant output, measured in terms of sales (column 4) or value-added (column 5). All dependent variables are log transformed. Each regression is again weighted to correct for the potential bias introduced by matching the NPRI to the ASM. Standard errors clustered at the CMA-industry level are reported in parentheses.

Table 2.—The Effects of the CWS on Manufacturing Plants

	(1) $\ln(PM_{2.5})$	(2) $\ln(\frac{PM_{2.5}}{Sales})$	(3) $\ln(\frac{PM_{2.5}}{VA})$	(4) ln(Sales)	(5) ln(VA)
$PM_{2.5}$ Std.	-0.232^{***}	-0.097	-0.043	-0.135^{***}	-0.190^{***}
	(0.080)	(0.093)	(0.110)	(0.047)	(0.066)
O_3 Std.	-0.282	-0.228	-0.216	-0.053	-0.065
	(0.193)	(0.195)	(0.206)	(0.072)	(0.088)
R^2	0.940	0.928	0.914	0.969	0.941
N	3,978	3,978	3,978	3,978	3,978

Table reports estimates of the CWS's effects on manufacturing plants. In column 1, the dependent variable is the natural log of pollution emissions. In columns 2 and 3, the dependent variables are the natural log of plant emission intensity (the emission-sales ratio and the emissions-value added ratio, respectively). The dependent variable in column 4 is the natural log of plant sales, while the dependent variable in column 5 is the natural log of plant value added. 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. Significant at *10%, **5%, and ***1%.

The estimates reported in column 1 of table 2 indicate that the CWS's $PM_{2.5}$ standard led to a significant reduction in emissions, causing a 26.1% reduction in $PM_{2.5}$ emissions from affected plants. The O_3 standard also had a relatively large negative effect on $PM_{2.5}$ emissions, although this effect is imprecisely estimated and not statistically significant at conventional levels. This is potentially a result of co-pollutant effects, as reducing O_3 pre-cursor emissions may also directly reduce $PM_{2.5}$ emissions.[30]

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, Holland, and Mansur (2012) find California's NO_X trading program reduced NO_X emissions from regulated plants by between 10% and 30% over the period 1990 to 2005. Similarly, Gibson (2019) finds that Clean Air Act regulation reduced PM emissions from regulated plants by 38% between 1987 and 2014. This indicates that the CWS had similar effects on facility pollution levels as the environmental policies that have been enacted elsewhere.

The remaining columns of table 2 highlight the mechanisms driving this emissions reduction. The estimates reported in columns 2 and 3 suggest that the CWS had a relatively small and imprecisely estimated effect on the emission intensity of the average $PM_{2.5}$ emitting plant, while the estimates reported in columns 4 and 5 show that $PM_{2.5}$ regulation is associated with significant reductions in output from these plants. For example, our preferred estimate, reported in

[29]In both specifications, we pool all CMA-years with air quality below the bottom bin's threshold together and those with air quality above 34 $\mu g/m^3$ in the top bin.

[30]Note that in our context, there is a high correlation between O_3 precursor emissions and $PM_{2.5}$ emissions: 72% for nitrogen oxide, 83% for carbon monoxide, and 64% for volatile organic compounds.

TABLE 3.—THE EFFECTS OF THE CWS ON THE NUMBER OF PLANTS IN OPERATION

Dependent Variable:	Panel A: $PM_{2.5}$ Emitters		Panel B: CAC Emitters	
Number of Plants	(1) OLS	(2) Poisson	(3) OLS	(4) Poisson
$PM_{2.5}$ Std.	-1.045^{***}	-0.304^{***}	-1.375^{**}	-0.312^{*}
	(0.342)	(0.117)	(0.663)	(0.164)
O_3 Std.	0.082	-0.052	0.025	-0.116
	(0.362)	(0.130)	(0.594)	(0.175)
R^2	0.485	0.335	0.522	0.471
N	2,498	2,759	3,711	3,888
Dependent Variable Mean	4.056		4.898	

Table reports estimates of the CWS's effect on the number of plants operating in an industry-CMA-year. Panel A shows plants that emit $PM_{2.5}$, while panel B shows plants that emit any criteria air contaminant. In each panel, the first column reports OLS estimates, and the second reports Poisson estimates. The first row reports the effects of $PM_{2.5}$ regulation, and the second row reports the effects of the O_3 regulation. The average number of plants in a regulated industry-CMA in the base year is shown in the final row. All regressions include industry-year and CMA-year fixed effects, and an indicator for industry-CMAs that are ever treated by the CWS. Standard errors clustered by CMA are reported in parentheses. Significant at $^{*}10\%$, $^{**}5\%$, and $^{***}1\%$.

column 4, indicates that CWS particulate matter regulation is associated with a 14.5% decrease in sales from the average $PM_{2.5}$ emitting plant.[31]

There are few existing estimates to which these results can be compared. In terms of output, Greenstone (2002) finds U.S. carbon monoxide regulation implemented under the Clean Air Act reduced the growth rate of plant shipments by approximately 15%. However, Greenstone finds no significant effect of O_3, sulphur dioxide, or particulate matter regulation. In terms of emission intensity, there are no previous estimates we can use for comparison. The closest work is that of Martin, de Preux, and Wagner (2014), who show a carbon tax levied in the United Kingdom led to an 18% drop in energy intensity at affected manufacturing plants.

