Title
Essays on Technology and the Environment from an International Perspective

Author
Barrows, Geoffrey

Publication Date
2015

Peer reviewed|Thesis/dissertation

Essays on Technology and the Environment from an International Perspective

by

Geoffrey Masters Barrows

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Zilberman, Chair
Associate Professor Meredith Fowlie
Professor Andrés Rodríguez-Clare

Spring 2015

Essays on Technology and the Environment from an International Perspective

Abstract

Essays on Technology and the Environment from an International Perspective

by

Geoffrey Masters Barrows

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor David Zilberman, Chair

In this dissertation, I present three essays that consider the environmental consequences of technological change, from an international perspective. The first two chapters use firm-level production data to estimate the response of CO_2 emission intensity to changes in competition in foreign markets. The first chapter estimates this response with respect to foreign demand shocks, i.e., a positive shock to exports. The second chapter exploits a specific liberalization episode to estimate the impact with respect to foreign competition shocks, i.e., a negative shock to exports. Both papers are co-authored with Hélène Ollivier. The final chapter analyzes the decision to adopt genetically engineered seeds in different countries around the world, and the attendant impacts on supply and land-use. This last chapter is co-authored with David Zilberman and Steven Sexton and was previously published in Environment and Development Economics.

The first chapter investigates the impact of exporting on the CO_2 emission intensity of manufacturing firms in India. Recent papers have argued that export market access encourages firms to upgrade technology, which lowers the emission intensity of production; however, data limitations confound previous attempts to separately identify productivity impacts from simultaneous changes in prices and product-mix. We present a model of how these alternative channels could also explain the results documented in the literature. Then, using a highly detailed production dataset of large Indian manufacturing firms that contains information on physical units of inputs and outputs by product, we are able to decompose the overall firm impact into three components – prices, product-mix, and technology. Export impacts at the firm level are identified from import demand shocks of foreign trading partners. We find that prices systematically bias down estimates of emission intensity in value, that firms adjust emission intensity in quantity through changing output shares across products, but that firms do not lower emission intensity within products over time (technology). The results imply that the productivity benefits from market integration alone are not enough to induce clean technology adoption.

The second chapter investigates the "third-party" impact of trade liberalization on the environmental performance of firms in countries that lose market share as a result of the liberalization. If competition matters for exporting (as previous research indicates), and exporting matters for emission intensity, then emission intensity reductions in liberalized markets may be offset by emission intensity increases in countries peripheral to the liberal-

ization. To test for this indirect effect, we exploit quasi-natural variation arising from the elimination of quota constraints on textile and apparel exports to the US between 1994 and 2007. Using a detailed panel of production and emission data at the firm-product level, we find that Indian exporters in Prowess lost on average 14% export sales as a result of liberalized trade between the US and India's competitors. This loss of export sales was accompanied by an increase in CO_2 intensity of 9%. The results do not appear to be due to fuel-switching, but there is suggestive evidence that capital investments and switching to higher emission intensity varieties may have played a role. Overall, the results support the importance of international competition for production and pollution decisions of firms around the world.

The final chapter uses aggregate data to estimate supply, price, land-use, and greenhouse gas impacts of genetically engineered (GE) seed adoption due both to increased yield per hectare (intensive margin) and increased planted area (extensive margin). An adoption model with profitability and risk considerations distinguishes between the two margins, where the intensive margin results from direct "gene" impacts and higher complementary input use, and the extensive margin reflects the growing range of lands that become profitable with the GE technology. We identify yield increases from cross-country time series variation in GE adoption share within the main GE crops- cotton, corn, and soybeans. We find that GE increased yields 34% for cotton, 12% for corn and 3% for soybeans. We then estimate quantity of extensive margin lands from year-to-year changes in traditional and GE planted area. If all production on the extensive margin is attributed to GE technology, the supply effect of GE increases from 5% to 12% for corn, 15% to 20% for cotton, and 2% to 40% for soybeans, generating significant downward pressure on prices. Finally, we compute "saved" lands and greenhouse gases as the difference between observed hectarage per crop and counterfactual hectarage needed to generate the same output without the yield boost from GE. We find that all together, GE saved 13 million hectares of land from conversion to agriculture in 2010, and averted emissions are equivalent to roughly 1/8 the annual emissions from automobiles in the US.

Contents

List of Figures

List of Tables

Acknowledgments

In preparing my dissertation, I benefited from discussions with many members of the academic community at Berkeley and elsewhere. To begin an inexhaustive list, Thibault Fally and Meredith Fowlie each contributed substantial input throughout the process. David Zilberman and Steven Sexton co-wrote chapter 3 with me, and were excellent collaborators. David Zilberman also chaired my dissertation committee and provided invaluable support and guidance throughout. Andrés Rodríguez-Clare, Larry Karp and Hélène Ollivier were all influential in chapters 1 and 2, with Hélène contributing as a co-author. Certainly, Hélène also deserves credit in a special overarching category of her own. Berkeley ARE faculty Jeremy Magruder, Ethan Ligon, Leo Simon, and Betty Sadoulet provided insight and general wisdom. I am further indebted to Amit Khandelwal, Ben Faber, James Harrigan, Ted Miguel, and Ann Harrison for helpful comments and encouragement. I'd also like to thank fellow graduate students Matthew Zaragoza-Watkins, Walter Graf, and Judson Boomhower for valuable input and discussions. Finally, I'd like to acknowledge the staff at the Center for Monitoring the Indian Economy, especially Sadhana Shah and Amanda Aversano for their help with the Prowess data.

Chapter 1

Does Trade Make Firms Cleaner? Theory and Evidence From Indian Manufacturing

With Hélène Ollivier

1.1 Introduction

Economists, policy-makers, and the general public have long been concerned about the environmental consequences of globalization. A primary fear is that free trade encourages dirty industries to relocate to poor countries, where environmental regulation is weak and production less efficient (Antweiler, Copeland, and Taylor, 2001; Frankel and Rose, 2005; Copeland and Taylor, 2004; Levinson, 2009). Yet, many argue that trade also encourages profit-maximizing firms to voluntarily increase productivity (Verhoogen, 2008; Lileeva and Trefler, 2010; Bustos, 2011; De Loecker, 2011; De Loecker et al., 2012), which may reduce emission intensity at the firm level (Forslid, Okubo, and Ulltveit-Moe, 2011; Cui, Lapan, and Moschini, 2012; Batrakova and Davies, 2012).[1]

Despite a large and decidedly mixed empirical literature on the former effect, we still know very little about the latter. Empirical estimates from both cross-sectional and panel data seem to support the hypothesis that trade (exporting, in particular) lowers the emission intensity of firms, but data is usually not detailed enough to identify the underlying channel. That is, we still don't now *why* firm-level emission intensity responds to trade shocks.[2] Recent theoretical models of trade's impact on emission intensity posit a technological upgrading mechanism à la Lileeva and Trefler (2010) and Bustos (2011), wherein the prospect of sizable

[1]The correlation between productivity and emission intensity is not necessarily negative. In general, the sign of this correlation depends critically on the underlying mechanism through which trade stimulates adjustments to the production process (a central point of this paper). However, most models (discussed below) feature Hicks-neutral productivity gains, which would imply trade-induced reductions in emission intensity as well.

[2]See Holladay (2010); Forslid, Okubo, and Ulltveit-Moe (2011); Cui, Lapan, and Moschini (2012); Galdeano-Gómez (2010); Batrakova and Davies (2012) for evidence in the corss-section and Gutiérrez and Teshima (2011); Martin (2012); Cherniwchan (2013) for within-firm estimates over time.

export sales induces firms to invest in productivity-enhancing, cleaner technology.[3] But with firm-level datasets usually denominated only in *value*, not physical quantities, and neither inputs nor outputs disaggregated to the product-level, the technology channel is usually not separately identified from other margins of adjustment at the firm-level (as discussed in De Loecker (2011); De Loecker et al. (2012)).

To fix ideas, consider an accounting decomposition of average emission intensity in value of a multi-product firm:

$$\frac{E_i}{V_i} = \sum_j \underbrace{\frac{E_{ij}}{Q_{ij}}}_{Technology} * \underbrace{\frac{Q_{ij}}{V_{ij}}}_{Price} * \underbrace{s_{ij}}_{Product-mix} \tag{1.1}$$

where E_i and V_i denote the environmental emissions and total sales generated by firm i, and E_{ij}, Q_{ij}, V_{ij}, s_{ij} correspond to emissions, sales, output, and (within-firm) sales share for firm-product ij. If export market access induces firms to adopt new, cleaner technology, then E_{ij}/Q_{ij} should fall with exports. However, the variables that are usually available in firm-level datasets include only E_i and V_i, so the best measure of emission intensity that can be constructed is just the left hand side of equation (1.1), E_i/V_i. Gutiérrez and Teshima (2011); Martin (2012); Cherniwchan (2013) each present evidence of how this measure (E_i/V_i) evolves within firms over time in response to trade shocks, and mostly find that it falls with increased export market access.[4] But note that if contemporaneous changes to prices and product-mix (captured in the second and third terms on the right hand side of equation (1.1)) also adjust endogenously with the export decision, then E_i/V_i could fall with exports without any change in technology. A growing list of papers illustrates the sensitivity of the latter two margins to trade shocks (De Loecker, 2011; Harrigan, Ma, and Shlychkov, 2011; De Loecker et al., 2012; Manova and Zhang, 2012; Bernard, Redding, and Schott, 2011; Mayer, Melitz, and Ottaviano, 2014), which gives reason to believe that other things could be happening within the firm beyond technological change.

In this paper, we present new evidence on the underlying channels through which changes in foreign market access impacts the emission intensity of firms. We first give a theoretical account of how both prices and product-mix could determine trade's impact on emission intensity. While these channels have been explored in other settings (De Loecker, 2011; De Loecker et al., 2012; Bernard, Redding, and Schott, 2011; Mayer, Melitz, and Ottaviano, 2014), they are novel to the trade and environment literature and bear directly on the interpretation of firm-level estimates. Next, we test the qualitative predictions of the model using a uniquely-detailed panel dataset of Indian manufacturers – the Prowess dataset – which allows us to compute firm-product emission intensity in quantity E_{ij}/Q_{ij} without imposing any functional form assumptions on production. Since these measures are net of price and product-mix effects by construction, we are able to perform the first direct test of

[3]Other mechanisms beyond the fixed-cost model of Lileeva and Trefler (2010) and Bustos (2011) could also explain trade-induced technological change, including bankruptcy costs (Durceylan, 2009), trapped factors (Bloom, Draca, and Van Reenen, 2011), and searching (Perla, Tonetti, and Waugh, 2012), but the mechanism mostly cited is the Bustos (2011) model.

[4]Empirical estimates usually deflate firm-level sales by an industry price index, but this procedure implicitly imposes functional form assumptions and neglects firm-specific deviations in price.

the clean technology hypothesis (i.e. $\frac{\partial E_{ij}/Q_{ij}}{\partial exports_i} < 0$). Additionally, we present evidence of the impact of the latter two channels– prices and product-mix.

In the model, firms have a core-competency product (as in Mayer, Melitz, and Ottaviano (2014)), and firms skew production towards this core-competency for the export market. Each product has a unique emission intensity associated with its production and unique destination-specific pricing. With export prices tending towards higher values than domestic prices, demand shocks in the foreign market increase average price, hence lowering E_i/V_i. However, demand shocks also skew production towards the core-competency, which increases (decrease) E_i/Q_i if core products are dirtier (cleaner) to produce than non-core products. This correlation is shown to depend on the relative magnitude of parameters of the production function. Thus, the sign of the contribution of product-mix to firm-average emission intensity is theoretically ambiguous, and hence, an empirical question.

In order to compute firm-product emission intensities, we rely on firm-product-specific energy input data contained in Prowess. It is extremely rare to observe inputs at the firm-product level because of the reporting burden on firms (in fact, to our knowledge, this is the only dataset that breaks down inputs by product), but concerns for energy security led the Indian government to require firms starting in 1988 to issue detailed product-specific energy input usage data in their publically available annual assessments.[5] Since the firm-product energy data have not been analyzed before, we describe them in detail in the main text (and in Appendix A.2) and perform several diagnostic checks. We find that the aggregate CO_2 trajectory implied by these firm-product energy intensities align with official India-wide estimates from manufacturing, that implied industry-average emission intensities correlate strongly with those computed from an independent report (the World Input Output Database), and that implied firm-level emissions match the reported firm-level emissions from the firm-level aggregate energy consumption data in Prowess. Finally, we can reject the hypothesis that firms merely apportion energy usage based on ready-to-hand measures like sales and output.

