Trade and the Environment: New Methods, Measurements, and Results

Jevan Cherniwchan,[1] Brian R. Copeland,[2] and M. Scott Taylor[3,4]

[1]Alberta School of Business, University of Alberta, Edmonton, Canada T6G 2R; email: jevan@ualberta.ca

[2]Vancouver School of Economics, University of British Columbia, Vancouver, Canada V6T 1L4

[3]Department of Economics, University of Calgary, Calgary, Canada T2N 1N4

[4]National Bureau of Economic Research, Cambridge, Massachusetts 02138

Annu. Rev. Econ. 2017. 9:59–85

The *Annual Review of Economics* is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-063016-103756

JEL codes: F6, F18, Q56

Keywords

globalization, trade, heterogeneous firms, pollution, environment

Abstract

We review recent research linking international trade to the environment, with a focus on new results and methods. The review is given structure by a novel decomposition linking changes in emissions to changes in productive activity at the plant, firm, industry, and national levels. Although some new results have emerged from the application of a Melitz-style approach to trade and the environment, the full potential of this approach has not yet been realized. We discuss existing empirical and theoretical work, introduce three new hypotheses, and suggest paths for future researchers to follow.

1. INTRODUCTION

The relationship between international trade and the environment has been viewed largely through the lens of comparative advantage. The focus of researchers on the emergence of pollution havens and the environmental impact of growth was natural given the policy concerns in the lead-up to the signing of the North American Free Trade Agreement (NAFTA) and the rapid growth in several developing countries. Today's concerns differ little; what has changed is the set of tools available to examine firm-level adjustments to trade liberalization. This review places these new methods, tools for measurement, and their results within the broader context of the literature and asks what this new focus has brought to the trade and environment debate.

A firm-level focus in answering trade and environment questions is very promising, but researchers have not yet fully exploited its potential. There are many new insights, but much remains poorly understood. Theories in which comparative advantage drives across-industry adjustment are often treated as competitors to theories based on within-industry adjustments, rather than as complements studying different units of analysis. As a result, there is little work that attempts to integrate empirical findings from the old and the new approaches. Moreover, there are several new and potentially important hypotheses that cry out for further study.

To help us organize the literature, we develop two tools. The first is a decomposition that attributes emissions to productive activity at the economy, industry, firm, and plant levels. The second is a simple partial equilibrium model of firm behavior. The decomposition identifies the set of possible adjustments to trade liberalization; the model then establishes causal connections to these same adjustments. Decompositions alone tell us nothing about causality; a model of firm and within-industry adjustments gives us an incomplete picture of economy-wide responses. Pairing the two allows us to reconcile a focus on new microlevel mechanisms with previously studied macrolevel channels.

Decompositions have been used to understand how international trade affects the environment since the work of Grossman & Krueger (1993) and Copeland & Taylor (1994); in this review, we move past the typical focus on national and industry-level aggregates to allow for both within-industry and within-firm changes in productive activity. We do so because researchers have shown that there is considerable heterogeneity in both emission levels and emission intensities across firms in even quite narrowly defined industries. By allowing for adjustments at the firm level, we can discuss how trade-inspired decisions to outsource production, invest in abatement, or offshore intermediates can affect measured emissions. Because empirical work now exploits plant-level observations, this level of detail seems unavoidable, and allowing for intermediate good trade opens a useful discussion of measurement issues related to final sales versus value added.

Our partial equilibrium model is deliberately simple. Firms are differentiated by productivity because across-firm heterogeneity is key to discussions of how selection and market-share reallocations affect emissions. We allow firms to choose abatement and produce intermediate goods at home or to outsource production abroad. This simple model guides our discussion of both theory and empirical work. Model details and some derivations are left to the **Supplemental Appendix** (to access these materials, follow the **Supplemental Material link** in the online version of this article or at **http://www.annualreviews.org/**).

Any review has to decide on its sins of omission. To maintain intellectual clarity, we focus on trade's effect on industrial pollution, leaving its impact on consumption-generated pollution,[1] resource use, and natural habitats untouched. We do not review the work on the effects of trade

[1]There is relatively little empirical work on the effects of trade on consumption-generated pollution. For an interesting example, the reader is referred to Davis & Kahn (2010). There is a larger literature on regulating consumption-generated pollution in open economies (see, for example, McAusland 2008, Copeland 2011).

on emissions from international transport, but we do discuss empirical studies employing data on carbon emissions (for recent empirical studies of the effects of international transport on carbon emissions, see Cristea et al. 2013, Shapiro 2016). To remain focused on empirical evidence, we do not discuss findings drawn from computable general equilibrium or simulation models. Although these exclusions are unfortunate, we provide a reasonably cohesive and constructive review of new theoretical and empirical research contributions to what has been the main body of research examining the environmental impact of trade.

Our review makes three contributions. We bring the interested reader up to speed on the latest research and glean from this new research its key insights. We then place these new contributions within the context of the original research program to understand how international trade affects environmental outcomes. Finally, we develop a set of new hypotheses to evaluate novel predictions coming from models with firm-level heterogeneity. The pollution reduction by rationalization hypothesis links market share reallocations and selection effects in the Melitz (2003) model to changes in industry emissions. The distressed and dirty industry hypothesis links changes in abatement and emission intensities to heightened foreign competition brought about by trade liberalization. Finally, the pollution offshoring hypothesis links firm-level decisions to offshore dirty intermediate inputs to trade liberalization with a partner that differs greatly in their pollution policy. There is no empirical work that explicitly addresses these hypotheses, although many existing papers contain evidence relevant to their evaluation.

The remainder of this review proceeds as follows. In Section 2, we develop our new decomposition. With this in hand, in Section 3 we revisit existing evidence and review new evidence on the core hypotheses that have guided the literature thus far. This review provides context for the study of theory and empirical work on firm-level adjustments in Section 4. Section 4 also includes a discussion of our three new hypotheses. Section 5 concludes.

2. THE MECHANICS OF POLLUTION EMISSIONS

We begin with a simple accounting exercise. Our objective is to decompose changes in the aggregate emissions of some pollutant Z into components reflecting changes in industrial activities within an economy. This type of decomposition has been influential in the literature on trade and the environment but has typically focused on industry-level outcomes. We build on this earlier work by allowing for changes in the mix of firms within each industry, the mix of productive activities at each firm, and the emission intensity of each activity.

Consider an economy with aggregate pollution emissions Z generated by N industries. Each industry i emits Z_i units of pollution. If we let S_i denote the scale of production in industry i (measured as domestic value added at base period prices), the economy's aggregate emissions can be written as

$$Z = \sum_{i=1}^{N} S_i E_i, \qquad 1.$$

where $E_i = Z_i/S_i$ is the emission intensity of industry i.

Taking logs and differentiating yield

$$\hat{Z} = \hat{S} + \sum_{i=1}^{N} \Theta_i \hat{\Phi}_i + \sum_{i=1}^{N} \Theta_i \hat{E}_i, \qquad 2.$$

where $S = \sum_{i=1}^{N} S_i$ is the economy-wide scale of output (measured by real GDP), $\Theta_i = Z_i/Z$ is the fraction of overall emissions Z coming from industry i, $\Phi_i = S_i/S$ is industry i's share of the economy's final output, and $\hat{Z} = \mathrm{d}Z/Z$, etc.

Equation 2 is an industry-level decomposition similar to that introduced by Grossman & Krueger (1993) and Copeland & Taylor (1994) and subsequently used by many authors (see Levinson 2009 for a recent example). Changes in pollution are decomposed into three channels. The first term in Equation 2 is the scale effect and represents the change in pollution due to a change in the overall level of economic activity. The second term is the composition effect, which reflects the change in pollution due to changes in the composition of economic activity across industries. The third term is the technique effect; it reflects changes in pollution due to changes in the emission intensities of each industry. Adding up these responses yields the full effect of a shock such as trade liberalization.

While the industry-level decomposition in Equation 2 has been influential in shaping our understanding of what drives changes in aggregate pollution emission levels, it tells us little about the microlevel adjustments generating change at the industry level. We therefore proceed with a firm-level decomposition. Referring to Equation 2, this means that we will decompose changes in industry-level emission intensities into effects determined by changes at the firm level.

Suppose each industry i has a continuum of firms on the interval $[0, n_i]$, where n_i is the marginal firm that is endogenously determined by the industry's profitability.[2] Let $z_i(n)$ denote the emissions produced by firm n. Aggregate industry emissions are given by

$$Z_i = \int_0^{n_i} z_i(n)\mathrm{d}n. \tag{3}$$

Denoting the value added produced by firm i as $v_i(n)$, we can define the scale of output in industry i as

$$S_i = \int_0^{n_i} v_i(n)\mathrm{d}n. \tag{4}$$

Using Equation 4, we can write the emission intensity of industry i as

$$E_i = \frac{Z_i}{S_i} = \int_0^{n_i} e_i(n)\varphi_i(n)\mathrm{d}n, \tag{5}$$

where

$$e_i(n) = \frac{z_i(n)}{v_i(n)} \tag{6}$$

is the emission intensity of firm n, and

$$\varphi_i(n) = \frac{v_i(n)}{S_i} \tag{7}$$

is the share of firm n in the value of production in industry i. Hence, the emission intensity of industry i is a weighted average of the emission intensities of the firms in the industry. Taking logs and differentiating yields, we obtain

$$\hat{E}_i = \int_0^{n_i} \hat{e}_i(n)\theta_i(n)\mathrm{d}n + \int_0^{n_i} \hat{\varphi}_i(n)\theta_i(n)\mathrm{d}n + n_i[\theta_i(n_i) - \varphi_i(n_i)]\hat{n}_i, \tag{8}$$

where $\theta_i(n) = z_i(n)/Z_i$ is firm n's share of emissions in industry i.

[2] We are not the first to develop a firm-level decomposition. The literature on productivity has used decompositions that allow for heterogeneous firms for many years (see, for example, Foster et al. 2008), and several authors in the trade and environment literature have recently adopted these methods (see, for example, Martin 2012, Cherniwchan et al. 2013, Barrows & Ollivier 2016).

Equation 8 shows how a change in industry i's average emission intensity reflects three distinct within-industry changes. The first term captures changes in firm-level emission intensities. The second term is an industry composition effect—industry average emission intensities rise if the market share of firms with higher emission intensity rises. The final term captures the impact of entry and exit. Its magnitude and sign depend on the difference between the marginal firm's share of industry emissions and its share of industry output. Industry average emission intensity rises if an entering firm is more pollution intensive than the average firm.