Next, we examine the net exit of plants in response to the CWS. To do so, we estimate equation (4) for two samples of plants. In panel A we consider plants that emit $PM_{2.5}$ and in panel B we consider plants that emit any criteria air contaminant (CAC). We include estimates from this sample to show that the results in panel A do not simply capture pollutant switching within plants.[32] As the dependent variable is a count variable, we estimate equation (4) using both ordinary least squares (OLS) and Poisson regression. In each panel, the OLS estimates are reported in the first column and the Poisson estimates in the second column. In all cases, standard errors clustered by CMA are reported in parentheses.

The estimates reported in table 3 suggest that the CWS had a significant effect on the number of plants operating in an industry-CMA-year. The estimates reported in column 1 of panel A, for example, show that $PM_{2.5}$ regulation reduced the number of $PM_{2.5}$ emitters in the average affected industry-CMA by 1.045 plants. This amounts to roughly 25% net exit in the average regulated industry-CMA. The magnitude of this effect is similar to previous estimates in the literature. For example, Becker & Henderson (2000) find that air quality regulation under the U.S. CAA reduced plant births in polluting industries by 26% to 45%.

The above results show that the CWS caused a significant net exit of $PM_{2.5}$-emitting plants and a significant reduction in output from surviving $PM_{2.5}$-emitting plants. The standard, however, had no significant effect on the pollution intensity of surviving plants.[33]

VI. Aggregate Implications

The goal of this paper is to quantify the degree to which air quality standards have contributed to the cleanup of manufacturing. Thus far, we have used quasi-experimental variation in regulatory stringency to show how a major Canadian air quality standard, the CWS, affected manufacturing plants. The CWS caused a net exit of $PM_{2.5}$-emitting plants and surviving plants to reduce output but had a relatively small effect on average plant pollution intensity. In this section, we develop an approach that uses these estimates to quantify the CWS's contribution to the cleanup of the Canadian manufacturing sector.

We start with an empirical analogue to equation (2), and build on the logic of a within-between decomposition, commonly used in both labor economics and productivity studies. Let t index time, such that industry i's pollution intensity at time t is $E_{it} = \int_0^{n_{it}} e_{it}(n)\lambda_{it}(n)dn$, where $e_{it}(n)$, $\lambda_{it}(n)$, and n_{it} are analogous to their counterparts in equation (2). Then the change in an industry's emission intensity between times $t-1$ and t can be expressed as

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

where we have assumed, for simplicity, that plants only exit. With some algebra, which we relegate to the online appendix, a convenient expression for the percentage change in an industry's emission intensity, $\dot{E}_{it} = \frac{E_{it}-E_{it-1}}{E_{it-1}}$, is

$$\dot{E}_{it} = \int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{\lambda}_{it}(n)dn - \int_{n_{it}}^{n_{it-1}} s_{zit-1}(n)dn + \int_0^{n_{it}} s_{zit-1}(n)\dot{e}_{it}(n)\dot{\lambda}_{it}(n)dn. \quad (9)$$

[31] It is worth noting that value added may provide a more accurate reflection of the level of productive activity that occurs in each plant (Cherniwchan et al., 2017). However, we focus our attention on sales as a measure of output to stay in line with the previous literature.

[32] That is, our estimates of net exit do not capture plants that switch from emitting $PM_{2.5}$ to some other pollutant so that they are still active producers but no longer $PM_{2.5}$ emitters.

[33] In the online appendix we also show that our results are not capturing large-emitter effects, within-firm reallocation or changes in the reporting of pollution data. In addition, we examine the potential mechanisms driving our results, and provide evidence that suggests large, fixed abatement costs are a key driver, as predicted by our model.

TABLE 4.—COUNTERFACTUAL ESTIMATES

	(1) Process Effect	(2) Reallocation Effect	(3) Selection Effect	(4) Interaction Effect	(5) Total
Share of Cleanup	0.081	0.243	0.067	−0.014	0.377

Table reports the process, reallocation, selection, and interaction effect shares of the change in manufacturing $PM_{2.5}$ intensity from 2004 to 2010. Columns 1 to 4 report estimates of each effect, and column 5 reports the total effect of the CWS.