Based on the energy reports, we compute firm-product CO_2 intensities following a standard procedure in the literature by multiplying physical quantities of energy source (e.g. coal, diesel, etc) by constant CO_2 per quantity coefficients and aggregating across energy source. We then relate these measures to exporting behavior of the firm over the period 1990-2011.

The key empirical challenge we face is that exporting is an endogenous decision that could correlate with other determinants of firm and firm-product emission intensity. Kellenberg (2009) finds that environmental regulation is jointly determined with strategic trade interests, which could drive both emission intensity and exports. Reverse causality could also play a role, if foreign consumers have a preference for green production techniques. Or, third-party consulting from either private firms or government extension officers may encourage both different production techniques and foreign market strategies simultaneously, which would mechanically link emission intensity to exporting. Differential changes in trade barriers alone is not enough to solve the endogeneity problem, because trade barriers often change gradually

[5]The Center for Monitoring the Indian Economy (CMIE) collected a large subset of these reports and digitized the information in the dataset Prowess. Energy is the only input reported at the firm-product level. All other inputs such as labor and capital are reported at the firm level (i.e., aggregated across all products). See Goldberg et al. (2010); De Loecker et al. (2012) for further descriptions of the output data.

over time, along with other macro factors, and can even change endogenously to strategic industry interests (Trefler, 1993).

To address these endogeneity concerns, we follow recent papers in the trade literature that identify firm-level trade impacts from macroeconomic fluctuations in trading partner markets (Brambilla, Lederman, and Porto, 2012; Hummels et al., 2014; Bernard, Moxnes, and Ulltveit-Moe, 2014; Bastos, Silva, and Verhoogen, 2014). Specifically, we instrument export sales of firms in Prowess with the weighted average import demand for goods from countries *other* than India in those foreign destinations that India exports to. The identification assumption is that foreign demand shocks are exogenous to unobservable factors that impact Indian firm-level emission intensity. Changes in weighted-average foreign import demands are shown to vary significantly across product categories, and thus deliver differential changes to export market opportunities for firms in Prowess operating in different product codes. As in Hummels et al. (2014); Bernard, Moxnes, and Ulltveit-Moe (2014), we find that foreign demand shocks have strong predictive power for firm-level exports: a 1% increase in foreign demand increases firm-level export value by 0.17% on average. The instruments are also shown to be uncorrelated with prior trends in emission intensity, hence exporting impacts can be identified from a difference-in-difference-like IV estimation.

We find in our sample that emission intensity in both value and quantity fall as firms exports more, both in the OLS and when instrumenting with foreign demand shocks. We estimate that E_i/V_i falls roughly 0.5% with a 1% increase in (instrumented) export value. This figure is broadly in line with estimates from Gutiérrez and Teshima (2011); Martin (2012); Cherniwchan (2013). Next, we decompose the firm-level estimate into a price effect and a quantity effect by replacing nominal sales V_i with real output Q_i, and find that 1/3rd the overall impact is due just to prices. This finding is consistent with the idea that higher export prices mechanically inflate the denominator in E_i/V_i. Netting out prices, we estimate that a 1% increase in (instrumented) export value lowers E_i/Q_i 0.38%.

These firm-level reductions in emission intensity represent real benefits for the environment, but they do not necessarily indicate that firms are adopting cleaner technology. Given that the model predicts product-mix could skew production towards cleaner or dirtier products, the within-product technology channel could, in principle, be either larger or smaller in magnitude than the 0.38% figure. To distinguish technological change from product-mix, we disaggregate further to the firm-product level and estimate the technological channel in isolation. At this level, we find that we can reject negative impacts at the 5% level. This implies that in this context, *all* of the real firm-level impact is channeled through changes in product-mix, not technological upgrading. The results hold for two different energy reports from the firm, and are robust to restricting the sample to short "quality ladder" industries (Khandelwal, 2010). We also present direct evidence that the foreign demand shocks increase average unit value and the share of production devoted to the largest product (i.e., product-mix). Thus, the empirical estimates support the model: while emission intensity falls with exports at the firm level, the driving mechanisms seem to be prices and product-mix, not technological upgrading.

The paper relates to the recent firm-level trade and environment literature discussed above, as well as the classic pollution haven literature (see Copeland and Taylor (2004) for a review). In Heckscher-Ohlin-based pollution haven models, trade can also impact firm-level emission intensity, but only through endogenous environmental regulation. The mechanisms investigated in this paper operate independent of regulation, thus represent a separate (potentially additional) firm-level mechanism. That is, the estimates in this paper and elsewhere in the recent literature represent a lower bound of trade's impact on emission intensity, since they abstract from longer-run endogenous changes to environmental regulation.

Beyond the trade and environment literature, the paper relates to the broader trade literature that connects the destination of exports to input and output choices of individual firms. Verhoogen (2008); Kugler and Verhoogen (2012); Brambilla, Lederman, and Porto (2012); Bastos, Silva, and Verhoogen (2014) find that exporting to high-income countries leads firms to increase the skill intensity of production, which thus increases wages. Manova and Zhang (2012) and Faber (2012) extend the result to imported material inputs. We find that demand shocks in high-income countries also impact energy inputs (via product-mix), showing that export destination also matters for environmental effects.

On the output side, several papers find that destination market matters for output prices and product-mix (Harrigan, Ma, and Shlychkov, 2011; Manova and Zhang, 2012; Goldberg et al., 2010; Eckel and Neary, 2010; Iacovone and Javorcik, 2010; Bernard, Redding, and Schott, 2011; Bernard, Redding, and Schott, 2013; Mayer, Melitz, and Ottaviano, 2014). Our results support these findings and provide the first link between the product-mix and input-mix directly.

Finally, the paper relates to a nascent literature that studies the causes of high emission intensity of firms in developing countries (Duflo et al., 2013; Greenstone and Hanna, 2014). Our results show that firms' environmental performance do respond to market incentives (product-mix), but not as strongly as they could (technological upgrading), absent pre-existing market failures.

The paper proceeds as follows. In section 2, we present a multi-product heterogeneous firm trade model similar to Mayer, Melitz, and Ottaviano (2014), which allows us to assess the price and product-mix impacts on firm-level emission intensity. In section 3, we present the Prowess energy data and the basic correlations with export orientation. In section 4, we discuss the identification strategy, the aggregate trade data from which we compute the instruments, and how we merge these statistics to the Prowess dataset. For this merge, we design a new mapping from CMIE product classification codes to Harmonized System (HS) 6-digit trade data, which we present as another contribution of the paper. Section 5 presents the results, and section 6 concludes.

1.2 Theoretical Framework

In this section, we present a heterogeneous-firm multi-product general equilibrium trade model that connects foreign demand shocks to firm-level exports, prices, product-mix and emission intensity. The model extends the Mayer, Melitz, and Ottaviano (2014)'s multi-product firm model (henceforth, MMO) to include two factors of production, where the second factor is emissions (or equivalently, energy). This extension allows us to compute emission intensity in both quantity and value, while allowing for substitution between factors.

The model delivers four comparative statics that we take to the data. First, foreign demand shocks increase firm-level exports. Second, foreign demand shocks alter the product-mix of the firm. In particular, the mix of the export basket skews towards "core competency" (i.e., best product). Third, this change in product-mix in turn alters firm-level average emission intensity in quantity. The sign of the impact depends on whether the firms' core competency is cleaner or dirtier than other varieties within each firm, and on relative market conditions. Fourth, foreign demand shocks influence firm-level average price, which in turn impacts firm-level emission intensity in value. In particular, demand shocks in high-environmental-regulation (rich) countries increase average prices, which lowers emission intensity in value.

These findings show that exporting can influence firm-level emission intensity in quantity through the product-mix, and firm-level emission intensity in value via prices, which are novel predictions in the literature. The full general equilibrium properties of the model are explored in Barrows and Ollivier (2014), so we move briskly through the model setup to discuss the firm-level impacts.

Setup of the Model

The world is comprised of H countries, indexed by $l \in \{1, ..., H\}$, each of which contains heterogeneous firms that decide whether to enter, what products to produce, and where to supply these products. Firms supply horizontally differentiated varieties that are substitutes in demand. Countries are asymmetric in terms of their market sizes (L), bilateral trade costs (θ), and environmental regulation (τ).

Preferences and Endowments.— The representative agent in each country l has preferences over a continuum of differentiated varieties indexed by $j \in \Lambda$, and a homogeneous good chosen as numeraire:

$$U = q_0^c + \alpha \int_{j\in\Lambda} q_j^c dj - \frac{1}{2}\gamma \int_{j\in\Lambda} (q_j^c)^2 dj - \frac{1}{2}\eta \left(\int_{j\in\Lambda} q_j^c dj \right)^2, \tag{1.2}$$

where q_0^c and q_j^c represent the individual consumption levels of the numeraire good and each differentiated variety j. The demand parameters α, γ, and η are all positive. An increase in α and a decrease in η both shift out the demand for the differentiated varieties relative to the numeraire. The parameter γ indexes the degree of product differentiation between the varieties. In the limit when $\gamma = 0$, varieties are perfect substitutes.

With quasi-linear preferences, demand for each differentiated variety is linear in prices. Let $\Lambda_l^* \in \Lambda$ be the subset of varieties that are consumed in country l. A linear market demand system for these varieties in country l is derived:

$$q_{jl} \equiv L_l q_{jl}^c = \frac{\alpha L_l}{\eta M_l + \gamma} - p_{jl}\frac{L_l}{\gamma} + \frac{\eta L_l \bar{p}_l M_l}{\gamma(\eta M_l + \gamma)}, \ \forall j \in \Lambda_l^*, \tag{1.3}$$

where M_l is the measure of consumed varieties in Λ_l^* and $\bar{p}_l = (1/M_l)\int_{j\in\Lambda_l^*} p_{jl} dj$ is their average price. A consumer may not have positive demand for any particular variety, but she has positive demand for the numeraire good by assumption. Thus, the set Λ_l^* is the largest subset of Λ that satisfies

$$p_{jl} \leq \frac{\alpha\gamma + \eta \bar{p}_l M_l}{\eta M_l + \gamma} \equiv p_l^{max}, \tag{1.4}$$

where the price bound p_l^{max} represents the price at which demand for a variety is driven to zero in country l. An increase in the number of varieties and a decrease in the average price reduces the price bound p_l^{max} in the destination market l. Both domestic and foreign producers face the same price bound, and any decrease in p_l^{max} can be interpreted as the competitive environment in country l becoming "tougher."

Production Technology.— The specification of entry and production follows Mayer, Melitz, and Ottaviano (2014) with the extension made by Barrows and Ollivier (2014). Each economy is composed of two sectors, one producing a non-polluting homogeneous good, which takes only labor ℓ as an input, and the other producing differentiated varieties, which requires both ℓ and energy z. The consumption of energy generates pollution, so without loss of generality, z, can also be thought of as emissions. Labor is mobile across sectors and is inelastically supplied in a competitive market. The numeraire good is produced under constant returns to scale at unit cost and sold in a competitive market. These assumptions imply a unit wage. The price of emissions depends on an exogenous environmental tax τ that is fixed by the national government.[6]

In order to begin production in the differentiated sector, firms must incur a sunk entry cost of $f_E > 0$ units of labor, whatever the country of location. This cost is associated with research and product development, which entails uncertain outcomes. Firms face uncertainty about their total factor productivity (TFP) denoted by φ. Entrants draw their firm ability from a known Pareto distribution with cumulative function $G(\varphi) = 1 - \varphi^{-k}$ with support on $[0, \infty]$. The shape parameter k indexes the dispersion of productivity draws. Since the entry cost is sunk, firms that can cover their marginal cost for at least one good survive and produce. If the firm is particularly efficient, it can decide to introduce multiple varieties, each variety being produced with a different technology. Each firm has one core variety with minimal marginal cost given the tax level τ, and new varieties can be added with higher marginal costs. We index by m the varieties produced by the same firm in increasing order of distance from its core variety $m = 0$.