When firms are identical, the second two terms in the decomposition drop out and the first term simplifies to $\hat{e}_i$, which is the emission intensity change experienced by all firms in industry i. When firms are not identical, all three terms come into play. If we do not account for these channels, then we may misidentify the way that abatement and emission intensities adjust to policy changes and other shocks.

For example, suppose environmental regulation is tightened and industry-level emission intensities fall. We may falsely conclude that firm emission intensities are very responsive to regulation when in fact they are not. Even if firm-level emission intensities do not change [$\hat{e}_i(n) = 0$ for all firms], the industry's emission intensity could have fallen because relatively dirty firms lost market share (the second term is negative) or the dirtiest firms left the industry (the last term is negative). This distinction is likely to be important. For many policy questions, we would like to understand whether the adjustment occurs primarily via within-firm abatement, across-firm market share reallocations, or changes in firm numbers and therefore industry competition.

In addition, it is common to relate industry-wide emission intensities (or pollution abatement costs) to outcome variables such as exports, imports, productivity, etc. However, once we admit that these industry-level aggregates are a function of the type, size, and number of firms present, then they are no longer primitive characteristics of the industries but rather endogenous outcomes that are likely related to exports, imports, productivity, etc. This endogeneity of the aggregates means that special care has to be taken in exploiting across-industry variation in these primitives to explain variation in outcome variables.[3]

Although a within-industry decomposition is useful in clarifying the set of potential adjustments within industries, should we drill down further to examine potential adjustments within firms and across plants? Empirical work suggests that there exists substantial plant-level heterogeneity within firms in terms of production processes, engagement in trade, etc., and that this may make within-firm adjustments relevant to pollution outcomes. Moreover, outsourcing production to other firms will also affect emission intensities at the firm level.

For these reasons, we now examine the potential role of firm-level adjustments in determining emission levels. Suppose each firm n produces output $y_i(n)$ by completing a set of $M_i(n)$ tasks (although M_i varies by firm, we suppress the n in what follows to economize on notation). Let T_i be a measure of the scale of task i. We then assume that the firm's output is

$$y_i(n) = f_{i,n}(T_1, T_2, \ldots, T_{M_i}), \qquad 9.$$

where $f_{i,n}$ is weakly increasing in all of its arguments. Each task may be performed entirely by the firm at one or more of its plants, or it may be outsourced to either domestic or foreign producers.

[3] Levinson & Taylor (2008) were the first to recognize these issues and discuss them within a neoclassical framework where sectors (the unit of analysis for empirical work) were themselves composed of many heterogeneous industries.

These tasks could be producer services contracted for by the firm, or task completion could be embodied in intermediate goods delivered to the firm as part of their production process.

Let $\lambda_{ij}^I(n)$ be the fraction of task j performed in-house by firm n, let $\lambda_{ij}^d(n)$ be the fraction of task j outsourced domestically, and let $\lambda_{ij}^*(n)$ be the fraction of task j completed offshore. We require

$$\lambda_{ij}^I(n) + \lambda_{ij}^d(n) + \lambda_{ij}^*(n) = 1. \qquad 10.$$

If the task is entirely offshored, then $\lambda_{ij}^* = 1$; alternatively, if it is carried out entirely within the firm, then $\lambda_{ij}^I = 1$. Let $p_i(n)$ denote the price used to value output, and let w_{ij} be the price used to value a unit of task j wherever it is completed. Let $\mu_i(n)$ be the rate at which firm n marks up the unit cost of tasks; that is, define $\mu_i(n)$ such that

$$p_i(n)y_i(n) = [1 + \mu_i(n)] \sum_{j=1}^{M_i} w_{ij}(n)T_{ij}(n). \qquad 11.$$

We can then write the value added produced by firm n as

$$\begin{aligned} v_i(n) &= p_i(n)y_i(n) - \sum_{j=1}^{M_i}[1 - \lambda_{ij}^I(n)]w_{ij}(n)T_{ij}(n) \\ &= \sum_{j=1}^{M_i} \lambda_{ij}^I(n)w_{ij}(n)T_{ij}(n) + \mu_i(n)\sum_{j=1}^{M_i} w_{ij}(n)T_{ij}(n). \end{aligned} \qquad 12.$$

Value added is the market value of final sales minus the cost of tasks completed elsewhere. This is equivalent to the value added of tasks produced in-house plus the return generated by the firm's markup (which is the value added created by the right to assemble, market, and brand the firm's unique product).

The completion of each task potentially generates some pollution (for simplicity, we assume that assembling the final good by aggregating tasks does not generate additional pollution). Denote the firm's domestic emission intensity of task j (measured as emissions per unit value generated by the task) as $e_{ij}(n)$. The level of the firm's domestic emissions from task j is then

$$z_{ij}(n) = e_{ij}(n)\lambda_{ij}^I(n)w_{ij}(n)T_{ij}(n). \qquad 13.$$

The total level of pollution emitted domestically by firm n in industry i is the sum of emissions from its domestic plants,

$$z_i(n) = \sum_{j=1}^{M_i} z_{ij}(n) = \sum_{j=1}^{M_i} \lambda_{ij}^I(n)e_{ij}(n)w_{ij}(n)T_{ij}(n), \qquad 14.$$

so the overall emission intensity of firm n (per dollar of value added) can be written as

$$e_i(n) = \frac{z_i(n)}{v_i(n)} = \frac{\sum_{j=1}^{M_i} \lambda_{ij}^I(n)e_{ij}(n)\sigma_{ij}(n)}{\sum_{j=1}^{M_i} \lambda_{ij}^I(n)\sigma_{ij}(n) + \mu_i(n)}, \qquad 15.$$

where

$$\sigma_{ij}(n) = \frac{w_{ij}(n)T_{ij}(n)}{\sum_{j=1}^{M_i} w_{ij}(n)T_{ij}(n)} \qquad 16.$$

is the share of task j in the total cost of all tasks.

To obtain a decomposition of the emission intensity of firm n in industry i, take logs and totally differentiate Equation 15:

$$\begin{aligned}\widehat{e}_i(n) =& \sum_{j=1}^{M_i} \theta_{ij}(n)\widehat{e}_{ij}(n) + \sum_{j=1}^{M_i} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\sigma}_{ij}(n) \\ &- \sum_{j=1}^{M_i} \frac{\lambda_{ij}^d(n)}{\lambda_{ij}^I(n)} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\lambda}_{ij}^d(n) \\ &- \sum_{j=1}^{M_i} \frac{\lambda_{ij}^*(n)}{\lambda_{ij}^I(n)} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\lambda}_{ij}^*(n) - \varphi_{i\mu}(n)\widehat{\mu}_i(n), \qquad 17.\end{aligned}$$

where $\theta_{ij}(n) = \lambda_{ij}^I(n) z_{ij}(n)/z_i(n)$ is the fraction of firm n's in-house emissions generated by task j, $\varphi_{ij}(n)$ is the share of the firm's in-house production of task j in value added, and $\varphi_{i\mu}(n)$ is the share of revenue from markups in value added.[4]

Equation 17 highlights how even a firm-level decomposition conceals some potentially important responses driving aggregate changes in emissions. As the equation shows, a change in any firm's emission intensity depends on changes along four different margins.

The first term in Equation 17 reflects changes in the emission intensity of each task conducted at each domestic plant. These changes could be precipitated by changes in environmental policy, input prices, technology, or abatement. This term is the firm's true technique effect because it represents a weighted average of its plant-level changes in the techniques of production.

The second term is a firm-level composition effect. It captures the changes in the importance of various tasks required for production. The average emission intensity of a firm will change depending on whether cleaner or dirtier tasks become more intensively used. We refer to this as the firm reorganization effect.

The third and fourth terms reflect changes in the extent of outsourcing of production to other domestic firms and to foreign producers, respectively. We refer to the third term as the domestic outsourcing effect and to the fourth as the offshoring effect. To interpret these terms, note that an increase in outsourcing of task j will reduce firm-level emission intensities if the share of in-house emissions from the outsourced task is greater than the share of in-house production of that task in value added. In other words, emission intensity falls if outsourced tasks are relatively emission intensive.

Firms may either outsource tasks to other domestic firms or offshore them to other countries. The implications for aggregate domestic emissions depend on whether the outsourcing is domestic or offshored. For domestic outsourcing, a fall in firm n's emission intensities may be offset by the emissions generated by an increase in production at some other domestic firm (possibly in a different industry, depending on what has been outsourced). This will appear elsewhere in the decomposition via potential changes in the scale, composition, and emission intensities of other domestic firms. In contrast, if the task is offshored, then the change in the firm's emission intensities will have a direct effect on aggregate emissions in the domestic economy. In either case, it is important to distinguish between emissions per unit of value added, which is given in Equation 15, and emissions per unit sales, which is given by the numerator of Equation 15 divided by the markup $1 + \mu_i(n)$. Emissions per unit of sales is always less than emissions per unit value

[4]For simplicity in presentation, we have shown the decomposition for the case where initial levels of $e_{ij}(n)$, $\sigma_{ij}(n)$, $\lambda_{ij}^d(n)$, $\lambda_{ij}^*(n)$, and μ are all positive. If any of these is zero, the corresponding term cannot be written in percent change form, and we would have to focus on the absolute rather than relative change.

added when outsourcing occurs. This gap is increasing in the extent of outsourcing, is increasing in the share of those inputs heavily outsourced, and is larger in industries with smaller markups. Any one of these three effects might occur with trade liberalization.

Finally, the last term in our emission intensity decomposition is the effect of a change in the firm's markups. We have assumed that the activities giving the firm the ability to charge markups do not generate pollution, and therefore we can think of markups as a nonpolluting activity. An increase in the share or level of markups in value added will lower the firm's emission intensity.

It should be apparent from our earlier discussion that ignoring these potential adjustments may come at some cost. Firms may get cleaner without emission intensities in any of their plants falling, and their emission intensities can respond in quite complicated ways to changes brought about by trade liberalization.