In equation (9), $s_{zit-1}(n)$ is plant n's share of industry i's pollution at time $t-1$, and dot notation is used to denote percentage changes. The first three terms of equation (9) are the process, reallocation, and selection effects that we discussed previously in section IIA. The final term is an interaction effect created by the interaction between the process and reallocation effects and can be thought of as the approximation error in equation (2) caused by focusing on small, rather than potentially large, changes.

Equation (9) provides guidance on how to perform simple counterfactual exercises that ask how much of the technique effect can be attributed to each of the channels caused by a policy change. Constructing an estimate of the process effect, for example, would require an estimate of how plant pollution intensity changed as a result of a particular policy, as well as information on plant pollution shares. Estimating the reallocation effect would require an estimate of changes in plant market shares and information on pollution shares. Estimating the selection effect would require an estimate of the pollution removed due to plant exit. In the online appendix, we formalize our approach and show how we perform this exercise using our estimates of the CWS's effect on plant pollution intensity and output and the net exit of plants.

A. *Estimates of the CWS's Contribution to the Technique Effect*

We compare our estimates of the process, reallocation, selection, and interaction effects to the observed change in manufacturing pollution intensity over our sample period (2004 to 2010). This allows us to calculate the fraction of the cleanup attributable to each effect and to quantify the CWS's contribution to Canada's manufacturing cleanup.[34] For simplicity, we focus solely on the direct effect of the CWS $PM_{2.5}$ standard on $PM_{2.5}$ emitters.

Our estimates of each effect's contribution to the observed change in manufacturing emission intensity are reported in table 4. Columns 1 to 4 report our estimates of the process effect, reallocation effect, selection effect, and interaction effect. Column 5 reports the fraction of the change in aggregate emission intensity that can be explained by the CWS.

The results of this exercise show that the $PM_{2.5}$ standard enacted under the CWS is responsible for approximately 38% of the reduction in manufacturing $PM_{2.5}$ intensity that occurred between 2004 and 2010. As this was the main policy regulating $PM_{2.5}$ pollution in Canada over this period, this suggests other factors played an important role in the cleanup.

Our results also shed light on the relative importance of each effect to the cleanup. $PM_{2.5}$ regulation primarily reduced aggregate emission intensity through a combination of reallocation and selection effects, which combined explain 31% of the cleanup. The process effect, in contrast, accounts for roughly 8% of the cleanup for $PM_{2.5}$. Estimates of the relative magnitude of these effects are important on their own right, because they signal both how environmental policy affects the broader economy and the incidence of these policies. In this case, the most pronounced effect of the CWS was to cause plants to reduce output (the reallocation effect) rather than to adopt cleaner production processes (the process effect).

VII. Conclusion

In this paper, we examine how plant-level responses to air quality standards have contributed to the cleanup of the manufacturing sector. To do so, we estimate the effects of a major revision to Canadian air quality regulation, the Canada-Wide Standard for Particulate Matter and Ozone (CWS), using a novel, confidential data set with pollution and production information for a panel of Canadian manufacturing plants.

We present a theoretical model to show how commonly used air quality regulations, such as those used by the CWS, reduce an industry's emission intensity by causing a reduction in output from regulated plants (a reallocation effect), plant entry and exit (a selection effect), and a reduction in the emission intensity of surviving plants (a process effect).

Next, we estimate the effects of the CWS on Canadian manufacturing $PM_{2.5}$-emitting plants and use the resulting empirical estimates to quantify how air quality regulations have contributed to the cleanup of $PM_{2.5}$ from Canadian manufacturing. These estimates represent the first complete characterization of how plant-level responses to air quality regulation have contributed to a cleanup. While our estimates are specific to Canada, given the similarity between the cleanups and regulatory structures in Canada, the United States, and Europe, we believe our results provide insights relevant for all three regions.

Our estimates imply that the CWS explains close to 40% of the reduction in the $PM_{2.5}$ intensity of the Canadian manufacturing sector from 2004 to 2010. Thus, our quasi-experimental evidence corroborates an important lesson from the quantitative model of Shapiro and Walker (2018): environmental regulations have played a key role in the cleanup of manufacturing. However, our results also suggest that a regulation-induced cleanup may involve significant economic disruption, as over 80% of the cleanup caused by the CWS was due to plant exit (the selection effect) and changes in output from regulated plants (the reallocation effect).

[34] Recall that 99% of Canada's $PM_{2.5}$ cleanup was caused by the technique effect, meaning this exercise is illustrative of CWS's contribution to the cleanup of manufacturing.

REFERENCES

Andersen, Dana C., “Accounting for Loss of Variety and Factor Reallocations in the Welfare Cost of Regulations,” *Journal of Environmental Economics and Management* 88 (2018), 69–94.