The production function of a variety m by a firm with total factor productivity φ is:

$$q(\varphi, m) = \varphi[(e^{-\sigma m}\ell)^{\epsilon} + (e^{-\nu m}z)^{\epsilon}]^{1/\epsilon}. \tag{1.5}$$

This function combines "effective" inputs in the standard CES structure where "effective input" equals actual input scaled by a distance function from core competency. Production function (1.5) is quasi-concave if $\epsilon \leq 1$, which is assumed in the rest of the paper. As demonstrated in Barrows and Ollivier (2014), we need to impose that $\nu > 0$ and $\sigma > 0$ to ensure that the core competency of a firm corresponds to $m = 0$. This implies that the unit cost function of variety m is increasing in m, where the unit cost function is given by

$$\Phi(\varphi, m) = w\ell + \tau z = \frac{1}{\varphi}\left[\tau^{\frac{\epsilon}{\epsilon-1}} e^{\frac{\nu m \epsilon}{\epsilon-1}} + e^{\frac{\sigma m \epsilon}{\epsilon-1}}\right]^{\frac{\epsilon-1}{\epsilon}}. \tag{1.6}$$

[6]alternatively, if z is energy, then τ is the price of energy inclusive of regulation. Either way, τ is fixed by the government.

We define the emission intensity of a variety m in *physical* output, $EQ(\varphi, m)$ (equivalently, E_{ij}/Q_{ij} from (1.1)) as

$$EQ(\varphi, m) = \frac{z}{q} = \frac{1}{\varphi}\left[e^{-\nu m\epsilon} + \left(\tau e^{m(\epsilon\nu-\sigma)}\right)^{\frac{\epsilon}{1-\epsilon}}\right]^{-1/\epsilon}. \tag{1.7}$$

and emission intensity in value as

$$EV_{lh}(\varphi, m) = \frac{z}{p_{lh}(\varphi, m)q_{lh}(\varphi, m)} = \frac{EQ(\varphi, m)}{p_{lh}(\varphi, m)}. \tag{1.8}$$

where $p(\varphi, m)$ denote the price, and subscripts lh denotes variables derived from producing in l and selling in h. Prices $p_{lh}(\varphi, m)$ vary with market conditions in h, hence, emission intensity in *value* depends on the conditions of the market where the product is sold. Note that $EQ(\varphi, m)$ is fixed conditional on environmental regulation τ. Thus, the model abstracts from technology upgrading at the firm-product level. We make this assumption for simplicity so we can assess the price and product-mix channels. No additional insight would be gained by allowing technological upgrading as in Forslid, Okubo, and Ulltveit-Moe (2011); Cui, Lapan, and Moschini (2012); Batrakova and Davies (2012), but notation would proliferate.

If we further impose $\epsilon > 0$ and $\nu > \sigma$, then we have the following lemma:

Lemma 1. *The emission intensity of variety m, $E(\varphi, m)$ is increasing (decreasing) in m if and only if*

$$\epsilon - (1-\epsilon)\left(\tau e^{m(\nu-\sigma)}\right)^{\frac{\epsilon}{\epsilon-1}} < (>)\frac{\sigma}{\nu}.$$

Proof.

$$\frac{dE(\varphi, m)}{dm} = \frac{1}{\varphi}\left[e^{-\nu m\epsilon} + \left(\tau e^{m(\epsilon\nu-\sigma)}\right)^{\frac{\epsilon}{1-\epsilon}}\right]^{\frac{-1-\epsilon}{\epsilon}}\left[\nu e^{-\nu m\epsilon} - \frac{\epsilon\nu-\sigma}{1-\epsilon}\left(\tau e^{m(\epsilon\nu-\sigma)}\right)^{\frac{\epsilon}{1-\epsilon}}\right],$$

which is positive if and only if the LHS of the inequality in Lemma 1 is less than the RHS. This holds if $\epsilon \leq 0$ or if $\sigma > \nu$. □

Lemma 1 suggests that the emission intensity and the unit cost of new varieties are not necessarily correlated. Whereas the unit cost is increasing in m if parameters ν and σ are positive, $E(\varphi, m)$ can still be either increasing or decreasing in m depending on production parameters. If either $\epsilon \leq 0$ or $\sigma > \nu$, then the LHS of the inequality in Lemma 1 is for certain less then the RHS, which means higher-m varieties are more emissions intensive.[7] However, if $\epsilon > 0$ and $\nu > \sigma$, then the equality could be reversed (depending on relative

[7]For example, when the elasticity of substitution between emissions and labor is positive and close to 1 ($\epsilon \to 0$), both unit cost and emission intensity are increasing in m. This corresponds to the standard framework in the trade and environment literature following Copeland and Taylor (2004) where emissions are a by-product of production, and where abatement requires labour in such a way that net output can be represented in a Cobb-Douglas function with emissions and labour as inputs. To illustrate, consider the following Cobb-Douglas function for variety m: $q(\varphi, m) = \varphi(e^{-\sigma m}\ell)^{\beta}(e^{-\nu m}z)^{1-\beta}$. It implies that both the unit cost function and the emission intensity function can be factorized by $e^{m[\beta\sigma+(1-\beta)\nu]}$. Thus these functions are both increasing in m for $\sigma, \nu > 0$.

magnitudes), implying that higher-m products are cleaner. We do not want to assume away the possibility that trade could increase emission intensity at the firm level, so we restrict attention to the case $\epsilon > 0$ and $\nu > \sigma$ for the rest of the paper. These conditions imply that labor and emissions exhibit a high-degree of substitutability, and higher-m varieties use even less efficiently emissions than labor. Under these conditions, firms substitute labour for emissions as they add higher-m products, which might reduce emission intensity.

Firm Behavior.— Firms engage in monopolistic competition on each destination market and markets are segmented, so their profit maximization problem treats each market separately, taking the average price level $\bar{p}_l$ and total number of varieties M_l as given. Firms that can cover at least the marginal cost of production of their core competency survive and produce. All other firms exit the industry. Surviving firms maximize their profits using the residual demand function (1.3) on each market, subject to a variable "iceberg" trade cost $\theta_{lh} > 1$ ($\theta_{ll} = 1$), which drives the delivered marginal cost of a variety m produced by firm φ in country l to the import country h to $\theta_{lh}\Phi(\varphi, m)$.

The profit maximizing price $p_{lh}(\varphi, m)$ and output level $q_{lh}(\varphi, m)$ of a variety with marginal cost $\Phi(\varphi, m)$ produced in country l and sold in country h must then satisfy

$$q_{lh}(\varphi, m) = \frac{L_h}{\gamma}\left[p_{lh}(\varphi, m) - \theta_{lh}\Phi(\varphi, m)\right]. \tag{1.9}$$

The variety is supplied to country h if and only if the maximizing price $p_{lh}(\varphi, m)$ is below the price bound p_h^{max} from (1.4). Let Φ_{lh} denote the unit cost of the marginal variety produced in country l and sent to country h achieving zero sales. Its demand level $q_{lh}(\Phi_{lh})$ is driven to zero as $p_{lh}(\Phi_{lh}) = \theta_{lh}\Phi_{lh} = p_h^{max}$. For a firm selling its varieties domestically, (1.9) becomes $q_{ll}(\varphi, m) = L_l\left[p_{ll}(\varphi, m) - \Phi(\varphi, m)\right]/\gamma$, which implies that the domestic cost cutoff is such that $\Phi_{ll} = p_l^{max}$. Therefore, $\Phi_{lh} = \Phi_{hh}/\theta_{lh}$: trade barriers make it harder for exporters to break even relative to domestic producers.

The domestic cutoff Φ_{ll} and the export cutoff Φ_{lh} summarize all the effects of market conditions relevant for each country l's firm performance measures. Firms in country l with marginal cost for their core competency $\Phi(\varphi, 0) > \Phi_{ll}$ cannot profitably produce their core variety for the domestic market and exit. This yields the cutoff productivity for firm survival: $\varphi_{ll} = \left[1 + \tau^{\frac{\epsilon}{\epsilon-1}}\right]^{\frac{\epsilon-1}{\epsilon}}/\Phi_{ll}$. Similarly, firms with marginal cost for their core product $\Phi(\varphi, 0) > \Phi_{lh}$ cannot profitably sell their products to country h, and this yields the cutoff productivity φ_{lh} for exporting to market h.

As in MMO, price, markup, revenue, and profit for a variety produced in cuntry l sold in country h can be written as functions of Φ_{lh} and $\Phi(\varphi, m)$:

$$p_{lh}(\varphi, m) = \frac{\theta_{lh}}{2}[\Phi_{lh} + \Phi(\varphi, m)], \tag{1.10}$$

$$\lambda_{lh}(\varphi, m) = \frac{\theta_{lh}}{2}[\Phi_{lh} - \Phi(\varphi, m)], \tag{1.11}$$

$$r_{lh}(\varphi, m) = \frac{L_h\theta_{lh}^2}{4\gamma}[\Phi_{lh}^2 - \Phi(\varphi, m)^2], \tag{1.12}$$

$$\pi_{lh}(\varphi, m) = \frac{L_h\theta_{lh}^2}{4\gamma}[\Phi_{lh} - \Phi(\varphi, m)]^2. \tag{1.13}$$

Lower marginal cost varieties have lower prices and earn higher profits than varieties with higher marginal costs. However, lower marginal cost varieties also have higher markups, which reveals an incomplete pass-through from firms to consumers.

For each producing firm with $\Phi(\varphi, 0) < \Phi_{ll}$, more profits can be earned by diversifying its product mix. A firm chooses endogenously its product mix by selecting the varieties it produces for a non-negative domestic profit ($\pi_{ll}(\varphi, m) \geq 0$), and the varieties it exports for a non-negative export profit ($\pi_{lh}(\varphi, m) \geq 0$). The total numbers of varieties produced and exported by a firm with productivity φ in country l are thus

$$M_{ll}(\varphi) = max \left\{ m | \left[\tau^{\frac{\epsilon}{\epsilon-1}} e^{\frac{\nu m \epsilon}{\epsilon-1}} + e^{\frac{\sigma m \epsilon}{\epsilon-1}} \right]^{\frac{\epsilon-1}{\epsilon}} / \Phi_{ll} \leq \varphi \right\} + 1 \quad \text{iff } \varphi \geq \varphi_{ll} \tag{1.14}$$

$$M_{lh}(\varphi) = max \left\{ m | \left[\tau^{\frac{\epsilon}{\epsilon-1}} e^{\frac{\nu m \epsilon}{\epsilon-1}} + e^{\frac{\sigma m \epsilon}{\epsilon-1}} \right]^{\frac{\epsilon-1}{\epsilon}} / \Phi_{lh} \leq \varphi \right\} + 1 \quad \text{iff } \varphi \geq \varphi_{lh}, \tag{1.15}$$

and zero otherwise.

Free Entry Condition and Equilibrium.— Entry is unrestricted in all countries. Firms decide where to locate prior to entry and paying the sunk entry cost. We assume that the entry cost f_E and the productivity distribution $G(\varphi)$ are common across countries, and that all countries produce the homogeneous good. A prospective entrant's expected profits in country l are then given by

$$\Pi_l = \sum_{h=1}^{H} \int_{\varphi_{lh}}^{\infty} \left[\sum_{\{m|\Phi(\varphi,m) \leq \Phi_{lh}\}} \pi_{lh}(\varphi, m) \right] dG(\varphi) - f_E,$$

which includes the expected profits made in the domestic market $h = l$ and in foreign markets $h \neq l$. The free entry condition in country l yields

$$\sum_{h=1}^{H} \Omega_h L_h (\theta_{lh})^2 \Phi_{lh}^{k+2} = 2\gamma(k+1)(k+2) f_E, \tag{1.16}$$

where $\Omega_h \equiv \sum_{m=0}^{\infty} \left[\tau_h^{\frac{\epsilon}{\epsilon-1}} e^{\frac{\nu m \epsilon}{\epsilon-1}} + e^{\frac{\sigma m \epsilon}{\epsilon-1}} \right]^{\frac{(1-\epsilon)k}{\epsilon}}$ is a sequence that depends exclusively on parameters. We find that this sequence converges if and only if $\epsilon \leq 1$, $\nu > 0$ and $\sigma > 0$, which are assumed above.[8] Hence, Ω_h is a constant.