Combining Equations 17 and 8 with Equation 2 yields our now much more detailed decomposition:

$$
\begin{aligned}
\hat{Z} = {} & \hat{S} + \sum_{i=1}^{N} \Theta_i \hat{\Phi}_i + \sum_{i=1}^{N} \Theta_i \int_0^{n_i} \hat{\varphi}_i(n)\theta_i(n)\mathrm{d}n + \sum_{i=1}^{N} \Theta_i n_i [\theta_i(n_i) - \varphi_i(ni)]\hat{n}_i \\
& + \sum_{i=1}^{N} \Theta_i \int_0^{n_i} \left[\sum_{j=1}^{M_i} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\sigma}_{ij}(n) \right] \theta_i(n)\mathrm{d}n \\
& - \sum_{i=1}^{N} \Theta_i \int_0^{n_i} \left[\sum_{j=1}^{M_i} \frac{\lambda_{ij}^d(n)}{\lambda_{ij}^I(n)} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\lambda}_{ij}^d(n) \right] \theta_i(n)\mathrm{d}n \\
& - \sum_{i=1}^{N} \Theta_i \int_0^{n_i} \left[\sum_{j=1}^{M_i} \frac{\lambda_{ij}^*(n)}{\lambda_{ij}^I(n)} \left[\theta_{ij}(n) - \varphi_{ij}(n)\right] \widehat{\lambda}_{ij}^*(n) \right] \theta_i(n)\mathrm{d}n \\
& + \sum_{i=1}^{N} \Theta_i \int_0^{n_i} \left[\sum_{j=1}^{M_i} \theta_{ij}(n)\widehat{e}_{ij}(n) \right] \theta_i(n)\mathrm{d}n - \sum_{i=1}^{N} \Theta_i \int_0^{n_i} [\varphi_i(n)\widehat{\mu}_i(n)]\, \theta_i(n)\mathrm{d}n. \qquad 18.
\end{aligned}
$$

The first term is the economy-wide scale effect; the second is the across-industry composition effect. Both appeared in the above decomposition in Equation 2. The remaining terms compose a decomposition of the classic technique effect. Industry emission intensities are affected by changes in the composition of firms within the industry and changes within each firm. The third and fourth terms capture the effects of changes in firm market shares and entry and exit, respectively. The other terms all capture within-firm adjustments: The fifth term is the within-firm reorganization effect; the next two terms capture the effects of domestic outsourcing and offshoring, respectively; the second-to-last term is the effect of direct changes in emission intensities of tasks within firms; and the final term captures the effects of changes in firm-level markups.

The key insight that comes from using this more detailed decomposition is that changes in pollution that have previously been attributed to changes in industry-level emission intensities are in fact influenced by within-industry composition effects, within-firm reorganization effects, and directly by trade via offshoring, in addition to plant- or task-level changes in emission intensities. A key theme throughout the remainder of the review is the extent to which a focus on these more detailed channels of adjustment helps us understand how trade affects the environment.

3. COMPARATIVE ADVANTAGE AND THE ENVIRONMENT

Early work assumed that comparative advantage was the key driver of interindustry trade flows; hence, the industry-level decomposition (Equation 2) was a natural starting point (for detailed expositions of the comparative advantage approach, see Copeland & Taylor 2004, Copeland 2011). This approach identifies three channels by which trade affects the environment:

1. Trade raises the scale of economic activity, which increases pollution.
2. Trade raises real income and the demand for environmental quality. If governments are responsive, policies are strengthened and pollution falls via a technique effect.
3. Controlling for incomes and scale, changes in the sectoral composition of clean and dirty industries affect emissions. These trade-created composition effects vary across countries depending on their comparative advantage.

Much of this work was motivated by the pollution haven hypothesis (PHH), which asserts that, following a reduction in trade barriers, pollution-intensive industries will contract in countries with relatively strong environmental regulation and expand in those where environmental policy is relatively weak, meaning that environmental policy differences serve as an important source of comparative advantage (for a detailed overview of the PHH, see Taylor 2005). Although this hypothesis grew out of popular concerns, it has a strong grounding in theory: Pollution havens can arise from income-induced differences in environmental policy (Copeland & Taylor 1994, 1995), differences in institutional capacity or property rights (Chichilnisky 1994, Brander & Taylor 1998), or differences in environmental carrying capacity (Copeland & Taylor 2003).

A necessary condition for the PHH to hold is that environmental policy differences translate into large differences in production costs. Therefore, an important research question is simply whether more stringent environmental policy adversely affects comparative advantage. We refer to this as the pollution haven effect (PHE). Only if this effect is strong enough to dominate other sources of comparative advantage will it determine the pattern of trade in dirty industries. Antweiler et al. (2001) and Copeland & Taylor (2003) study a model where policy differences across countries interact with capital abundance to determine comparative advantage. If pollution-intensive industries are also capital intensive, then a capital-abundant country with relatively stringent environmental policy may export the pollution-intensive good. In this case, a PHE exists (tightening up environmental policy weakens comparative advantage in polluting industries), but the PHH nevertheless fails.

3.1. Do Environmental Regulations Affect Trade Flows?

A large empirical literature studies the PHE or, in other words, sets out to test the hypothesis that more stringent environmental policy reduces competitiveness. Initially, this literature generated a puzzle: Despite the seemingly weak assumptions needed to generate a PHE, most of the empirical work prior to the late 1990s (and some subsequent work) found either a zero or a positive effect of more stringent environmental policy on net exports (see, for example, Jaffe et al. 1995). Subsequent work attributed these paradoxical results to problems arising from endogenous policy. For example, successful, large, and competitive industries may draw greater scrutiny and regulation, whereas politicians will be reluctant to tighten environmental policy in sectors facing heavy import competition. Empirical work accounting for unobserved heterogeneity and endogenous policy has found that more stringent environmental policy adversely affects competitiveness (for reviews, see Brunnermeier & Levinson 2004, Copeland & Taylor 2004, Levinson 2010).

In this section, we provide a selective overview of recent work on the effects of environmental policy, with the aim of highlighting key empirical issues and promising strategies used in the literature.

Levinson & Taylor (2008) examine the effects of environmental regulations on bilateral trade between the United States and Mexico and between the United States and Canada over the period

1977–1986. They develop a simple partial equilibrium model to illustrate how the use of indirect measures of environmental regulations [such as pollution abatement costs (PACs)] may imply downward-biased estimates. This suggests that earlier findings of small or nonexistent PHEs were largely due to the presence of omitted variables and measurement error.

To address these issues, they adopt a panel instrumental variables (IV) approach and construct instruments using geographic variation in factors that affect pollution demand and supply. Their results (using the panel IV approach) suggest that environmental regulations have a large, significant effect on trade flows: A 1% increase in PACs increases net imports into the United States from Mexico and Canada by 0.4% and 0.6%, respectively. Moreover, their panel IV estimates are much larger in magnitude than the corresponding ordinary least squares estimates, highlighting the downward bias arising from using PACs as a measure of regulatory stringency.

The usefulness of a close connection between theory and empirical work is further illustrated by Kellenberg (2009), who studies the effects of environmental policy on the activities of US multinational firms in 50 countries over the period 1999–2003. He outlines a simple game in which production decisions of multinational firms are affected by strategic environmental policy choices of governments. This model forms the basis of his research design: He uses the model to both derive an estimating equation and motivate the use of neighboring country characteristics as instruments to address the potential endogeneity of environmental policy. Using this IV approach, he finds that weak environmental policy is associated with increased activity by US multinationals, providing further evidence of the PHE.

One issue with the aforementioned studies is that they rely on research designs that employ model-based arguments for identification. This makes it difficult to ensure that the resulting estimates are causal; if the theoretical model is misspecified, it is likely that corresponding identification assumptions will not hold. One way to address this issue is to use policy changes or other shocks as sources of identifying variation.

Several researchers have used the US Clean Air Act as a source of such variation (see, for example, Becker & Henderson 2000). The Clean Air Act mandates federal air quality standards, which are enforced at the county level. Counties whose air quality does not meet these standards (nonattainment counties) must tighten up environmental regulations. This yields a source of policy variation across locations. In addition, nonattainment status is reevaluated each year, which results in temporal variation. Hanna (2010) uses a panel of firm-level data on US manufacturing firms from 1966 to 1999 and finds that the Clean Air Act caused US multinationals to increase their foreign assets by 5% and their foreign sales by 9%. Larger firms account for most of this effect. Her results provide support for the PHE: More stringent US regulation shifts production in affected industries out of the United States. A natural question is, where does that production move to? If the PHH held, we would expect production to move to countries with weak environmental regulations and, in particular, to developing countries. Hanna (2010) finds no evidence to support this, which is consistent with the view that pollution regulation is one of only many factors that affect comparative advantage.

Aichele & Felbermayr (2015) exploit a different source of policy variation. They study the effects of the Kyoto Protocol on the carbon content of trade for 15 industries in 40 countries over the period 1995–2007.[5] They sketch out a simple model that yields a gravity equation linking the quantity of bilateral trade between two countries to cross-country differences in carbon prices. In the absence of actual data on environmental regulation, they exploit differences in country commitments under the Kyoto Protocol as a source of exogenous variation in carbon prices.

[5]The effects of the Kyoto Protocol on trade are also examined by Aichele & Felbermayr (2012, 2013).

They find evidence of a significant PHE: For a bilateral pair with a committed importer and noncommitted exporter, the Kyoto Protocol led to a 5% increase in imports.

These and other recent studies find that environmental regulations have a significant effect on trade flows. A natural question is how large these effects are relative to traditional determinants of comparative advantage. This is addressed by Broner et al. (2015), who build on an empirical approach developed by Romalis (2004) to examine the relative importance of environmental regulation in determining success in exporting to the United States in 85 industries from 101 countries for the year 2005. Using a novel IV approach they find a strong PHE. The impact of weak environmental regulations on comparative advantage is similar in magnitude to the effects of physical and human capital.

3.2. Does International Trade Affect Environmental Outcomes?

Although the work discussed in the previous section has found evidence that environmental policy affects trade flows, there is still little support for the PHH. From the perspective of trade being driven by comparative advantage, this suggests that trade has a relatively small effect on the environment.