Anoulies, Lisa, “Heterogeneous Firms and the Environment: A Cap-and-Trade Program,” *Journal of Environmental Economics and Management* 84 (2017), 84–101.

Barrows, Geoffrey, and Hélène Ollivier, “Cleaner Firms or Cleaner Products? How Product Mix Shapes Emission Intensity from Manufacturing,” *Journal of Environmental Economics and Management* 88 (2018), 134–158.

Becker, Randy, and Vernon Henderson,“Effects of Air Quality Regulation on Polluting Industries,” *Journal of Political Economy* 108 (2000), 379–421.

Bertrand, M, E Duflo, and S Mullainathan, “How Much Should We Trust Differences-in-Differences Estimates?” *Quarterly Journal of Economics* 119 (2004), 249–275.

Brunel, Claire, “Pollution Offshoring and Emissions Reductions in E.U. and U.S. Manufacturing,” *Environmental and Resource Economics* 68 (2017), 621–641.

Bustos, Paula, “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms,” *American Economic Review* 101 (2011), 304–340.

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

CCME, “Joint Initial Actions to Reduce Pollutant Emissions That Contribute to Particulate Matter and Ground-level Ozone” (2000), http://www.ccme.ca/files/Resources/air/pm_ozone/pmozone_joint_actions_e.pdf/.

Cherniwchan, Jevan, “Trade Liberalization and the Environment: Evidence from NAFTA and U.S. Manufacturing,” *Journal of International Economics* 105 (2017), 130–149.

Cherniwchan, Jevan, Brian R. Copeland, and M. Scott Taylor, “Trade and the Environment: New Methods, Measurements, and Results,” *Annual Review of Economics* 9 (2017), 59–85.

Environment Canada, “Multi-Pollutant Emission Reduction Analysis Foundation (MERAF) for the Iron and Steel Sector,” unpublished report (2002). [Available from the authors.]

Environment Canada and FPAC,“Towards More Innovative Air Quality Management: Proposal for a Pulp and Paper Air Quality Forum” (2004), http://publications.gc.ca/site/eng/274270/publication.html.

Environment and Climate Change Canada, “Summary of National Pollutant Release Inventory Reporting Requirements” (2016), https://www.ec.gc.ca/inrp-npri/default.asp?lang=en&n=629573FE-1.

European Parliament, “Directive 2010/75/EU of the European Parliament and of the Council of 24 November 2010 on Industrial Emissions (Integrated Pollution Prevention and Control)” (2010), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32010L0075.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur, “What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NO_x Trading Program,” *American Economic Review* 102 (2012), 965–993.

Gibson, Matthew, “Regulation-Induced Pollution Substitution,” this REVIEW 101 (2019), 827–840.

Government of Canada, “Clean Air in Canada: Progress Report on Particulate Matter and Ozone” (2003), http://publications.gc.ca/pub?id=9.664805&sl=0.

Graff Zivin, Joshua, and Matthew Neidell, “Environment, Health and Human Capital,” *Journal of Economic Literature* 51 (2013), 689–730.

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

Grether, Jean-Marie, Nicole A. Mathys, and Jaime de Melo, “Scale, Technique and Composition Effects in Manufacturing SO_2 Emissions,” *Environmental and Resource Economics* 43 (2009), 257–274.

Grossman, Gene M., and Alan B Krueger, “Environmental Impacts of a North American Free Trade Agreement” (pp. 13–56) in Peter M. Garber, ed. *The Mexico-US Free Trade Agreement* (Cambridge, MA: MIT Press, 1993).

Henderson, J. Vernon, “Effects of Air Quality Regulation,” *American Economic Review* 86 (1996), 789–813.

Levinson, Arik, “Technology, International Trade, and Pollution from U.S. Manufacturing,” *American Economic Review* 99 (2009), 2177–2192.

——— “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 (2015), 43–56.

Martin, Ralf, Laure B. de Preux, and Ulrich J. Wagner, “The Impact of a Carbon Tax on Manufacturing: Evidence from Microdata,” *Journal of Public Economics* 117 (2014), 1–14.

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

Melitz, Marc J., and Giancarlo I. P. Ottaviano, “Market Size, Trade, and Productivity,” *Review of Economic Studies* 75 (2008), 295–316.

Ryan, Stephen P., “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica* 80 (2012), 1019–1061.

Shapiro, Joseph S., and Reed Walker, “Why Is Pollution from U.S. Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade,” *American Economic Review* 108 (2018), 3814–3854.

Tombe, Trevor, and Jennifer Winter, “Environmental Policy and Misallocation: The Productivity Effect of Intensity Standards,” *Journal of Environmental Economics and Management* 72 (2015), 137–163.

World Bank Group, “Pollution Prevention and Abatement Handbook 1998” (1998), http://documents.worldbank.org/curated/en/758631468314701365/pdf/multi0page.pdf.