Using $\Phi_{lh} = \Phi_{hh}/\theta_{lh}$ and the symmetry across countries that gives a system of equations, we obtain

$$\Phi_{ll} = \left(\frac{2\gamma(k+1)(k+2) f_E}{\Omega_l L_l} \Psi(\boldsymbol{\theta}) \right)^{\frac{1}{k+2}}, \tag{1.17}$$

where $\boldsymbol{\theta}$ is a vector of all country-pair trade costs, and the function $\Psi(.)$ varies depending on the number of countries. If $H = 2$ for instance, we have $\Psi(\theta_{lh}, \theta_{hl}) = (1 - \theta_{lh}^{-k})/[1 -$

[8]We have $\left[\tau_h^{\frac{\epsilon}{\epsilon-1}} e^{\frac{\nu m \epsilon}{\epsilon-1}} + e^{\frac{\sigma m \epsilon}{\epsilon-1}} \right]^{\frac{(1-\epsilon)k}{\epsilon}} \longrightarrow_{m \to \infty} 0$ if and only if $\nu, \sigma > 0$ and $\epsilon \leq 1$.

$\theta_{lh}^{-k}\theta_{hl}^{-k}]$. The domestic cost cutoff thus depends on the domestic market size, on the country's environmental regulation, and on the trade costs incurred for imports and exports. Since the distribution of the exporters' delivered unit cost $\theta_{lh}\Phi(\varphi, m)$ to country h matches the distribution of country h's domestic firms' cost, it results in a matching price distributions for both domestic firms in country h and exporters to that country. In Appendix A.1, we show that $\Phi_{lh} < \Phi_{ll}$ in the non-specialized equilibrium, so that only a subset of relatively more productive firms export, and firms only export their more profitable varieties. This selection into exporting implies that on average exporters are more efficient than non-exporters.

Impacts of Product-Mix and Price Effects on Firm Emission Intensity

In this subsection, we derive comparative static relationships between foreign demand shocks, firm-level exports, product-mix, and firm-level emission intensity in value and quantity, which we then take to the data.

Exports.— Consider an exogenous increase in the size of the foreign market: $dL_h > 0$. In the model, $dL_h > 0$ implies the population of h literally increases, but one could also think of it as any increase in the purchasing power of that population, e.g., a shock to income or exchange rates. With an increase in L_h, there are two effects. First, higher demand leads country h to buy more of everything, including products manufactured in country l. However, this demand shock also attracts more firms to sell in h, driving up competition – i.e., lowering Φ_{hh} given equation (1.17) – and in particular, lowering the export cost cutoff Φ_{lh} for country l's firms. These two forces – a direct effect and a competition effect – pull in opposite directions, and affect country l's heterogeneous firms differently. In particular:

Prediction 1. More productive exporters see their export revenues increase when facing a foreign demand shock whereas less productive exporters see them decrease.

Proof: See Appendix A.1

The average impact in the population of firms depends on the distribution of productivities, along with trade costs and market conditions of trading partners, though the evidence suggests that the direct effect tends to dominate the competition effect in aggregate (Hummels et al., 2014; Bernard, Moxnes, and Ulltveit-Moe, 2014). This would suggest that $dL_h > 0$ is a positive shifter of exports.

Product-Mix.— Within the firm, the import demand shock impacts varieties differently depending on their marginal cost:

Prediction 2. For a firm-product with unit cost $\Phi(\varphi, m)$ manufactured in country l and sold in country h, an increase in L_h lowers (increases) export revenues for less (more) profitable products with unit cost $\Phi(\varphi, m) > (<)\sqrt{\frac{k}{k+2}}\Phi_{lh}$.

Proof: See Appendix A.1

This impact on export revenues is mostly driven by changes in exported quantities. With an increase in L_h, country l's exporters adjust by dropping some of their most expensive varieties (extensive margin), as well as by changing the relative output share of each variety (intensive margin). Multi-product models with CES demand would also feature the extensive margin (e.g., Bernard, Redding, and Schott 2011), but they miss the intensive margin effect because markups are fixed with CES demand. As in MMO, both the intensive and extensive margin impacts skew production towards the core varieties, whose relative markups and profits rise.

The impact of product-mix on the firm average emission intensity of production, denoted $\overline{EQ}_{lh}(\varphi) = \sum_{m=0}^{M_{lh}(\varphi)-1} EQ(\varphi,m) q_{lh}(\varphi,m) / \sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\varphi,m)$, depends on whether $EQ(\varphi,m)$ is increasing or decreasing in m. If $EQ(\varphi,m)$ is increasing in m, so core varieties are cleaner, then these adjustments make firms' production cleaner; whereas if $EQ(\varphi,m)$ is decreasing in m, so core varieties are dirtier, then the reverse is true. Hence, we have:

Prediction 3. An increase in L_h reduces (increases) firm-destination-level emission intensity in *quantity* $\overline{EQ}_{lh}(\varphi)$ if and only if $EQ(\varphi,m)$ is increasing (decreasing) in m.

Proof: see Appendix A.1

While the impact of $dL_h > 0$ on $\overline{EQ}_{lh}(\varphi)$ is straight-forward once we know the correlation between m and EQ, it is unfortunately unobservable in most cases. That is, we usually do not observe inputs broken down by export destination. Instead, we usually observe

$$\overline{EQ}(\varphi) = \sum_{h=1}^{H} \left(\frac{\sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\varphi,m)}{\sum_{h=1}^{H} \sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\varphi,m)} \right) \overline{EQ}_{lh}(\varphi), \tag{1.18}$$

i.e., firm-level average emission intensity in quantity, averaged over all destinations markets. At this level of aggregation, the impact is not as clear because we must aggregate over multiple products as well as over multiple destinations.

To see the different forces at work, suppose there are only two markets ($H = 2$): a domestic market l and a foreign market h. If core varieties are cleaner (i.e., $EQ(\varphi,m)$ is increasing in m), an increase in L_h lowers $\overline{EQ}_{lh}(\varphi)$ by Prediction 3. Because there is selection into exporting, the set of exported varieties is smaller than the set of varieties sold domestically, and product mix in the export basket is more skewed toward core varieties than in the domestic basket. Therefore, we can infer that $\overline{EQ}_{lh}(\varphi) < \overline{EQ}_{ll}(\varphi)$. In this case, an increase in exporting would decrease the firm emission intensity in quantity $\overline{EQ}(\varphi)$. Conversely, if core varieties are dirtier, we expect that $\overline{EQ}_{lh}(\varphi) > \overline{EQ}_{ll}(\varphi)$, and an increase in L_h increases $\overline{EQ}_{lh}(\varphi)$ even further. Exporting more would thus raise the aggregate firm emission intensity in quantity.

Generalizing this result to many destination markets requires comparing the average emission intensity of the basket of goods exported to country h where the demand shock occurs with the average emission intensity of all other baskets. If the former remains lower (higher) than the latter, even after the demand shock, then the average emission intensity $\overline{EQ}(\varphi)$ of efficient firms decreases (increases) with the demand shock, whereas the impact of the demand shock on less efficient firms is ambiguous.

Prices.— Next, we consider the impact of prices on emission intensity in value for individual firm-products (equivalently, single-product firms), and then for multi-product firms. Price effects derive from both pass-through of the demand shock into pricing, as well as aggregating sales over destinations. Since the same firm-product is sold at different prices across markets, an increase in L_h changes firm-product average price by increasing the share of exported goods.

Again, for simplicity, suppose there are two markets: l and h. The average emission intensity in value for a firm-product with unit cost $\Phi(\varphi, m)$ is

$$\overline{EV}(\varphi, m) = \frac{r_{ll}(\varphi, m)}{r_{ll}(\varphi, m) + r_{lh}(\varphi, m)} EV_{ll}(\varphi, m) + \frac{r_{lh}(\varphi, m)}{r_{ll}(\varphi, m) + r_{lh}(\varphi, m)} EV_{lh}(\varphi, m), \quad (1.19)$$

where $EV_{ll}(\varphi, m)$ and $EV_{lh}(\varphi, m)$ are the emission intensity in value in the domestic and foreign markets, respectively, as defined by (1.8). We observe that the emission intensity in value of the export basket is lower than the emission intensity of the domestic basket if and only if the export price is higher than the domestic price. Using (1.10), we have:

Prediction 4. Firms sell the same products at higher prices in countries with stricter environmental regulations.

Proof: see Appendix A.1

Countries with stricter environmental regulations have weaker competition because of higher factor prices. If the foreign market is therefore less competitive than the domestic market ($\Phi_{ll} < \Phi_{hh}$), the export price will be higher, yielding $EV_{lh}(\varphi, m) < EV_{ll}(\varphi, m)$. Following the trade and environment literature (Copeland and Taylor, 2003), environmental regulation differences across countries are generally induced by income differences. Thus, our prediction is consistent with many empirical studies that find that firms sell at higher prices in wealthier destinations (Harrigan, Ma, and Shlychkov, 2011; Bastos and Silva, 2010; Manova and Zhang, 2012).

To study the impact of exporting on firm emission intensity in value, we must compare the prices across markets and the share of products sold in each market, given (1.19). When exporting increases, the relative share of sales coming from exports increases, whereas sales in the domestic market are unaffected. A positive shock to L_h also lowers the export price, which increases $EV_{lh}(\varphi, m)$. Therefore, if the export price remains higher than the domestic price (i.e., $EV_{lh}(\varphi, m) < EV_{ll}(\varphi, m)$), an increase in exporting would reduce $\overline{EV}(\varphi, m)$ as long as the compositional shift (toward country h export basket) outweighs the export price decrease. Otherwise, a positive demand shock can increase $\overline{EV}(\varphi, m)$. In the context of multiple destination markets, the impacts of a demand shock from country h depends on whether the export price to country h is higher or lower than the average price over all other destinations.

Finally, for multi-product firms, we must combine the price and product-mix effects to assess how exporting impacts the emission intensity in value. The impacts of an import demand shock depend on relative market conditions, relative firm productivity, and whether core products are cleaner or dirtier than higher-m varieties. In general, the impacts are ambiguous. However, if we assume that core varieties are cleaner (dirtier) and that the export price to country h is higher (lower) than the average price over all other markets, then exporting more decreases (increases) the firm emission intensity in value as both product-mix and price channels reinforce each other. If these channels have opposite impacts (when core products are dirtier and export price to country h is higher, for instance), then the net effect on the firm emission intensity is ambiguous.

1.3 Data and Preliminary Correlations

To test the predictions of the model, we need production data with both inputs and outputs reported at the product-level in physical quantities and values. While some datasets exist that report *outputs* at the product level (e.g., the US Census of Manufacturers used in Bernard, Redding, and Schott 2011 or the French Customs data used in Mayer, Melitz, and Ottaviano 2014), firms rarely report *inputs* at the product level as well. The production dataset we use – Prowess – offers precisely this unusual feature for energy inputs, thus offering a unique opportunity to separately identify the components of equation (1.1). Additionally, Prowess reports the export share in revenue, which enables us to control for endogenous selection into exporting. We present the basic correlations between exporting and firm emission intensity in this section, and discuss identification in Section 4.

Prowess Data Description

As part of the Indian Companies Act of 1956, Indian firms above a given size threshold are required to issue annual reports on a wide array of economic activity, which the Center for Monitoring the Indian Economy (CMIE) collected and digitized in the dataset Prowess. The reporting requirements make the sample well-suited for analyzing trade impacts, since exporting is generally dominated by large, productive firms (Bernard, Redding, and Schott, 2011; Bernard, Redding, and Schott, 2013; Mayer, Melitz, and Ottaviano, 2014). In the annual reports, firms list quantity and value of sales by product, thus we observe both prices (unit values) and product-mix of the firm. Additionally, Prowess contains rich information on energy use, from which we compute CO_2 intensities. While other researchers have exploited the multi-product dimension of Prowess (De Loecker et al., 2012; Goldberg et al., 2010), the energy data – especially the product-specific data – have not been analyzed before, so we describe it in some detail here.[9]

In accordance with the 1988 amendment to the Companies Act, firms are required to report energy use in two ways in their annual reports. First, firms report consumption *per unit of production* for each product manufactured in each of 140 possible energy sources (see Appendix Figure A.1 for the legislative language). That is, firms report physical quantity of each energy source used to generate a single unit of each output product. We provide a sample report in Figure A.2. From these energy data, we compute emission intensity in physical amounts of CO_2 per physical unit of output by merging CO_2 intensity coefficients and summing over energy sources. We refer to the resulting dataset as the "product-specific" dataset.