In early work, Antweiler et al. (2001) use an international panel of data on air quality in cities to estimate the scale, technique, and composition effects of openness to trade on sulfur dioxide concentrations. Consistent with the comparative advantage–based theory discussed at the beginning of Section 3, they find that scale effects raise pollution, income-induced technique effects lower pollution, and the sign of composition effects varies across countries. These composition effects were, however, very small, whereas the measured technique effects were large. Somewhat surprisingly, they find that openness tends to raise sulfur dioxide concentrations in rich countries (which have more stringent environmental policy) and lower concentrations in poor countries (which have weaker environmental policy). This is consistent with the view that traditional determinants of comparative advantage (e.g., factor endowments, technology differences) have a much larger effect than environmental policy on trade flows. These conclusions echo those of Grossman & Krueger (1993). Subsequent empirical work by Cole & Elliott (2003) and others (see, e.g., Managi et al. 2009) also finds small composition effects, again suggesting a small impact of trade on the environment.[6]

An alternative to estimation is to measure composition effects directly. This approach is used by Levinson (2009) to study the large decrease in pollution emitted by the US manufacturing sector between 1987 and 2001. Levinson adopts the industry decomposition given by Equation 2 and uses it to calculate the scale, composition, and technique effects directly. The scale and composition effects are calculated from observed data; the technique effect is calculated as a residual by subtracting the scale and composition effects from observed changes in emissions. Levinson's estimates suggest that the reduction is largely a product of the technique effect. The composition effect accounts for only approximately 12% of the reduction.[7] Levinson then asks if the size of the composition effect can be explained by trade. He employs another accounting identity to calculate the pollution displaced by US imports and compares it with the magnitude of the observed

[6] One potential concern is that trade and income are endogenous, so that estimates could be biased. Frankel & Rose (2005), Chintrakarn & Millimet (2006), and McAusland & Millimet (2013) adapt an IV approach to study the effects of trade on the environment in various settings. They all find that the effects of trade are relatively small.

[7] Levinson (2015) addresses potential errors created by the residual method by instead calculating the technique effect directly using data from the National Emissions Inventory maintained by the US Environmental Protection Agency (EPA). His estimates again show that technique effects account for the majority of the cleanup of the US manufacturing sector.

composition effect. He finds that trade plays almost no role: It can account for only approximately 4% of the overall reduction in emissions.

Subsequent research has found similar patterns in other jurisdictions and time periods. Shapiro & Walker (2015) use product-level data to study the fall in emissions from US manufacturing for the period 1990–2015. They find very large technique effects and very small composition effects. Grether et al. (2009) decompose emissions of sulfur dioxide for 62 countries between 1990 and 2000 and find that, despite a 10% increase in the scale of production, emissions have fallen by close to 10%. This is primarily due to a significant negative technique effect: Compositional changes are less than one-fifth the size of the technique effect. Similarly, Brunel (2016) shows that technique effects are the primary driver of reductions in EU emissions over the 1995–2008 period. Interestingly, Brunel documents a small positive composition effect in the European Union, meaning production was shifting toward relatively dirty industries during her period of study. However, like Levinson (2009), Brunel finds that trade explains little of this compositional change.

One exception to this pattern is the work of Barrows & Ollivier (2016), who study carbon emissions in India. They find that composition effects are similar in magnitude to technique effects. This raises the possibility that explanations for evolution of emissions over time may differ between developing and developed countries. Using a very different research strategy, Bombardini & Li (2016) try to directly identify the effect of trade on pollution in China by adopting an approach similar to that used by Autor et al. (2013) to look at the effects of trade on labor markets. Bombardini & Li (2016) exploit regional variation in comparative advantage within China to identify the effects of trade on regional pollution levels and the subsequent effects of these pollution changes on infant mortality rates. They find robust evidence that increased regional participation in export markets increases infant mortality rates due to increased ambient pollution concentrations. Their results are suggestive of potentially large effects of trade on pollution driven by comparative advantage.

The findings of small composition effects and large technique effects, especially for the United States and European Union but also in large cross-country studies, have led many to conclude that international trade has had little to no effect on pollution emissions via composition effects. Emissions from US manufacturing have fallen quite dramatically over the past 25 years, but the evidence from aggregate industry-level decompositions points to large drops in emission intensities rather than changes in the composition of production as the explanation. This is puzzling, particularly in light of the work discussed in the previous section, which indicates that environmental regulations are a significant determinant of trade flows. This work also raises an interesting related question: Why have emission intensities fallen so dramatically?

One possible explanation is that the effect of environmental policy on trade patterns is not economically significant and that emission intensities have fallen because of more stringent environmental policy. Shapiro & Walker (2015), for example, find that a doubling of the implied pollution tax from 1990 to 2008 would be sufficient to explain the observed fall in emission intensities, and they argue that this is not unreasonable in light of the increase in regulatory stringency in the United States during this time.

A second potential explanation for the puzzle is that the measurement of technique effects is in many cases flawed. A failure to account for important within-industry and within-firm adjustments created by trade may be driving emission intensities downward. To see how this could be true, consider a simple example. Take a finely defined industry that uses no domestic intermediates, and suppose it experiences a reduction in pollution over some time period. Over this same time period, imports rise by 20,000 units, and exports rise by 10,000 units. Our data also give us emissions per unit of output prior to the change, output pre- and post-change, and aggregate emissions pre- and post-change. Multiplying output changes by emission coefficients gives us the scale effect, there is no measured composition effect, and the impact of trade in this exercise would simply be the change in net exports times emission intensities. Taking all these changes into account and subtracting from the actual changes in emissions yield a remainder we would call the technique effect. It is tempting and common to conclude that trade's contribution to the emission reduction is simply the reduction in net imports of 10,000 units times the preexisting emissions per unit output.

As we show in the following sections, if the change in net imports were due to trade liberalization, then much could be wrong with this calculation. For example, if the reduction in production for domestic consumption caused by imports came from the dirtiest firms, then the reduction in emissions attributed to trade should be higher because imports are displacing the dirtiest of home production. If the increase in domestic production for exports was produced by the cleanest firms, then the increase in emissions attributed to trade should be smaller because exports are encouraging expansion by the cleanest firms. Standard industry-level decomposition methods will miss a fall in emissions resulting from a trade-induced reallocation of output across clean and dirty firms. As a result, these methods can also underestimate the effects of trade and misclassify them as technique effects.

4. TRADE, FIRM HETEROGENEITY, AND THE ENVIRONMENT

Recently, researchers have used plant- and firm-level data sets to investigate how international trade may affect the environment. We develop a simple theoretical model of a representative firm inspired by Melitz (2003) to highlight key mechanisms[8] and then discuss empirical findings (for model details and derivations, see the **Supplemental Appendix**).

4.1. Technology and Costs

We consider a firm producing a differentiated good in a monopolistically competitive industry. Preferences are the usual constant elasticity of substitution (CES) Dixit-Stiglitz specification. Firms must pay a fixed entry cost to obtain a productivity draw and a fixed cost to engage in any production. Firms have to pay an additional fixed cost F_e (in terms of labor) to export and incur variable shipping costs for delivery to foreign markets.

The firm produces final goods by using a Leontief technology to assemble a continuum of intermediates $x(j)$, $j \in [0, 1]$. Each intermediate good j is produced with a CES production technology using clean and dirty inputs:

$$x(L, D; j) = \gamma [a_j^{1-\delta} L^\delta + b_j^{1-\delta} D^\delta]^{1/\delta}, \tag{19.}$$

where $\delta < 1$, $a_j > 0$, and $b_j > 0$. L is a clean input (such as labor) with factor price w, D is a dirty input available at price r, and $\gamma > 0$ is a productivity parameter. We choose the index j to be increasing in the dirty input intensity of intermediates (b_j/a_j is increasing in j).

Pollution emissions are directly proportional to the use of the dirty input, but firms can pay a fixed cost A to invest in abatement technology; this reduces emissions generated by the dirty

[8] Several recent papers have used the Melitz model to study the interaction between trade and the environment. The reader is referred, for example, to the work of Yokoo (2009), Cui (2014), Cole et al. (2014), Kreickemeier & Richter (2014), Ravetti & Baldwin (2014), Forslid et al. (2015), Konishi & Tarui (2015), and Barrows & Ollivier (2016).

input. Pollution emissions are given by

$$z = g(A)D, \tag{20.}$$

where $g(A)$ is decreasing in A and $0 \le g(A) \le 1$.[9]

For simplicity, any investment A in abatement reduces the emission intensity of the dirty input for all intermediates produced by the firm at home. This would be the case if, for example, abatement was capturing pollution particles/effluent from a common combustion/discharge chamber, or if an investment in knowledge was making all processes cleaner.

The government regulates pollution with an emission tax τ, so the full price τ_D to firms for the dirty input is given by

$$\tau_D = r + \tau g(A). \tag{21.}$$

Firms decide to produce intermediates in-house or offshore them to a foreign producer[10] by comparing relative costs. To generate a motive for offshoring, we assume that the pollution charges at home are relatively high in the sense that

$$\frac{\tau_D}{w} > \frac{\tau_D^*}{w^*}, \tag{22.}$$

where an asterisk denotes foreign variables.

The domestic firm compares the costs of producing intermediate good j in-house and abroad and will offshore intermediates for which

$$[1+\kappa]c^*(w^*, \tau_D^*; j) < c(w, \tau_D; j), \tag{23.}$$

where c is the cost function corresponding to Equation 19 and where $\kappa > 0$ is a parameter that reflects the variable costs of outsourcing. The condition in Equation 23 is equivalent to

$$\frac{w}{w^*} > [1+\kappa]T_j, \tag{24.}$$

where

$$T_j \equiv \frac{\gamma}{\gamma^*}\left[\frac{1+\frac{b_j}{a_j}[\tau_D^*/w^*]^{1-\sigma}}{1+\frac{b_j}{a_j}[\tau_D/w]^{1-\sigma}}\right]^{1/[1-\sigma]} \tag{25.}$$

and $\sigma = 1/[1-\delta]$. The curve $T(j)$ measures the role of environmental policy and emission intensities in determining the cost of foreign production relative to home production. It is downward sloping because Equation 22 holds: Pollution charges are relatively high at home. For the given abatement, the outsourcing decision based on Equation 24 is illustrated in **Figure 1**. Intermediates on the interval $(j_0, 1]$ are offshored because of the relatively stringent environmental policy at home.

Finally, the firm minimizes costs by choosing abatement (taking into account offshoring). Higher emission charges increase the incentive to abate, as do higher output levels. Because more productive firms use less of the dirty input to produce output, they invest less in abatement; however, if their output were to increase more than in proportion to an increase in productivity, then abatement would rise. When the firm produces less in-house, it has a smaller incentive to abate.[11]

[9]This specification has been influenced by Bustos (2011), who studies the effects of trade on technology upgrading, and by Girma et al. (2008), Batrakova & Davies (2012), and Forslid et al. (2015), who also specify an abatement technology that can be upgraded with a fixed-cost investment.

[10]For simplicity, we do not consider domestic outsourcing.

[11]The role of outsourcing as a substitute for abatement is highlighted by Cole et al. (2014).

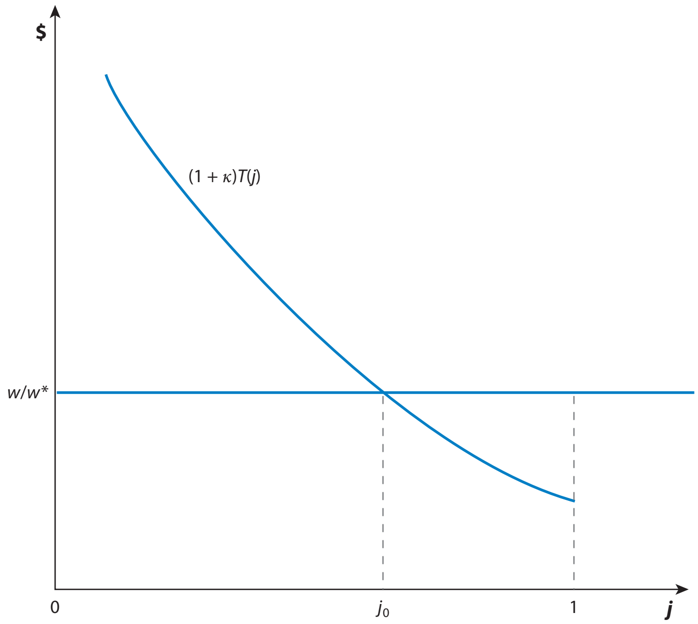

Figure 1

The decision to offshore. The downward-sloping curve $(1+\kappa)T(j)$ depicts the relative cost of obtaining intermediate good j from foreign and domestic sources given cross-country differences in environmental policy. The horizontal line w/w^* depicts the relative price of clean inputs across countries. Intermediates on the interval $(j_0, 1]$ are outsourced.

4.2. Production Choices and the Decision to Export

Firms can produce for only the domestic market or choose to export but suffer additional costs. If firms sell in the domestic market, then they earn profits π^d. If they sell in both markets, then they earn π^e, which accounts for the costs of exporting. The incremental profits from exporting are given by

$$\tilde{\pi}^e = \pi^e - \pi^d. \tag{26}$$

It is straightforward to show that all profit functions are increasing in productivity, and thus, in **Figure 2**, we depict them in a typical Melitz (2003) diagram. First consider the profit functions labeled $\tilde{\pi}_0^e$ and π_0^d.

Low-productivity firms with $\gamma < \gamma_0^d$ are not viable because $\pi_0^d < 0$. They must exit the industry. High-productivity firms with $\gamma \geq \gamma_0^e$ can cover the fixed costs of exporting to earn positive incremental profits ($\tilde{\pi}_0^e > 0$). Firms with intermediate productivity serve only the domestic market. The upper envelope π_0^e represents profits to firms conditional on exporting decisions. More productive firms are larger, more profitable, and more likely to be exporters.

4.3. The Effects of Trade Liberalization

Suppose there is bilateral trade liberalization between the domestic and foreign markets. Environmental policy is fixed. There are adjustments across firms within a given industry, as well as adjustments within firms in a given industry that may involve offshoring and abatement decisions. We discuss these using both our decomposition and specific examples from the theory sketched above.

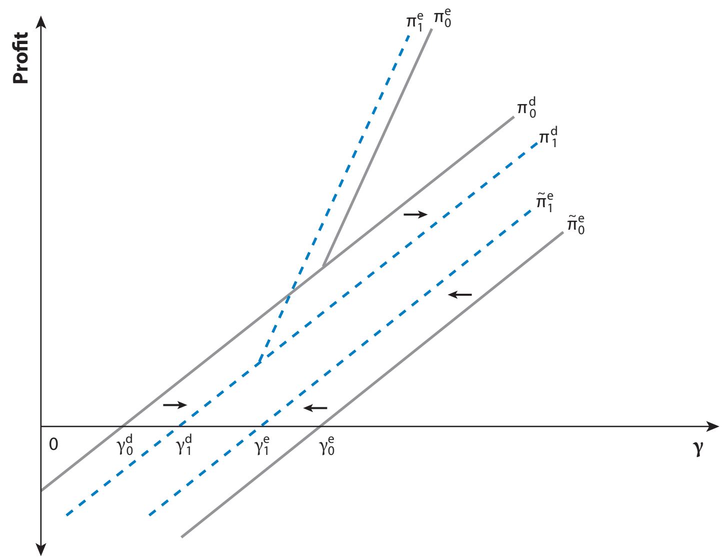

Figure 2

The decision to export and the effects of trade liberalization. The lines denoted π^d, π^e, and $\tilde{\pi}^e$ depict the profits from selling in domestic markets, the profits from selling in both domestic and export markets, and the incremental profits from exporting, respectively. The effects of trade liberalization are given by the move from solid to dashed lines.

4.3.1. Across-firm adjustments to trade liberalization. To focus on adjustments across firms, assume that firm-level emission intensities are constant (abatement is fixed, outsourcing is not possible, and trade does not affect the price of dirty inputs) and that there is one industry. In this case, the effects of trade liberalization are very similar to those found in Melitz (2003). The reduction in trade costs changes the profitability of serving the domestic and foreign markets, alters the set of active firms, and redistributes outputs across firms.

These effects are illustrated by the move from solid to dashed lines depicted in **Figure 2**. Increased foreign competition lowers the profits from the domestic market and raises the productivity cutoff for domestic firms from γ_0^d to γ_1^d. Low-productivity firms exit the industry. Exporters face more competition in the local market but have more lucrative export opportunities because trade costs have fallen. This raises the incremental profits from exporting, and domestic firms with productivity between γ_1^e and γ_0^e enter the export market. The most productive firms experience profit increases; the least productive experience profit decreases.

Because emission intensities are constant, the effect of this trade liberalization on emissions is given by a drastically simplified version of Equation 18:

$$\mathrm{d}Z = \mathrm{d}S + \int_0^{n_0} e(n)\mathrm{d}\varphi(n)\mathrm{d}n + \varphi(n)\left[e(n) - E\right]\mathrm{d}n_0, \qquad 27.$$

where n_0 is the marginal firm. We have written the decomposition in terms of absolute changes to highlight the key new forces introduced by the Melitz framework—the potential for market share reallocations (the second term in Equation 27) and selection effects (the third term in Equation 27) to affect emissions.

In the model developed above, emission intensities are decreasing in firm productivity even without an active abatement decision.[12] Because trade liberalization redistributes market share toward firms at the upper end of the productivity spectrum, the second term is therefore negative, driving down emissions. Because low-productivity, dirty firms exit when liberalization occurs, the selection effect is also negative. The only element pushing up emissions is the scale effect. The net impact on pollution comes from weighing the positive scale effect against the emission-reducing effects created by industry rationalization.

The possibility of reducing emissions solely from trade-inspired industry rationalization generates a testable hypothesis, which we refer to as the pollution reduction by rationalization (PRR) hypothesis. When the PRR hypothesis is true, trade liberalization lowers industry emissions because a mix of exit, entry, and market share reallocations overwhelms the positive scale effects created by new export opportunities.

Evaluating the PRR hypothesis is an empirical task, but in an interesting recent paper, Kreickemeier & Richter (2014) provide some insight. They develop a Melitz-style setting much like the one we developed above and consider the environmental implications of a unilateral trade liberalization by one of two identical countries. The liberalization is marginal, there is no abatement, but emission intensities vary with the productivity of the firm in a flexible manner. They highlight the interaction between the scale and industry rationalization effects and find that if emission intensities fall more than in proportion to productivity, then trade reduces aggregate emissions and the PRR hypothesis holds.

If the PRR hypothesis is true, this then presents some interesting and confusing possibilities for research. For example, an industry's average emission intensity will fall with trade liberalization, and this fall may be misinterpreted as a technique effect. Trade liberalization does not cause any firm to become cleaner, even though we observe relatively clean exporting firms, relatively dirty domestic producers, and a fall in industry emissions with trade. Trade causes emissions to fall, but only because of trade-induced rationalization.

4.3.2. Within-firm adjustments to trade liberalization. In this section, we consider within-firm adjustments. For ease of exposition, we focus on a single firm and suppress both n and i. This firm's total emissions z are thus the product of the scale of its value added v times its emission intensity e. That is, we find that $z = ve$ and hence

$$\widehat{z} = \widehat{v} + \widehat{e}. \qquad 28)$$

We have already discussed how trade liberalization affects a firm's scale, so we focus here on changes within firms affecting emission intensities. In the model developed above, we made three simplifications relative to the more general decomposition. We assumed the firm assembles intermediates in fixed proportions, and thus there is no scope for emissions to change from a reorganization of production. We assumed all outsourcing is to foreign firms (offshoring), so if a firm's emissions fall from outsourcing production, they cannot be offset by increases elsewhere at home. Finally, we assumed task production is specialized so that a change in outsourcing is always an adjustment on the extensive margin via j_0.

[12]This follows because we assumed that variable production is homothetic and that increases in productivity are equivalent to neutral technical progress. Similar results are obtained by Cui et al. (2012). If productivity variations across firms are non-neutral and more productive firms use a more energy-intensive technology, then firms may be both more productive and dirtier.

With these simplifications, the decomposition of our firm's change in emission intensity reduces to

$$\widehat{e} = \int_0^{j_0} \theta_j \widehat{e}_j dj + [\theta_{j_0} - \varphi_{j_0}]\mathrm{d}_{j_0}. \tag{29}$$

The firm's emission intensity is affected by changes in the emission intensities of various tasks and by offshoring. We discuss these effects in the following sections.