Second, firms also report total physical consumption and expenditure from each energy source aggregated across all products each year. From these reports, we compute firm-level CO_2 again by converting each energy source with a CO_2 intensity coefficients and summing over energy sources. The "firm-level" reports merge easily with the output data on firm name, providing a means to distinguish between impacts on E_i/V_i from impacts on E_i/Q_i, while the "product-specific" reports provide independent measures of E_{ij}/Q_{ij} – a free-standing

[9] Lipscomb (2008) also investigates the environmental implications of endogenous portfolio choice of firms in Prowess, but she characterizes products as either "clean" or "dirty" depending on regulatory criteria, not the energy data contained in Prowess itself.

measure of technological efficiency. We will use the two reports in tandem to separate the components of equation (1.1). For a more detailed descripton of the data, see Appendix A.2.

Descriptive statistics for the two datasets are reported in Table 3.2. The output data covers 103,451 firm-product-year observations between 1990-2011. For each observation, we know the 16-digit Prowess product category code, output quantity, output units, and value in millions of current year rupees. In Panel A, we aggregate outputs to the firm-level and merge to firm-year CO_2 emissions and export revenue shares, as well as firm-level energy prices. After aggregating to the firm level, we have 42,026 firm-year observations, representing 4,982 firms and 14,958 distinct firm-products. The average firm generates 1.16 billion rupees of revenue, 0.1 MT (Mega Tons) of CO_2 in a given year and earns 11.64% of sales from exports.

A caveat to mention is that while the data matches reasonably well with other better-known datasets (see below), outlier observations are a significant problem in Prowess, as in many production datasets. Upon inspection of outlier values, it appears in many cases as if decimals have been transposed or units mis-reported. We adopt the standard approach of dropping the top and bottom 1% of values for emission intensity for most of the analysis. Additionally, we drop firms that exhibit excessive variation in emission intensity over the period. If a firm's emission intensity changes by several orders of magnitude over the sample, we assume this is due to reporting errors (misplaced decimal, misreported units) and drop it from the analysis (see Appendix A.2). This second data cleaning procedure drops approximately 3% of the data.

Table 1.1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Panel A: Firm-level Dataset					
Sales Value (Bill of Rs)	1.16	2.59	0	26.97	42026
Production (Various Units)	34.48	775.09	0	99599	42026
CO_2 Emissions (MT)	0.1	0.44	0	6.60	42026
Export Share	0.11	0.22	0	1	42026
Energy Use (Million mmBTU)	0.95	4.95	0	117.69	42026
Panel B: Product-specific Dataset					
Log CO_2 Intensity (Kg CO_2/unit)	5.87	2.76	-4.33	12.68	48037
Log Energy Intensity (mmBTU/unit)	0.66	2.87	-10.46	8.58	48037

Notes: Annual production data from Prowess covering years 1990-2011. Sales value and production in Panel A have been aggregated to the firm-level. Production data are reported in various units, e.g. tonnes, number, etc. CO_2 emissions are imputed by multiplying physical quantities of energy source by a source-specific CO_2 intensity coefficients and summing over energy sources (see Appendix A.2 for details). Export share is defined as the export revenue free on board (f.o.b.) divided by total revenues. In Panel A, we have dropped 3% of observations with very high variation in emission intensity within the firm over time, and the top and bottom 1% of sales value and CO_2 emissions observations (see Appendix A.2 for details). In Panel B, we have dropped the top and bottom 1% of CO_2 intensity values.

Panel B reports Log CO_2 emission intensity and Log mmBTU per quantity of outputs (various units) for 48,037 firm-product-year observations between 1990-2011, covering 3,483 firms and 6,237 distinct firm-products. Here, we also trim the top and bottom 1% of values

Table 1.2: Multi-product Firms by Industry

(1) Industry	(2) # Firms	(3) 1-Product	(4) 2-Product	(5) 3-Product	(6) ¿3-Product
Panel A: Firm-level Dataset					
Agricultural products	348	186	66	42	54
Mineral Products	88	54	21	5	8
Food products, beverages & tobacco	368	116	102	58	92
Textiles	980	428	284	161	107
Wood, Pulp & Paper Products	193	134	33	18	8
Chemicals	988	349	225	142	272
Plastics & Rubbers	387	139	91	77	80
Non metallic mineral products	264	141	69	17	37
Base Metals	620	243	170	93	114
Machinery	459	132	101	53	173
Transport equipment	215	67	56	33	59
Misc. Manufactured Articles	77	8	24	12	33
Total	4987	1997	1242	711	1037
Panel B: Product-specific Dataset					
Agricultural products	252	222	20	9	1
Mineral Products	51	46	3	0	2
Food products, beverages & tobacco	236	193	27	9	7
Textiles	709	546	121	24	18
Wood, Pulp & Paper Products	172	145	22	3	2
Chemicals	580	356	106	52	66
Plastics & Rubbers	263	184	45	14	20
Non metallic mineral products	196	127	46	10	13
Base Metals	537	370	88	33	46
Machinery	257	154	43	20	40
Transport equipment	106	70	20	6	10
Misc. Manufactured Articles	192	146	28	11	7
Total	3551	2559	569	191	232

Notes: Total number of firms by industry along with breakdown by "1-product", "2-product", etc. Firms are assigned to an industry based on the product that accounts for the greatest aggregate sales over the entire period (1990-2011). Firms are then allocated to "1-product", "2-product", etc designations based on the number of Prowess product categories the firm operates in over the entire period.

and throw out firm-products for which reporting error seems likely. The reported unit of output varies from product to product, but we restrict the sample to common units within the firm-product over time so that emission intensities are comparable across periods. Almost all the firm-products were already reported in consistent units, so this restriction drops very few observations.

The distributions of firms and products by industry are reported in Table **??**. Industry descriptions based on the Prowess product classification system are listed in the first column. Column 2 reports the firm count by industry, and columns 3-6 breakdown the total firm count into 1-product firms, 2-product firms, etc. A "1-product" firm means the firm reports at most a single product over the period, while a "2-product" firms reports at most 2 distinct products etc. In the firm-level dataset (Panel A), we find that about 60% of firms are multi-product firms. By contrast, in the product-specific dataset (Panel B), only 30% of the firms are multi-product. Both datasets have a fairly broad coverage across all manufacturing industries.

Diagnostic Checks of Product-specific Energy Data

The emissions calculated in Table 3.2 constitute the most-detailed measures of environmental performance of firms available, to our knowledge. With emission intensity in physical quantity of output at the firm-product level, we can measure changes in pollution intensity via technology directly from the data. While incredibly detailed, the drawback of the Prowess data is that it is self-reported. Duflo et al. (2013) shows that Indian firms systematically under-report pollution emissions, which suggests our emissions estimates could be biased. While firms obviously have an incentive to under-report local pollutants like NO_x and PM_{10}, it is not clear that firms benefit from under-reporting energy consumption. Even if there were systematic measurement error in the energy data, it is not obvious how the bias would correlate with the export decision. Nonetheless, we perform several diagnostic tests in Appendix A.2 to assess the quality of the data.

First, we compute aggregate emissions by energy source and compare to other external measures of aggregate emissions. In Figure A.3, we calculate that total CO_2 emissions in Prowess amount to 467 MT in 2009. By comparison, the recently constructed World Input Output Database (WIOD) database reports total CO_2 emissions estimates from the same 12 industry groupings in Prowess to equal 586 MT of CO_2 in the same year, so we estimate that Prowess covers 80% of of manufacturing-based emissions in WIOD. This is a reasonable figure since Prowess contains most of the large manufacturers in India. Second, we compare emission intensities by industry in Prowess to the WIOD data in Figure A.4, and find strong agreement in most cases. Third, we cross-check the firm-level energy reports against the product-specific energy reports (aggregated to the firm level) in Figure A.5 and find that the separate reports yield a consistent picture of firm-level emissions.

Finally, even though aggregate and firm-level emissions seem to match external checks, one might be concerned that the product-specific energy intensity estimates are based not on actual energy consumption, but some convenient heuristic employed by the firm. The concern is that perhaps the cost of measuring product-specific energy intensity is prohibitively high. If production of distinct outputs occurs on the same site on the same machines, then it might be quite difficult for a firm to assign energy use to each production process. In this case, a likely candidate explanation for how the firm generates product-specific energy statistics is that they divide total energy consumption by the revenue shares of each product, not the actual energy use. Under this hypothesis, the aggregate of product-specific energy use would match the firm-level energy use, but the product-specific reports would still not reflect true emission intensity. In Figure A.6, we test for this behavior by estimating the correlation between energy shares in the product-specific data and revenue share of different products in the output data. Under the alternative hypothesis that firms merely divide energy share according to the revenue share, the correlation between these variables should be 1. In Appendix A.2 we document that we can reject this hypothesis.

Preliminary Correlations

Previous research into trade's impact on firms' environmental performance either reports cross-sectional correlations between exporting and emission intensity (Holladay, 2010; Forslid, Okubo, and Ulltveit-Moe, 2011; Cui, Lapan, and Moschini, 2012), or connects changes in tariffs to changes in firm-level emission intensity over time, without confirming that impacts channel through the export decision (Gutiérrez and Teshima, 2011; Martin, 2012; Cherniwchan, 2013). With our detailed emissions and export panel data from Prowess, we can test both the cross-sectional and the within-firm correlations. While these estimates should not be interpreted as causal, they serve as a useful benchmark for subsequent IV estimation. Additionally, the within-firm results are novel, since previous work lacks either the panel structure or the export data.

We begin with the cross-section. Previous research has found that exporters have lower emission intensity than non-exporters, which is consistent with a technology upgrading model in which firms adopt more-efficient, cleaner technology when they start exporting (Holladay, 2010; Forslid, Okubo, and Ulltveit-Moe, 2011; Cui, Lapan, and Moschini, 2012). Using the firm-level dataset, we can test for this relationship in Prowess by estimating

$$EV_{iyt} = \gamma_1 W_{iyt} + \beta_1 X_{iyt} + \epsilon_{iyt}, \tag{1.20}$$

where EV_{iyt} is the log emission intensity in value of firm i operating in industry y in year t, X_{iyt} denotes the exports of firm i, which are measured either in export share (from 0 to 1), export dummy (taking the value 1 if the firm exports positive value in the year and 0 otherwise) or log export value. W_{iyt} are controls such as year fixed effects and industry fixed effects, and ϵ_{iyt} is an idiosyncratic error term.

Results are reported in Table 1.3. In columns 1 and 4, we find that exporters are indeed cleaner than non-exporters. Column 1 reports the unconditional correlation, while column 4 reports estimates that control for primary industry and year fixed effects. Standard errors are clustered at the firm-level. In column 1, we find that exporters have 31% lower emission intensity in value than non-exporters. The point estimate is statistically significant at the 1% level. Controlling for industry and year effects, we find that the point estimate attenuates somewhat (falling to -0.20), though exporters are still significantly cleaner than non-exporters. In columns 2 and 5, we replicate the analysis using export share (from 0 to 1) as the measure of export intensity, and in columns 3 and 6, we take the log of export value. In all specifications, greater export participation is associated with lower emission intensity in value.

According to the model, the correlation documented in Table 1.3 could owe to many factors. First, exporters tend to be more efficient than non-exporters (Bernard and Bradford Jensen, 1999), so one would expect that they are also more efficient at transforming energy (and hence, CO_2 emissions) into outputs. Second, exporters earn a different price on their goods than non-exporters, since they sell in different markets. If they command a higher price on average, then emission intensity in value would fall mechanically for exporters due to the price effect. Third, exporters tend to sell a wider range of products than non-exporters. If follow-on products tend to be cleaner to produce than core products, this might also generate a lower emission intensity in value for exporters. These possibilities militate in

Table 1.3: Exporting and Emission Intensity in the Repeated Cross-Section

Dep Var:	Log (E/V)					
	(1)	(2)	(3)	(4)	(5)	(6)
Exporter	-0.31***			-0.20***		
	(0.04)			(0.04)		
Export Share		-0.80***			-0.70***	
		(0.08)			(0.08)	
Log Export Value			-0.05***			-0.03***
			(0.01)			(0.01)
Year FE	N	N	N	Y	Y	Y
Industry FE	N	N	N	Y	Y	Y
Num of Obs	35665	26381	19359	35665	26381	19359
R squared	0.01	0.02	0.01	0.10	0.10	0.09
Mean Dep. Var	3.11	3.05	2.97	3.11	3.05	2.97

Notes: All regressions include years 1995-2011. Top and bottom 1% of emission intensity values have been dropped. Observations have also been dropped if emission intensity in value is several orders of magnitude above or below the rest of the observations for the same firm. Firms are assigned to an industry based on the product that accounts for the greatest aggregate sales over the period. Firms Standard errors are clustered on the firm. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

favor of a fixed effect model to control for time-invariant differences between exporters and non-exporters.