4.3.3. Emission intensities and abatement. A firm's emission intensities are affected by trade in various ways, but a key question is whether firms' endogenous abatement choices reinforce the rationalization and selection effects that operate at the industry level. At first blush, it may seem to be so. In our simple model (treating outsourcing as given), abatement investment rises with firm productivity as long as the output rises more than in proportion to productivity. In a CES framework where markups are fixed, a fall in costs lowers prices proportionately and raises output relatively more. Therefore, output rises more than proportionately to productivity changes, and more productive firms abate more. The emission-reducing forces we found in the PRR effect appear to be strengthened by abatement. There are, however, several complications that alter this simple story and open up an entirely new possibility.

First, this result is sensitive to the demand structure. Cao et al. (2016) use a preference structure developed by Melitz & Ottaviano (2008) and show that more productive firms may invest less in abatement. Because there is no longer constant elasticity demand, highly productive firms operate on the relatively inelastic portion of their demand curves. Markups are higher, and output responds less than proportionately to productivity. The simple channel linking higher firm productivity to greater abatement and lower emission intensity is now broken.

A second complication is that investment in abatement need not lower emission intensities. An increase in abatement has two effects on emission intensity. There is the direct effect that emissions per unit of the dirty input fall. However, there is also a classic rebound effect (see Gillingham et al. 2016). When abatement lowers emissions per unit of the dirty input, pollution charges per unit of the dirty input fall, reducing the cost to firms of using the dirty input. This encourages substitution toward the dirty input. If the elasticity of substitution is not too large ($\sigma \leq 1$ is sufficient), an increase in abatement lowers emission intensities. However, for sufficiently large σ (that is, if the dirty and clean inputs are very close substitutes), an increase in abatement expenditure can raise emission intensity (for details, see the **Supplemental Appendix**). Therefore, the response of emission intensities to abatement can depend quite delicately on technology.

A third complication arises because trade creates both winners and losers. Suppose the rebound effect is weak, so that abatement reduces emission intensities. Exporters then become cleaner because they increase their abatement as their output expands, but domestic producers who survive see their output fall and reduce their abatement. Their emission intensities will rise because of trade. Although industry output does shift away from these firms, there is no reason to believe this generates a net reduction in pollution for the industry as a whole.

Forslid et al. (2015) study the interaction between these effects in a two-country symmetric world within a Melitz-type framework. Trade liberalization induces exporting firms to increase abatement and become cleaner, whereas nonexporting firms invest less in abatement and become dirtier. The net effect is to reduce emissions in each country with a symmetric trade liberalization. The question of which features of the model ensure this result remains unresolved. For example, consider a simple case where technologies are either clean or dirty and there is a fixed cost for using the clean technology. Suppose that, initially, all firms are large enough to pay the fixed cost.

If trade liberalization pushes some firms below the threshold at which the clean technology is cost-effective, then those firms would revert to the dirty technology and emissions would rise.

Allowing for endogenous abatement complicates considerably the set of adjustments we may expect and, at worst, introduces a decidedly negative but novel potential outcome from trade liberalization. Out of these complications come two implications for empirical work.

The first is a new hypothesis. With active abatement, trade can indeed cause some firms to become cleaner, and this induced technique effect may lower industry emissions. However, trade can also cause some firms to become dirtier. Industries hit hard by trade liberalization may feature a set of dirty and distressed domestic firms who have forsaken abatement expenditures. Even exporters who face increased competition in local markets may lower their abatement expenditures if the liberalization is severe enough (see Forslid et al. 2015). Therefore, endogenous abatement choices introduce a new possibility we refer to as the distressed and dirty industry (DDI) hypothesis. It holds when trade liberalization increases industry emissions because of reductions in pollution control or abatement expenditures by firms that downsize as a result of trade.

A second implication comes from considering a simple difference-in-differences methodology to measure trade's causal effect on firm-level emission intensities. Suppose we compare the emission intensities of exporters to those of domestic firms both pre- and postliberalization. Domestic firms are the control group. If we were to conduct such an evaluation on the stripped-down model of our previous section with no abatement and fixed emission intensities, then this method would reveal the correct answer—trade has no causal effect on firm-level emission intensities—because trade merely reallocates output and selects firms. However, if we conduct the same experiment with endogenous abatement, we may find a very large effect—for the wrong reasons. Exporters become cleaner relative to domestic firms because trade makes them cleaner and their control group dirtier.

4.3.4. Offshoring dirty inputs. The interaction between trade liberalization, offshoring, and aggregate pollution emissions is complex and currently not well understood. A trade liberalization that reduces the cost of offshoring shifts down the $T(j)$ schedule in **Figure 1** and increases the range of intermediates offshored. Because the marginal outsourced intermediates are the dirtiest, this will reduce firm-level emission intensities. This effect is captured by the final term in our decomposition above (Equation 29). This is a potentially important channel because domestic firms become cleaner not because they have reduced the emission intensity of their activities but because they have shifted the dirtiest parts of their production out of the country. We refer to this as the pollution offshoring hypothesis (POH). It is reminiscent of the PHH, but in that hypothesis the focus is typically on dirty final good producers moving production or plants to countries with weak environmental policy. The POH is more subtle in that it leads to fragmentation of production in countries with stringent environmental policy—only the dirtiest parts of the production process are shifted abroad.

The POH interacts with abatement in interesting ways. Abatement investments and offshoring dirty intermediates are substitute channels for firms to respond to more stringent environmental regulation. Both lead to lower measured emission intensities in firms subject to regulation; however, whereas abatement leads to real reductions in pollution, offshoring shifts the incidence of pollution elsewhere.[13] If the POH holds, one possibility is that trade liberalization may reduce the

[13] Cole et al. (2014) develop a simple model in which firms can pay either an abatement cost or a fixed cost to offshore all of their polluting activity. The model predicts that large firms will outsource polluting activity. Using Japanese data, they find

incentives for pollution abatement investments because the option of fragmentation has become more cost-effective.

4.4. Empirical Evidence

The literature providing microlevel evidence on international trade and the environment is still in its infancy. It can be usefully divided into two branches: studies examining the relationship between export status and pollution emitted by firms and plants and studies examining the effects of trade liberalization.

4.4.1. Exporting and the environment. To date, researchers have focused on export status as a key determinant of pollution emissions. This can largely be attributed to the work of Holladay (2016), who examines how export status and import competition affected the level of toxic pollution emitted by US manufacturing plants over the period 1990–2006. Holladay employs emissions data reported in the US EPA's Risk-Screening Environmental Indicators database and data on plant characteristics from the National Establishment Time-Series (NETS) database. This yields an unbalanced panel data set with information on toxic pollution emissions and several plant characteristics, including an indicator of the plant's export status in the last year it was observed.

Holladay estimates several versions of the following:

$$\ln Z_{ijt} = \alpha + \beta \, 1[\text{Exporter}]_i + \pi W_{ijt} + \gamma_j + \delta_t + \epsilon_{ijt}, \qquad 30.$$

where Z_{ijt} is pollution emitted by plant i in industry j at time t, $1[\text{Exporter}]_i$ is a time-invariant indicator of plant i's export status, W_{ijt} are additional controls, γ_j and δ_t are industry and year fixed effects, and ϵ_{ijt} is the error term. The coefficient of interest, β, measures the average difference in log pollution emissions between exporters and nonexporters.

Holladay finds that, conditional on log sales and industry, state, and year fixed effects, exporters emit 10% less pollution than nonexporters. Holladay suggests that productivity differences across these groups of plants are driving the result, although he does not provide productivity evidence directly.

There are several reasons to be skeptical of Holladay's finding. When he estimates Equation 30 for 20 subsamples, each corresponding to one of the 20 major industrial groups (two-digit Standard Industrial Classification categories) that make up US manufacturing, the estimates are negative and statistically significant for only 7 of 20 industry groups. Exporters are significantly more pollution intensive than nonexporters in 4 of the 20 industry groups.

One potential explanation for this fragility is that these estimates are capturing a different mechanism. One possibility is offshoring. If offshoring requires the payment of fixed costs, then large productive plants that are more likely to export may be more likely to offshore dirty intermediate inputs. This is a material possibility because Harrison & McMillan (2011) report increased use of offshoring by US manufacturing plants during this time.

Another potential explanation is that this fragility reflects problems arising from endogeneity. Export status is not an exogenous characteristic of plants (e.g., Lileeva & Trefler 2010), which means the negative relationship between pollution emissions and exporting may be capturing a more primitive characteristic that drives the decision to export.

a positive relation between their measure of the stringency of environmental regulations and outsourcing and between firm size and outsourcing. However, their data do not allow them to determine whether the outsourced activities are relatively pollution intensive.

These concerns do not mean the negative relationship documented by Holladay is uninformative, and we can exploit theory to understand the broader implications of this result.

Suppose, for example, that the negative correlation between export status and pollution emissions is driven solely by productivity differences, as suggested by Holladay. This case resembles the scenario discussed in Section 4.4.1; only the most productive plants are able to afford the fixed cost required to export, and these plants pollute less because they require fewer inputs per unit of output. The effects of trade liberalization are then immediate. Trade will redistribute market share to the most productive plants that have lower emission intensities and lead to the exit of dirty plants from the industry, meaning that trade liberalization will reduce emissions through industry rationalization. Hence, Holladay's findings can be interpreted as suggestive evidence of a necessary condition for the PRR hypothesis.

Holladay also examines the effects of import competition on pollution emissions and plant entry and exit. He does so by adopting a version of the estimating Equation 30 that includes measures of import competition. These estimates indicate that plants in import-competing industries emit more pollution on average, and import-competing industries are characterized by less entry and more exit than non-import-competing industries. If we again assume that a negative relationship between emissions and productivity exists and allow for endogenous abatement decisions, then we can interpret these results with the aid of our theory. In this case, Holladay's estimates suggest that import competition increases emissions from the average plant even though the dirtiest plants are exiting. This can be attributed to the loss in market share caused by foreign competition, meaning that the remaining plants are getting dirtier as a result of decreased pollution abatement. Thus, the effects of import competition documented by Holladay are suggestive of the mechanism underlying the DDI hypothesis.

Altogether, the results presented by Holladay suggest that exporting and import competition have significant effects on the pollution emitted by US manufacturing plants. These findings are, however, subject to significant caveats, some of which have been addressed in subsequent research.

One example of this is the work of Cui et al. (2016), who also examine the effects of exporting on the pollution emitted by US manufacturing plants. Although they use a similar research design and also employ data from the NETS, the analysis presented by Cui et al. differs from that of Holladay along two key dimensions. First, Cui et al. obtain pollution data from the US National Emissions Inventory for the years 2002, 2005, and 2008, which allows them to examine four common pollutants: sulfur dioxide, carbon monoxide, ozone, and particulate matter. Second, they construct a productivity estimate for each plant. Although the lack of information on capital stocks and other inputs in the NETS data prevents them from constructing a standard total factor productivity measure, this approach allows them to provide some preliminary evidence of the relationship between export status, plant productivity, and pollution emissions.

Cui et al. find that exporters are less pollution intensive than nonexporters for each of the four pollutants they study.[14] However, these effects are estimated conditional on productivity, which indicates that the differences are not simply due to productivity differences. Instead, some other mechanism is likely driving the relationship. One potential channel is plant abatement; recall from our earlier discussion that, for a given productivity level, plant abatement decisions will depend on market size. In this case, the effects of exporting documented by Cui et al. can be interpreted as capturing the effects of market size differences. Indeed, their estimates are consistent with this conjecture; once they control for plant employment, a plant characteristic that is determined in

[14]These effects are substantially larger than the effects documented by Holladay (2016). Cui et al. (2016) find that exporting is associated with a 26.2–29.5% reduction in emission intensity depending on the pollutant.

part by market size, the effect of exporting is no longer statistically significant. Of course, given that Cui et al. share a similar empirical approach to Holladay, these findings are subject to the same identification concerns expressed above.

The relationship between export status and pollution emissions is also examined by Forslid et al. (2015), who study the relationship between export status and the emissions of carbon dioxide, sulfur dioxide, and nitrous oxide from Swedish manufacturing firms over the period 2000–2011. Forslid et al. use data from a statistical agency (Statistics Sweden), which allows them to address key measurement issues. First and foremost, the data contain information on export sales by year, meaning that Forslid et al. are able to accurately track export participation over time. In addition, the data contain detailed information on value added and input use, which allows Forslid et al. to accurately measure the productive activity at each plant and calculate firm productivity using the approach developed by Levinsohn & Petrin (2003).

To examine the relationship between exporting and pollution emissions, Forslid et al. exploit the fact that they observe firm exports in all years and adopt a research design that has been used elsewhere in the trade literature. This design treats firms that do not export as a counterfactual for those that do, using a logic similar to difference-in-differences; the average pollution emissions from exporting firms (the treatment group) are compared to the average pollution emissions from nonexporters (the control group) before and after an exporting episode begins.

Using this approach, Forslid et al. show that export status and emission intensity are negatively related for each of the pollutants they study. Exporters are 11.4% less carbon dioxide intensive, 18.7% less nitrous oxide intensive, and 26.7% less sulfur dioxide intensive than nonexporters conditional on firm productivity and industry and year fixed effects. However, dividing the sample into energy-intensive and non-energy-intensive industries shows no effects in the energy-intensive industries, which is worrisome.

Forslid et al. also provide evidence that exporting affects firm abatement decisions. They show that exporting is associated with a 62–73% increase in investment in abatement activities for firms in non-energy-intensive industries.

Although Forslid et al. cite their abatement findings as evidence in support of their main finding of a large negative relationship between emission intensity and exporting, our theory suggests another potential explanation. As we discussed in Section 4.3.3, when abatement is endogenous, trade can make exporters cleaner and nonexporters dirtier. In the present setting, this means the control group is contaminated, and Forslid et al.'s estimates overstate the effects of exporting.

4.4.2. The effects of trade liberalization. A second set of studies has begun to move beyond simply documenting the relationship between export status and pollution emissions to examine the effects of trade liberalization on firms. This branch of the literature has exploited tariff changes that occur during episodes of trade liberalization as a source of identifying variation.

The use of tariffs as a source of identifying variation is typified by the work of Martin (2012), who examines the effects of India's trade liberalization in 1991 on the greenhouse gases emitted by the Indian manufacturing sector. She constructs a firm-level panel data set from India's Annual Survey of Industries that includes detailed information on the output and inputs of manufacturing firms over the period 1985–2004. Martin supplements these data with yearly estimates of firm greenhouse gas emissions that are constructed from the observed energy use of firms using time-invariant emissions factors.

Martin begins her analysis by adopting a decomposition similar to the one developed in Section 2 to divide aggregate changes in greenhouse gas emissions and fuel use into across-industry, within-industry, and within-firm changes. She then uses industry tariff changes to examine how trade liberalization affects the changes she observes. Her underlying research design is similar in spirit to difference-in-differences; the average change in outcome for industries that are affected by a change in trade policy (the treatment group) is compared to the average change in outcome for industries that are unaffected by the policy (the control group) before and after the policy change occurs.

Given that India's trade liberalization occurred as part of a larger set of economic reforms, Martin also examines the effects of changes in regulations on foreign direct investment and industrial licensing. She finds that decreases in intermediate input tariffs increased the energy efficiency of affected firms by 23%, while the relaxation of industrial licensing requirements shifted market share to more fuel-efficient firms. There is a possibility, however, that these effects capture differences in outcomes due to both the effects of trade policy changes and the effects of other political economy–driven policy changes that occurred during India's liberalization. Overall, this is a potentially very useful approach using elements from an emissions decomposition as inputs in a regression framework.

A related approach is adopted by Cherniwchan (2017), who examines the effects of trade liberalization between the United States and Mexico following NAFTA on the pollution emitted by US manufacturing plants over the period 1991–1998. Cherniwchan constructs a plant-level panel data set using data on the emissions of particulate matter and sulfur dioxide, constructed from the Toxic Release Inventory maintained by the EPA, and data on plant characteristics from the NETS database.

To identify NAFTA's effects on plant pollution emissions, Cherniwchan develops a triple difference research design that exploits two sources of variation in the costs of trade: tariffs and the trade costs created by geography. His approach is based on the idea that geographic variation in trade costs will cause the effects of a tariff reduction to differ across states: A tariff reduction will have little effect on plants that are located in states where the costs of moving goods to and from foreign markets are very high but will have a large effect on plants in states where these costs are low. This variation means that a change in trade policy will only affect (or treat) the subset of plants that are located in states with low geographic trade costs.

Using this approach, Cherniwchan shows that NAFTA led to substantial reductions in the emissions of particulate matter and sulfur dioxide from affected US manufacturing plants. These reductions are driven by two aspects of the liberalization: increased foreign market access and decreases in the cost of importing intermediate inputs. Cherniwchan finds that, for the average plant, increased foreign market access reduced particulate matter and sulfur dioxide emissions by approximately 1.7% annually, whereas reductions in the cost of importing inputs reduced sulfur dioxide emissions by close to 1.3% annually. Moreover, he finds that these reductions are primarily due to within-plant changes in the emission intensity of production. He also presents some evidence suggestive of the POH: He finds that the plant emission intensities are falling in part due to changes in access to relatively dirty intermediate inputs. Altogether, his estimates suggest that NAFTA played a large role in the cleanup of the US manufacturing sector during the 1990s.

5. CONCLUSION

We have reviewed recent evidence linking international trade and the environment and found significant advances in many areas. In particular, our focus on evidence brought by heterogeneous firm models of international trade revealed new and interesting possibilities, but there is still much work to be done.

We found considerable new and convincing evidence linking tighter environmental regulations to reduced net exports (or increased net imports) in polluting sectors. This new evidence is clear support for the PHE. Although there is certainly value added in examining how its strength varies across instruments and industries, we view its existence as a settled question. Despite this finding, there remains little evidence that trade liberalizations shift dirty good production to low-income or weak-regulation countries, as suggested by the PHH. How to reconcile strong PHEs with little evidence for the PHH remains an open question.

We found that there has been significant application of heterogeneous firm models to questions concerning trade and the environment. There is some evidence that exporters are cleaner than other firms, but this relationship varies not only in strength across industries but in direction as well. There is better, but limited, evidence that trade liberalizations lower firm and perhaps even industry emissions, although the exact mechanisms by which this occurs need further study. Much of this research is hamstrung by lack of data on capital stocks and value added, making productivity measurement either impossible or heroic; much of the research is also hampered by identification problems. Although some researchers have adopted clever strategies to identify the causal impact of trade, others have been less careful. Despite these limitations, this research agenda is extremely valuable and should be pursued vigorously. The set of studies is still very small, and there is still much in dispute; more importantly, researchers have yet to realize the full potential of this new approach to answer important questions—both new and old.

Two important and old empirical puzzles in the literature are the apparent weakness of trade liberalization in shifting the composition of national output toward dirty (or clean) products and the apparent strength of technique or technology effects to drive emissions downward. Weak composition effects have typically been explained by recourse to offsetting forces governing comparative advantage in dirty products; large technique effects have been attributed to some combination of surprisingly strong policy responses and technological progress.

A simple and attractive alternative explanation for both puzzles is provided by models with heterogeneous firms. The logic is simple: If much of dirty good trade is intra-industry, then the pollution consequences of trade will be felt most strongly within industries and not across them. Small composition effects arise because trade primarily causes specialization within dirty industries. With heterogeneous firms, specialization begets rationalization, and with only slight additional assumption, trade can drive industry-level emissions downward as the dirtiest firms exit and output is reallocated to the cleanest firms. This is not a necessary outcome, but it is a possible one, and we have named it the PRR hypothesis. If the PRR hypothesis holds in each and every industry, and if most trade in dirty products is intra-industry and not driven by the forces of comparative advantage, then changes in the overall dirtiness of any one country's production will be limited, but declines in emissions may well be large. Thus, both puzzles are solved.

This neat explanation, however, side-steps many complications. As we have shown, the DDI hypothesis would work against this outcome, if true, and the POH, if true, would mean that success in lowering emissions at home might come with failure abroad. Moreover, the simple explanation also assumes that within-industry differences in emission intensities are more important to outcomes than across-industry differences. Exploring the strength of these new hypotheses and their workings in a world where interindustry trade is driven by comparative advantage will greatly improve our understanding of the mechanisms by which international trade affects the environment.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We would like to thank an anonymous reviewer and the editorial team for comments and suggestions that improved the manuscript. We would also like to thank Fruzsina Mayer for excellent research assistance. Funding from the Economics and Environmental Policy Research Network (EEPRN) is gratefully acknowledged. M.S.T. also acknowledges support from the Canada Research Chairs Program and SSHRC. B.R.C. conducted part of this research while visiting Université Paris-Est Marne la Vallée.