Exploiting the panel dimension of our data, we estimate:

$$Y_{it} = \alpha_i + \alpha_t + \gamma_2 W_{it} + \beta_2 X_{it} + \epsilon_{it}, \tag{1.21}$$

where Y_{it} is firm-year outcome and X_{it} is exporting activity, ϵ_{it} is an idiosyncratic error term. Controls include year fixed effects, firm fixed effects, and firm-specific energy prices. Columns 1-4 of Table 1.4 present impacts on firm-level emission intensity in value. The impact of exporting is identified from within-firm changes in export value over time. Whether we measure export activity as the log of export value (columns 1-2) or the export share in revenue (column 3-4), we find that exporting is associated with lower emission intensity in value, as in the cross-section. Point estimates are consistently negative and statistically significant at the 1% level. With the log-log specification in column 1, the point estimate is directly interpretable as an elasticity. We find that a 1% increase in export value is associated with 0.14% lower firm emission intensity. These results confirm that, not only are exporters cleaner, but firms become cleaner (in terms of E_i/V_i) when they export more. These results are in line with work from Gutiérrez and Teshima (2011); Martin (2012); Cherniwchan (2013), but previous estimates were based on changes in tarriffs. I.e., Gutiérrez and Teshima (2011); Martin (2012); Cherniwchan (2013) did not observe exports of the firm. Here, we can see directly that exporting correlates with lower emission intensity.

Furthermore, using our detailed production data, we can distinguish impacts on E_i/V_i from impacts on E_i/Q_i. The model predicts that exporting also impacts prices (see prediction

Table 1.4: Exporting and Emission Intensity Within the Firm

Dep Var:	Log (E/V)				Log (E/Q)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Export Value	-0.14***	-0.12***			-0.09***	-0.06***		
	(0.02)	(0.02)			(0.02)	(0.02)		
Export Share			-0.20**	-0.32***			-0.03	-0.19
			(0.10)	(0.12)			(0.10)	(0.15)
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y	Y	Y
Includes MP	Y	N	Y	N	Y	N	Y	N
Num of Obs	12255	4767	17236	6689	12170	4640	17148	6553
R squared	0.09	0.09	0.03	0.06	0.05	0.05	0.02	0.03
Mean Dep. Var	2.89	2.82	3.01	2.97	-1.04	-1.35	-0.99	-1.13

Notes: All data aggregated to the firm-year level. Sample includes only firms that export some positive value over the period. Odd-numbered columns include single and multi-product firms ("MP"), while even-numbered columns include only single-product firms. Sample includes only firms that report outputs in consistent units across products and over time. Columns 1-4 report estimates for emission intensity in value (E_i/V_i), while columns 5-8 report estimates for emission intensity in quantity (E_i/Q_i). All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the firm. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

4), which would mechanically influence emission intensity in value without any change in emission intensity in quantity. To control for prices, in columns 5-8, we re-estimate (1.21) taking emission intensity *in quantity* as the outcome variable. When estimating impacts in quantities, it is necessary that output units are constant within the firm both over time and across products. Thus, in columns 5-8, we restrict the sample to firms for which output units are constant across all products over the period. For comparison, we make the same restriction in columns 1-4. In column 5, we find that emission intensity in quantity at the firm-level also falls with exports: a 1% increase in export value lowers E_i/Q_i by 0.09%. The estimate is statistically significant at the 1% level. The point estimate using export share is also negative (column 7-8), though statistically insignificant. Comparing the results in columns 1 and 5, we find that the impact on emission intensity in quantity (β_2 = -0.09) is smaller in magnitude than the impact on emission intensity in value (β_2 = -0.14). We can reject the null hypothesis of equality between the two coefficients with a p-value ¡0.01. This finding is consistent with the hypothesis that firms charge a higher price on the export market, so that when they export more, emission intensity in value falls mechanically.

Estimates in odd-numbered columns of Table 1.4 are based on the pooled sample of both multi-product and single-product firms. Thus, estimates could be influenced by endogenously changing product shares within the firm. While the firm-level dataset is not well-suited to address product-mix (since we do not know emission intensity by product in this dataset), we can indirectly address the question by conditioning on single-product firms. If firms only produce a single product throughout the period, then product-mix does not change (by construction), so could not influence emission intensity. We return to distinguishing product-mix from technology later with the product-specific dataset; but here we find preliminary evidence using the firm-level energy reports.

Table 1.5: Exporting and Prices and Core Share Within the Firm

Dep Var:	Log (UV)				Core Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Export Value	0.06***	0.05***			0.00	
	(0.01)	(0.01)			(0.00)	
Export Share			0.15**	0.09		-0.05
			(0.06)	(0.10)		(0.04)
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y
Includes MP	Y	N	Y	N	Y	Y
Num of Obs	12091	4609	17073	6534	7630	10720
R squared	0.07	0.11	0.06	0.09	0.14	0.15
Mean Dep. Var	-3.97	-4.17	-4.02	-4.08	0.64	0.66

Notes: All data aggregated to the firm-year level. Sample includes only firms that export some positive value over the period. Columns 1 and 3 include single and multi-product firms ("MP"), columns 2 and 4 include only single-product firms, and columns 5 and 6 include only multi-product firms. Sample includes only firms that report outputs in consistent units across products and over time. Columns 1-4 report estimates for average unit value (total sales over total production), while columns 5-6 report estimates for the "core share" of production - the share of sales devoted to the highest-sales product (0 to 1). All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the firm. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

We report results in even-numbered columns in Table 1.4. In column 6, we find that emission intensity in quantity at the firm-product level also falls with exporting: a 1% increase in export value lowers emission intensity in quantity 0.06%. The point estimate is statistically significant at the 1% level. Column 6 rules out both price effects and product-mix by construction, so the change in emission intensity is purely attributable to a technological change. As the estimates in column 5 derives from the full sample, and thus combines both technological and product-mix effects, a comparison of column 6 to column 5 reveals the product-mix effect. Since the point estimate in column 6 (β_2 = -0.06) is smaller in magnitude than the one in column 5 (β_2 = -0.09), it suggests that the product-mix effect skews the firm towards cleaner products. This could happen if the core products tend to be cleaner and the firm sells more of the core product when it increases exports. Additionally, comparing columns 1 and 3, we also find that the impact on E_i/V_i is smaller in magnitude when the product-mix channel is ruled out. Both comparisons point to a negative impact on emission intensity via the product-mix channel.

Finally, we examine the price and product-mix channel directly by regressing unit values (a proxy for price) and share of revenue derived from core product ("core share") on the export measures. We report results in Table 1.5. In the first two columns, we find that prices rise with exports both at the firm-level (column 1) and the firm-product level (after restricting to single-product observations) (column 2). This is precisely what we expect from the comparison of columns 1 and 5 in Table 1.4: higher export prices mechanically inflate the denominator of E_i/V_i. The result also holds when using export share as the measure of export orientation (columns 3-4), though the point estimate for single-product observations is statistically insignificant (column 4). Finally, in columns 5-6 we regress the core share (our proxy for product-mix) on the endogenous export variables. Here, we cannot reject the null of no impact on the share of production devoted to the core product. We will return to the question of product-mix in the next section when we address endogeneity in the export decision.

1.4 Identification Strategy

Tables 1.3 and 1.4 confirm the endogenous relationship of interest: exporting is associated with lower emission intensity. The result holds in the cross-section, as well as within the firm over time. Prices also increase with exports, so part of the effect on emission intensity in value is due to prices, but even when emission intensity is measured in quantity, firms appear to improve environmental performance when they export more. While these correlations support the model predictions, our ability to infer causality from the OLS estimates are confounded by the fact that exporting is an endogenous decision that could correlate with many other unobserved factors that also determine emission intensity. To establish the causal connection between exporting and emission intensity, in this section we compute product-code-specific instruments for exporting from aggregate trade statistics. We then merge to Prowess and report evidence in support of the identification assumption.

Exogenous Variation in Exports

The primary endogeneity concern with (1.21) is that time-varying omitted variables could impact both the emission intensity of production and the export decision of the firm (see above). A secondary concern stems from observations with zero export value. Since export volumes cannot be negative, the distribution is censored on the left. This generates selection bias if unobservable determinants of export participation correlate with export volume. With exogenous determinants of exporting, and information on the export decision, we can correct for selection using a two-step Heckman procedure, in which export participation in a given year is first predicted via probit estimation on the exogenous regressors, and then the inverse mills ratio is included in the export volume regression. This procedure, while common in the trade literature (Harrigan, Ma, and Shlychkov, 2011), obviously requires data on the export decisions of firms.

To generate exogenous determinants of exporting, we follow recent papers that exploit fluctuations in import demand of foreign trading partners to predict firm-level exports. Hummels et al. (2014); Bernard, Moxnes, and Ulltveit-Moe (2014) show that import demand of foreign trading partners from third-party exporters correlates with firm-level exports in the country of study. That is, as foreign trading partners increase imports from third-party countries, they tend also to increase imports from firms in the studied country. Changes in import demand can be attributed to sector specific income shocks, which are likely exogenous to time-varying unobservable determinants of firm-level productivity in foreign countries. Thus, foreign import demand from countries other than India provide an exogenous source of export variation for Indian firms, with which we can identify the causal impact on emission intensity. The strategy is to compute product-code-specific import demand shocks in destinations served by India (from countries other than India) and then aggregate shocks across the different destinations. This procedure delivers product-code-specific demand shocks that can be matched to firm-level export volumes in Prowess.

Formally, we define total import demand for a destination d in product j in time t as $D_{djt} = \sum_{o \in \Delta_o} D_{odjt}$, where D_{odjt} is the import value from origin o to import destination d in product j and year t, and Δ_o is the set of all origin countries that export to d in year t other than India. Aggregating across destinations, we compute:

$$\widetilde{D}_{jt} = \sum_{d \in \Delta_d} x_{dj0} D_{djt}, \tag{1.22}$$

where export shares $x_{dj0} \equiv \frac{X_{dj0}}{\sum_{d \in \Delta_d} X_{dj0}}$ with X_{dj0} representing Indian exports in product j to destination d in base year $t = 0$, and Δ_d is the set of destinations importing j from India in year $t = 0$. We use base-year shares to weight the demand shocks because current-year export shares may respond to time-varying omitted variables like cost shocks, while base-year shares likely do not.[10] The weighted-average demand shocks $\widetilde{D}_{jt}$ represent time-varying idiosyncratic shocks to demand for exports for Indian firms producing in the different industries j. We also separate demand shocks for low-income (*LI*) and high-income (*HI*) countries to separately estimate impacts depending on destination environmental regulation (with income proxying for regulation):

$$\widetilde{D}^{HI}_{jt} = \sum_{d \in \Delta_d} x_{dj0} D_{djt} \mathbb{1} \left[d \in \{HI\} \right] \tag{1.23}$$

$$\widetilde{D}^{LI}_{jt} = \sum_{d \in \Delta_d} x_{dj0} D_{djt} \mathbb{1} \left[d \in \{LI\} \right] \tag{1.24}$$

where destinations d belong to either the "high income" set or the "low income" set, which are both held constant over time.

To compute the demand shocks, we use the CEPII BACI dataset of bilateral national trade flows. The data is reported at the six-digit harmonized system (HS6) level for the years 1995-2011 for 240 countries.[11] For each of the 5,108 HS6 codes, in each year, we compute $\widetilde{D}_{jt}$,

[10] The shares x_{dj0} may still be endogenous to firm unobservables. For example, Melitz (2003) predicts that only the most efficient firms can serve distant and more competitive markets; but to the extent that these differences are time-invariant, they are captured by product or firm-level fixed effects in the regressions.