LITERATURE CITED

Aichele R, Felbermayr G. 2012. Kyoto and the carbon footprint of nations. *J. Environ. Econ. Manag.* 63:336–54

Aichele R, Felbermayr G. 2013. Estimating the effects of Kyoto on bilateral trade flows using matching econometrics. *World Econ.* 36:303–30

Aichele R, Felbermayr G. 2015. Kyoto and carbon leakage: an empirical analysis of the carbon content of bilateral trade. *Rev. Econ. Stat.* 97:104–15

Antweiler W, Copeland BR, Taylor MS. 2001. Is free trade good for the environment? *Am. Econ. Rev.* 91:877–908

Autor DH, Dorn D, Hanson GH. 2013. The China syndrome: local labor market effects of import competition in the United States. *Am. Econ. Rev.* 6:2121–68

Barrows G, Ollivier H. 2016. *Emission intensity and firm dynamics: reallocation, product mix and technology in India*. Work. Pap. 245, Grantham Res. Inst. Clim. Change Environ., London Sch. Econ. Polit. Sci., London

Batrakova S, Davies RB. 2012. Is there an environmental benefit to being an exporter? Evidence from firm level data. *Rev. World Econ.* 148:449–74

Becker R, Henderson V. 2000. Effects of air quality regulations on polluting industries. *J. Polit. Econ.* 108:379–421

Bombardini M, Li B. 2016. *Trade, pollution and mortality in China*. NBER Work. Pap. 22804

Brander JA, Taylor MS. 1998. Open access renewable resources: trade and trade policy in a two-country model. *J. Int. Econ.* 44:181–209

Broner F, Bustos P, Carvalho V. 2015. *Sources of comparative advantage in polluting industries*. Unpublished manuscript, Cent. Recer. Econ. Int., Barcelona, Spain

Brunel C. 2016. Pollution offshoring and emissions reductions in EU and US manufacturing. *Environ. Resour. Econ.* In press. **http://doi.org/10.1007/s10640-016-0035-1**

Brunnermeier SB, Levinson A. 2004. Examining the evidence on environmental regulations and industry location. *J. Environ. Dev.* 13:6–41

Bustos P. 2011. Trade liberalization, exports, and technology upgrading: evidence on the impact of MERCOSUR on Argentinian firms. *Am. Econ. Rev.* 101:304–40

Cao J, Qiu LD, Zhou M. 2016. Who invest more in advanced abatement technology: theory and evidence. *Can. J. Econ.* 49:637–62

Cherniwchan J. 2017. Trade liberalization and the environment: evidence from NAFTA and U.S. manufacturing. *J. Int. Econ.* In press. **http://doi.org/10.1016/j.jinteco.2017.01.005**

Cherniwchan J, Copeland BR, Taylor MS. 2013. *Environmental regulations and international competitiveness: a review of recent evidence*. Unpublished manuscript, Univ. Alberta, Edmonton, Can.

Chichilnisky G. 1994. Global environment and north-south trade. *Am. Econ. Rev.* 84:851–74

Chintrakarn P, Millimet DL. 2006. The environmental consequences of trade: evidence from subnational trade flows. *J. Environ. Econ. Manag.* 52:430–53

Cole MA, Elliott RJ. 2003. Determining the trade-environment composition effect: the role of capital, labor, and environmental regulations. *J. Environ. Econ. Manag.* 46:363–83

Cole MA, Elliott RJ, Okubo T. 2014. International environmental outsourcing. *Rev. World. Econ.* 150:639–64

Copeland BR. 2011. Trade and the environment. In *Palgrave Handbook of International Trade*, ed. D Bernhofen, R Falvey, D Greenaway, U Kreickemeier, pp. 423–96. Basingstoke, UK: Palgrave Macmillan

Copeland BR, Taylor MS. 1994. North-south trade and the environment. *Q. J. Econ.* 109:755–87

Copeland BR, Taylor MS. 1995. Trade and transboundary pollution. *Am. Econ. Rev.* 85:716–37

Copeland BR, Taylor MS. 2003. *Trade and the Environment: Theory and Evidence*. Princeton, NJ: Princeton Univ. Press

Copeland BR, Taylor MS. 2004. Trade, growth, and the environment. *J. Econ. Lit.* 42:7–71

Cristea A, Hummels D, Puzzello L, Avetisyan M. 2013. Trade and the greenhouse gas emissions from international freight transport. *J. Environ. Econ. Manag.* 65:153–73

Cui J. 2014. *Induced clean technology adoption and international trade with heterogeneous firms*. Unpublished manuscript, Wuhan Univ., Wuhan, China

Cui J, Lapan H, Moschini G. 2012. *Are exporters more environmentally friendly than non-exporters? Theory and evidence*. Work. Pap. 12022, Dep. Econ., Iowa State Univ., Ames, IA

Cui J, Lapan H, Moschini G. 2016. Productivity, export, and environmental performance: air pollutants in the United States. *Am. J. Agr. Econ.* 98:447–67

Davis LW, Kahn ME. 2010. International trade in used vehicles: the environmental consequences of NAFTA. *Am. Econ. J. Econ. Policy* 2:58–82

Forslid R, Okubo T, Ulltveit-Moe KH. 2015. *Why are firms that export cleaner? International trade, abatement and environmental emissions*. Unpublished manuscript, Dep. Econ., Stockholm Univ., Stockholm, Swed.

Foster L, Haltiwanger J, Syverson C. 2008. Reallocation, firm turnover, and efficiency: selection on productivity or profitability? *Am. Econ. Rev.* 98:394–425

Frankel JA, Rose AK. 2005. Is trade good or bad for the environment? Sorting out the causality. *Rev. Econ. Stat.* 87:85–91

Gillingham K, Rapson D, Wagner G. 2016. The rebound effect and energy efficiency policy. *Rev. Environ. Econ. Policy* 10:68–88

Girma S, Hanley A, Tintelnot F. 2008. *Exporting and the environment: a new look with microdata*. Work. Pap. 1423, Kiel Inst. World Econ., Kiel, Ger.

Grether JM, Mathys NA, de Melo J. 2009. Scale, technique and composition effects in manufacturing SO_2 emissions. *Environ. Resour. Econ.* 43:257–74

Grossman G, Krueger A. 1993. Environmental impacts of a North American Free Trade Agreement. In *The Mexico-U.S. Free Trade Agreement*, ed. PM Garber, pp. 13–56. Cambridge, MA: MIT Press

Hanna R. 2010. US environmental regulation and FDI: evidence from a panel of US-based multinational firms. *Am. Econ. J. Econ. Policy* 2:158–89

Harrison A, McMillan M. 2011. Offshoring jobs? Multinationals and US manufacturing employment. *Rev. Econ. Stat.* 93:857–75

Holladay JS. 2016. Exporters and the environment. *Can. J. Econ.* 49:147–72

Jaffe AB, Peterson SR, Portney PR, Stavins RN. 1995. Environmental regulation and the competitiveness of U.S. manufacturing: What does the evidence tell us? *J. Econ. Lit.* 33:132–63

Kellenberg D. 2009. An empirical investigation of the pollution haven effect with strategic environment and trade policy. *J. Int. Econ.* 78:242–55

Konishi Y, Tarui N. 2015. Emissions trading, firm heterogeneity, and intra-industry reallocations in the long run. *J. Assoc. Environ. Resour. Econ.* 2:1–42

Kreickemeier U, Richter P. 2014. Trade and the environment: the role of firm heterogeneity. *Rev. Int. Econ.* 22:209–25

Levinsohn J, Petrin A. 2003. Estimating production functions using inputs to control for unobservables. *Rev. Econ. Stud.* 70:317–41

Levinson A. 2009. Technology, international trade and pollution from US manufacturing. *Am. Econ. Rev.* 99:2177–92

Levinson A. 2010. Offshoring pollution: Is the US increasingly importing polluting goods? *Rev. Environ. Econ. Policy* 4:63–83

Levinson A. 2015. A direct estimate of the technique effect: changes in the pollution intensity of US manufacturing 1990–2008. *J. Assoc. Environ. Resour. Econ.* 2:43–56

Levinson A, Taylor MS. 2008. Unmasking the pollution haven effect. *Int. Econ. Rev.* 49:223–54

Lileeva A, Trefler D. 2010. Improved access to foreign markets raises plant-level productivity . . . for some plants. *Q. J. Econ.* 125:1051–99

Managi S, Hibiki A, Tsurumi T. 2009. Does trade openness improve environmental quality? *J. Environ. Econ. Manag.* 58:346–63

Martin LA. 2012. *Energy efficiency gains from trade: greenhouse gas emissions and India's manufacturing sector*. Unpublished manuscript, Dep. Agr. Res. Econ., Univ. Calif., Berkeley, CA

McAusland C. 2008. Trade, politics, and the environment: tailpipe versus smokestack. *J. Environ. Econ. Manag.* 55:52–71

McAusland C, Millimet DL. 2013. Do national borders matter? Intranational trade, international trade, and the environment. *J. Environ. Econ. Manag.* 65:411–37

Melitz M. 2003. The impact of trade on aggregate industry productivity and intra-industry reallocations. *Econometrica* 71:1695–725

Melitz M, Ottaviano GI. 2008. Market size, trade and productivity. *Rev. Econ. Stud.* 75:295–316

Ravetti C, Baldwin RE. 2014. *Emissions, exporters and heterogeneity: asymmetric trade policy and firms' selection*. Work. Pap. CTEI-2014-2, Cent. Trade Econ. Integr., Geneva, Switz.

Romalis J. 2004. Factor proportions and the structure of commodity trade. *Am. Econ. Rev.* 94:67–97

Shapiro JS. 2016. Trade costs, CO_2, and the environment. *Am. Econ. J. Econ. Policy* 8:220–54

Shapiro JS, Walker R. 2015. *Why is pollution from U.S. manufacturing declining? The roles of trade, regulation, productivity, and preferences.* NBER Work. Pap. 20879

Taylor MS. 2005. Unbundling the pollution haven hypothesis. *Adv. Econ. Anal. Policy* 3. **https://doi.org/10.2202/1538-0637.1408**

Yokoo H. 2009. *Heterogeneous firms, Porter hypothesis and trade*. Unpublished manuscript, Inst. Econ. Res., Kyoto Univ., Kyoto, Jpn.