[11] Years 1995-1997 are classified according to the 1992 revision, while 1998-2011 are classified according to 1996 revision. We utilize the UN mapping to convert between the two. http://unstats.un.org/unsd/trade/conversions/HS%20Correlation%20and%20Conversion%20tables.htm

Table 1.6: Descriptive Statistics of Weighted Average Demand Shocks

	Billions USD				N	# Countries
	Mean (1)	St. Dev (2)	Min (3)	Max (4)	(5)	(6)
Total Demand ($\widetilde{D}_{jt}$)	0.057	0.416	0	41.32	86,836	222
High Income Demand ($\widetilde{D}_{jt}^{HI}$)	0.053	0.402	0	41.32	86,836	47
Low Income Demand ($\widetilde{D}_{jt}^{LI}$)	0.004	0.043	0	6.94	86,836	175

Notes: Statistics derived from CEPII BACI dataset. Total demand ($\widetilde{D}_{jt}$) indicates weighted average demand of India's trading partners in a given HS6-year, where weights correspond to 1995 India export shares. "High Income Demand" and "Low Income Demand" assign positive weights only to those destinations in the designated income group, as defined by the world bank categories. All three measures cover 5,108 HS6 categories over the years 1995-2011.

$\widetilde{D}_{jt}^{HI}$, and $\widetilde{D}_{jt}^{LI}$ taking 1995 as the base year ($t = 0$). I.e., in each year, we aggregate income shocks over destinations using 1995 Indian export shares as weights. Income groups conform to the World Bank Development Indicator categories. The list of high-income destinations are reported in Table A.2.[12]

Table 1.6 reports descriptive statistics for $\widetilde{D}_{jt}$, $\widetilde{D}_{jt}^{HI}$, and $\widetilde{D}_{jt}^{LI}$ based on imports of 222 countries. We identify 47 countries as *HI*, and the remaining 175 as *LI*. Average world demand across the 5,108 HS6 codes and 17 years equals 0.057 Billion USD, with almost all coming from high-income imports (column 1). Table 1.7 investigates variation in demand shocks across HS6 codes. For each measure of demand ($\widetilde{D}_{jt}$, $\widetilde{D}_{jt}^{HI}$, and $\widetilde{D}_{jt}^{LI}$), we compute within each HS6 the high-low spread and total growth over the period, and the year-on-year percentage growth. High-low spreads are computed as the difference between the highest value and the lowest value over the period, normalized by the mean value over the period. Total growth is the percentage change between 2011 and 1995. We report mean and standard deviations (below in parenthesis) by quartile and overall. The overall average growth in total weighted average demand is 1.97, or nearly triple. Variation within the HS6 is even higher if we look at the normalized high-low spread. Comparing the minimum to the maximum, the difference is more than 4x. There is also substantial variation in these growth rates across HS6 categories. In terms of total growth, the bottom quartile (lowest values) shrank on average by 63%. By contrast, the top quartile grew 675%. There is also substantial variation in year-to-year growth. Demand can fall by nearly 50% (bottom quartile average) or more than triple (top quartile average) in a year. These patterns are qualitatively the same when we break out demand by destination income. All together, the statistics in Table 1.6 and 1.7 indicate significant variation in growth across HS6 categories, which will differentially impact Indian firms operating in different product-codes and aid in identifying the causal impact of exporting.

[12] The World Bank defines a set of OECD and non-OECD countries as "high income." For our measure of HI, we take the union of the two sets. All other countries are classified as LI.

Table 1.7: Variation in Demand Shocks

	Quartile				
	IV	III	II	I	All
Total Demand ($\widetilde{D}_{jt}$)					
High-Low Spread	0.92	1.60	2.56	7.34	3.11
	(0.24)	*(0.20)*	*(0.42)*	*(3.97)*	*(3.22)*
Total Growth	-0.63	0.30	1.44	6.75	1.97
	(0.26)	*(0.28)*	*(0.40)*	*(12.56)*	*(6.91)*
Year-over-Year Growth	-0.46	-0.04	0.13	2.19	0.46
	(0.27)	*(0.05)*	*(0.05)*	*(5.51)*	*(2.94)*
Total Demand High Income ($\widetilde{D}_{jt}^{HI}$)					
High-Low Spread	0.89	1.54	2.42	6.64	2.88
	(0.22)	*(0.19)*	*(0.38)*	*(3.72)*	*(2.92)*
Total Growth	-0.58	0.35	1.47	7.63	2.23
	(0.27)	*(0.27)*	*(0.39)*	*(25.39)*	*(13.10)*
Year-over-Year Growth	-0.42	-0.03	0.12	1.45	0.28
	(0.26)	*(0.05)*	*(0.05)*	*(3.09)*	*(1.71)*
Total Demand Low Income ($\widetilde{D}_{jt}^{LI}$)					
High-Low Spread	1.49	2.42	3.78	9.57	4.32
	(0.35)	*(0.26)*	*(0.62)*	*(3.94)*	*(3.73)*
Total Growth	-0.62	0.64	2.43	15.76	4.55
	(0.30)	*(0.40)*	*(0.69)*	*(27.87)*	*(7.57)*
Year-over-Year Growth	-0.57	-0.06	0.22	5.32	1.24
	(0.25)	*(0.08)*	*(0.09)*	*(14.07)*	*(7.47)*

Notes: This table describes the variation within and across HS6 categories in weighted average total world demand, high-income demand, and low-income demand. Each cell reports the mean for the quartile or the overall mean with standard deviations listed below. "High-Low Spread" is the difference between the highest and lowest value within HS6 over the period, normalized by the mean value over the period. "Total Growth" is the growth rate between 2011 and 1995 within the HS6. "Year-over-year Growth" is the percentage change from one year to the next within the HS6. The top and bottom 1% of year-over-year growth rates are excluded.

Mapping to Prowess

Finally, having computed product-specific trade shocks, we merge the aggregate statistics to Prowess to test for firm-level impacts. Merging trade data to Prowess is problematic because CMIE classifies products according to its own 16-digit codes, which do not map directly to any other classification system at a disaggregated level. Previous researchers map from HS trade data to 4-digit National Indian Classification (NIC) through the Debroy and Santhanam (1993) mapping, and then to the 16-digit Prowess ID codes via a NIC mapping supplied by CMIE (e.g., De Loecker et al. 2012; Goldberg et al. 2010; Chakraborty 2012). This strategy does not fit our purpose because the 4-digit NIC codes are fairly coarse, and thus obscure much of the variation at the HS6 level in demand shocks.[13] To exploit the detailed variation in the Prowess data, we generate our own mapping that connects the Prowess ID code directly to HS trade classifications. This mapping allows for a tighter link between (HS) product-specific shocks and production activity in Prowess firms. As more researchers are increasingly interested in analyzing the Prowess data, our mapping represents another contribution of the paper. The details can be found in Appendix A.2.

Testing the Parallel Trends Assumption

The identification strategy is to compare the trend in emission intensity for firms that operate in product-codes that see large demand changes compared to firms that do not. If the firms follow common trends before the trade shocks, and if HS6-specific demand shocks abroad are uncorrelated with unobservable determinants of firm-level emission intensity in Prowess, then demand shocks identify the causal impact of exporting on emission intensity. While the latter condition is in principle untestable, the former can be evaluated by regressing prior trends in emission intensity on future demand shocks, if pre-period data is observed. Since the production data spans the years 1990-2011, and the trade data starts in 1995, we can test for correlations between emission intensity over the period 1990-1995 and future trade shocks without compromising statistical power in the IV estimation. The strategy is similar to the one proposed by Topalova and Khandelwal (2011) in which the authors test for correlation between changes in tariffs and industry characteristics.

In particular, we regress the percentage change in emission intensity in firm-products in the product-specific dataset on the percentage change in demand shocks $\widetilde{D}_{pt}$, where $\widetilde{D}_{pt}$ is the import shock faced by Prowess product-category p, after passing the HS6 shocks $\widetilde{D}_{jt}$ through the mapping. We thus estimate

$$\Delta_{t,r} Y_{ip} \quad = \quad \alpha_0 + \alpha_1 \Delta_{s,t} \widetilde{D}_p + \epsilon_{ip} \tag{1.25}$$

where $\Delta_{t,r} Y_{ip}$ is the percentage change from year $r < 1995$ to year $t = 1995$ in firm-product emission intensity, α_0 a constant, and $\Delta_{s,t} \widetilde{D}_p$ is the percentage change in the trade shock in product p between year $t = 1995$ and future year $s > 1995$. We compute $\Delta_{t,r} Y_{ip}$ at different intervals depending on whether we observe the firm-product in the base year r or not. Results are reported in Table 1.8.

[13] Additionally, the Debroy and Santhanam (1993) mapping relates much older revisions of the NIC and the HS, and thus must be passed through several other mappings (between newer and older revisions of the HS) to generate usable translations.

Table 1.8: Testing the Parallel Trends Assumption

	Δ Shock 95-96					Δ Shock 96-97				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ E/Q 90-95	0.04					-0.13*				
	(0.03)					(0.07)				
Δ E/Q 91-95		0.04					-0.17			
		(0.03)					(0.10)			
Δ E/Q 92-95			0.03					-0.06		
			(0.03)					(0.09)		
Δ E/Q 93-95				0.03					-0.03	
				(0.02)					(0.08)	
Δ E/Q 94-95					0.04					0.02
					(0.04)					(0.11)
Num of Obs	625	870	1111	1483	1838	615	855	1094	1465	1821
R squared	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00

Notes: This table tests for differential trends in emission intensity in the years prior to 1995 for firm-products in product-codes that see larger foreign demand shocks after 1995. Dependent variable is listed in the left-most column, with the independent variable reported above the column numbers. Each cell reports the point estimate α_1 from regressing the percentage change in emission intensity between the years specified in the left-most column on the year-over-year percentage demand shock reported above. The percentage change in weighted average demand shocks between years 1995 and 1996, $\frac{\widetilde{D}_{j,1996}-\widetilde{D}_{j,1995}}{\widetilde{D}_{j,1995}}$ and 1996 and 1997, $\frac{\widetilde{D}_{j,1997}-\widetilde{D}_{j,1996}}{\widetilde{D}_{j,1996}}$, respectively. Emission intensities are specified in quantity of output at teh firm-product level, based on the product-specific dataset. Top and bottom 1% of emission intensity values have been dropped. All regressions include industry fixed effects. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

In columns 1-5, the independent variable is the year-over-year percentage change in foreign demand between 1995 and 1996, $\Delta_{1996,1995}\widetilde{D}_p$. Each row reports the regression coefficient α_1 from estimating (1.25) taking $\Delta_{1995,r}Y_{ip}$ as the dependent variable for different values of r. For all the firm-products that we observe in both 1990 and 1995, we compute $\Delta_{1995,1990}Y_{ip}$ and report α_1 in the top row. For all the firm-products that we observe in both 1991 and 1995, we compute $\Delta_{1995,1991}Y_{ip}$ and report α_1 in the second row, and so on. In all columns 1-5, we find that the correlation between future shocks $\Delta_{1996,1995}\widetilde{D}_p$ and pre-1995 trends are not statistically significant. In columns 6-10, we repeat the exercise taking $\Delta_{1997,1996}\widetilde{D}_p$ as the future shocks. Again, the correlations are not statistically significant. We take these regressions as confirming evidence that the parallel trends assumption holds.

1.5 Results

In this section, we analyze the relationship between emission intensity, firm exports, and exogenous foreign demand shocks. When comparing emission intensity in quantity to emission intensity in value, we use the firm-level dataset, while we use the product-specific dataset to estimate technological impacts directly. Using the two datasets in tandem, we can separate the different component channels of equation (1.1).

Firm-level Evidence

We begin with an analysis of the firm-level dataset. In the previous section, we computed foreign demand shocks at the HS6 product level and mapped to Prowess codes via the concordance described in Appendix A.2. Following a standard approach from the literature (e.g., De Loecker (2011); De Loecker et al. (2012); Bernard, Redding, and Schott (2011)), we then compute firm-level demand shocks by averaging over the shocks experienced by different products within the firm, $\widetilde{D}_{it} = \sum_{p \in \Delta_i} s_{p0} \widetilde{D}_{pt}$; where s_{p0} corresponds to the share of firm i's total value in product p in base year $t = 0$ – the first year that the firm appears in the dataset – Δ_i is the set of goods produced by firm i. As in the aggregate exports, while year-to-year changes in s_{pt} may be endogenous to demand shocks $\widetilde{D}_{pt}$, base year shares likely are not.

With foreign demand shocks computed at the firm-level, we estimate the first-stage impact on firm-level exports:

$$\text{Log } X_{it} \quad = \quad \alpha_i + \alpha_t + \beta_3 \text{Log } \widetilde{D}_{it} + \gamma_3 W_{it} + \epsilon_{it} \tag{1.26}$$

where X_{it} is export value of firm i in year t and W_{it} is a vector of controls including firm-specific energy prices. Since exports are censored at zero, we also include the inverse mills ratio in W_{it}, which is recovered from probit estimation of export participation in a given year.

How should we expect $\widetilde{D}_{it}$ to impact X_{it}? Recall that an increase in foreign income could have two effects. On one hand, the foreign market consumes more of all varieties, including those from India, which pushes Log X_{it} up. On the other hand, increased foreign demand drives up competition, which hurts less-productive suppliers, and could push Log X_{it} down. The coefficient β_3 captures the average net effect. Using the universe of Danish and Norwegian exporters respectively, Hummels et al. (2014) and Bernard, Moxnes, and Ulltveit-Moe (2014) both estimate a positive elasticity with respect to foreign demand shocks, which suggests that the direct effect dominates. Since Prowess is a sample of large firms, for which we expect the direct effect to dominate (prediction 1), we also expect $\beta_3 > 0$.

Next, we estimate the causal impact of exporting on emission intensity

$$\text{Log } Y_{it} = \alpha_i + \alpha_t + \beta_4 \widehat{\text{Log } X}_{it} + \gamma_4 W_{it} + \epsilon_{it} \tag{1.27}$$

where Y_{it} is either emission intensity in value or in quantity, and $\widehat{\text{Log } X}_{it}$ are instrumented exports from the first-stage, and W_{it} is as above. Having purged X_{it} of the influence of unobservable co-determinants of export and production efficiency in (1.26), $\widehat{\text{Log } X}_{it}$ represents the exogenous portion of X_{it} resulting from foreign demand shocks and can be taken as orthogonal to the error term ϵ_{it}. Thus, β_4 can be recovered via OLS. When Y_{it} is emission intensity in value E_i/V_i, β_4 includes price effects, product-mix effects, and technological effects. When Y_{it} is emission intensity in quantity E_i/Q_i, β_4 includes just product-mix effects and technology. If prices respond to foreign demand shocks, we would expect the two estimates to differ. In particular, if the export price is higher than the domestic prices, we would expect β_4 to be more negative when emission intensity is measured in values.

Finally, we also break out foreign demand shocks by destination market income:

$$\text{Log } X_{it} \quad = \quad \alpha_i + \alpha_t + \beta_5 \text{Log } \widetilde{D}_{it}^{HI} + \beta_6 \text{Log } \widetilde{D}_{it}^{LI} + \gamma_5 W_{it} + \epsilon_{it} \tag{1.28}$$

The point of breaking out demand shocks by destination is to see how price and product-mix effects vary with the income of the importing market. The model predicts that the export price is higher for high-regulation destinations (see prediction 4). Since high regulation countries also have higher income (Copeland and Taylor, 2003), we take income as a proxy for environmental regulation. The price effect should be comparatively stronger and more negative for high income countries. We can test for this asymmetry across destination markets by estimating separate coefficients.

Results are reported in Table 1.9. Columns 1-3 report impacts on E_i/V_i while columns 4-6 report impacts on E_i/Q_i. We restrict the sample to firms with constant unit reporting throughout the period so that emission intensities in quantity are comparable over time. Units need not be constant *across* firms, because as long as units do not vary over time or across products within the firm, the influence of differential unit reporting will be captured by the firm fixed-effect α_i. We also exclude the top and bottom 1% of emission intensities in value and quantity. With these restrictions, we have 1587 firms and 9008 firm-year observations.

Panel B reports the first stage (equation 1.26). In column 1, we find that a 1% increase in foreign demand shocks increases export value at the firm-level by 0.17%. The point estimate is statistically significant at the 1% level, and the F-stat for joint significance of all independent regressors is 8.70. Standard errors here, and throughout, are clustered at the HS6 level to allow for arbitrary correlation in the error term within product-code (potentially across firms). To interpret the magnitude of this point estimate, we multiply by the median year-over-year percentage growth of an HS6 product code (3.4%), and find for the median product code, demand shocks are responsible for $3.4 * 0.17 = 0.6\%$ export growth year-to-year. For comparison, median year-over-year export value growth over the period was 4.2%, so foreign demand shocks explain about 14% of the median growth rate. Column 2 breaks out demand shock impacts by market income (as a proxy for environmental regulation). In Panel B, we find that most of the impact on exporting is channeled through high-income demand shocks. The point estimate on $\widetilde{D}_{it}^{HI}$ is estimated to be 0.17, statistically significant at the 1% level, while the coefficient on $\widetilde{D}_{it}^{LI}$ is only 0.03, and statistically insignificant.

Panel A reports the second-stage impact of exporting on emission intensity. Columns 1 and 2 instrument the export decision with either $\widetilde{D}_{it}$ (column 1) or the disaggregated shocks $\widetilde{D}_{it}^{HI}$ and $\widetilde{D}_{it}^{LI}$ (column 2). Column 3 reports the endogenous (uninstrumented) correlation for the same set of firm-years. In column 1, we find that a 1% increase in exporting lowers E_i/V_i by 0.57%. Compared to column 3, the IV impact is significantly larger in magnitude than the OLS estimate. If reverse causality or omitted variables biased the OLS estimates, we would expect that the absolute value of the point estimate in column 3 would be larger than in column 1. The fact that it is not suggests that measurement error is a bigger problem for the OLS than endogeneity. To interpret the magnitude of the IV estimate, consider that the average exporter in the sample earns 9.7% of its revenue from exporting. If exports double, with no change in its domestic revenues, the export share would increase from 9.7% to 17.7%. Applying the point estimate from column 1, this increase in exports generates a 57% decline in emission intensity in revenue. A similar reduction in emission intensity applied to the median firm-year observation would relocate the firm to the first quartile (lowest values) of E_i/V_i.

Table 1.9: IV Estimates of Exporting Impacts on Emission Intensity at the Firm Level

Dependent Variable:	Log (E/V)			Log (E/Q)		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Second Stage						
Log Export Value	-0.57***	-0.47***	-0.09***	-0.38***	-0.33***	-0.05***
	(0.10)	(0.07)	(0.01)	(0.09)	(0.07)	(0.01)
Panel B: First Stage						
Log $\widetilde{D}_{it}$	0.17***			0.17***		
	(0.04)			(0.04)		
Log $\widetilde{D}_{it}^{HI}$		0.17***			0.17***	
		(0.03)			(0.03)	
Log $\widetilde{D}_{it}^{LI}$		0.03			0.03	
		(0.03)			(0.03)	
R^2	0.086	0.089		0.086	0.089	
F-stat	8.70	9.96		8.70	9.96	
Panel C: Reduced Form						
Log $\widetilde{D}_{it}$	-0.10***			-0.06***		
	(0.02)			(0.02)		
Log $\widetilde{D}_{it}^{HI}$		-0.08***			-0.07***	
		(0.02)			(0.02)	
Log $\widetilde{D}_{it}^{LI}$		0.03			0.00	
		(0.03)			(0.01)	
R^2	0.057	0.057		0.035	0.035	
Selection Correction	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y
Num of Obs	9008	9008	9008	9008	9008	9008
Num of Firms	1587	1587	1587	1587	1587	1587

Notes: Panel A reports the second stage impact of exporting on emission intensity, Panel B reports the corresponding first-stage impacts of demand shocks on log export value, and Panel C reports the reduced form. Columns 1-2, 4-5 report instrumental variable estimates, while columns 3 and 6 report the OLS. All regressions control for the inverse mills ratio of exporting in a given year. All data aggregated to the firm-year level. Demand shocks are averaged over products produced within the firm using base year product sales shares (first year the firm appears in the dataset). Emissions values computed from firm-level energy reports. Sample includes only firms that export some positive value over the period and firms that report outputs in consistent units across products and over time. Columns 1-3 report estimates for emission intensity in value (E_i/V_i), while columns 4-6 report estimates for emission intensity in quantity (E_i/Q_i). All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the HS6 category of the core product in the base year (first year the firm appears in the dataset). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Next, in columns 4 and 6, we estimate the IV and the OLS impacts on E_i/Q_i. While price effects could mechanically influence estimates in columns 1-3, columns 4-6 net prices out. In column 4, we find that a 1% increase in export value from exogenous foreign demand shocks lowers emission intensity in quantity 0.38%. Compared to the OLS estimate on the same sample (column 6), the magnitude of the IV is again larger, which again implies that endogeneity bias is not a major concern. The same counterfactual as above (increasing export share from 9.7% to 17.7%) would generate a 38% fall in emission intensity in quantity. While this is a large reduction in emission intensity, it is 50% smaller in magnitude than the estimate for E_i/V_i. Controlling for output price makes a large difference. We can reject equality of the coefficients with a p-value under 1%. This result is consistent with the hypothesis that export prices are higher on average than domestic prices, so that when Indian firms export more, average price increases and E_i/V_i fall mechanically.

Panel C estimates the reduced-form, i.e. the direct impact of the foreign demand shocks on emission intensity. In column 1, we find that a 1% increase in foreign demand lowers emission intensity in value 0.1%, statistically significant at the 1% level. In column 2, we confirm that the impact is channeled exclusively through the high-income country shocks: a 1% increase in import demand from a high-income country lowers E_i/V_i 0.08%, while the same demand shock from a low-income country *increases* E_i/V_i 0.03%, and only the former is statistically significant at conventional levels.

While we find in columns 4-6 of Table 1.9 that exporting causes firms to lower emission intensity in quantity, we cannot tell from these estimates if the channel is product-mix or technological upgrading (or some mixture of the two). The estimates in Table 1.9 are based on firm-level averages for a sample that includes multi-product firms, so both channels could play a role. As before, we isolate technology from product-mix by restricting the sample to single-product firms and re-estimate equations (1.26)-(1.28). With this restriction, the sample size drops to 716 firms and 3751 firm-year observations. Results are reported in Table 1.10. Here, while we find negative impacts in the OLS on both E_i/V_i (column 3) and E_i/Q_i (column 6), we cannot reject the null of no impact in the IV in either measure (columns 1 and 4, respectively). It appears that the negative emission intensity in quantity effects at the firm-level (Table 1.9, column 4) do not survive disaggregation to the firm-product level. This suggests that *all* of firm-level effect can be attributed to product-mix. We will investigate further with the product-specific dataset to establish the null results more firmly.

Next, we explore the price and product-mix channels directly in Table 1.11. Columns 1-3 report impacts on the firm average log unit value while columns 4-6 report impacts on the core share of sales. In Panel A, we find in column 1 that the a 1% increase in exporting induced by foreign demand shocks translates into 0.19% higher unit values, statistically significant at the 1% level. Again, the impact is stronger than the OLS impact (column 3).

For the product-mix channel, we find in column 4 that an increase in exports leads to a higher sales share for the core product, statistically significant at the 1% level. To interpret the point estimate, note that the average sales share for the core product is 85%. If exports increase 10%, the point estimate in column 4 implies that the core share increases 0.26 * Log 1.1 * 100 = 2.5 percentage points, or from 85% to 87.5%. That is, a 10% increase in exports increases core share by 3%, for an elasticity of 0.3. Since emission intensity in quantity falls at the firm level with no corresponding fall in E_{ij}/Q_{ij} (Table 1.10, column 4), we can infer that the increase in core share generates the reductions in E_i/Q_i. This result is consistent

Table 1.10: IV Estimates of Exporting Impacts on Emission Intensity for Mono-product Firm

Dependent Variable:	Log (E/V)			Log (E/Q)		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Second Stage						
Log Export Value	-0.16	-0.08	-0.11***	0.01	0.09	-0.05**
	(0.11)	(0.11)	(0.02)	(0.11)	(0.12)	(0.02)
Panel B: First Stage						
Log $\widetilde{D}_{it}$	0.28**			0.28**		
	(0.11)			(0.11)		
Log $\widetilde{D}_{it}^{HI}$		0.20**			0.20**	
		(0.09)			(0.09)	
Log $\widetilde{D}_{it}^{LI}$		0.05			0.05	
		(0.05)			(0.05)	
R^2	0.068	0.068		0.068	0.068	
F-stat	6.05	5.58		6.05	5.58	
Panel C: Reduced Form						
Log $\widetilde{D}_{it}$	-0.04			0.00		
	(0.04)			(0.04)		
Log $\widetilde{D}_{it}^{HI}$		-0.02			0.00	
		(0.04)			(0.03)	
Log $\widetilde{D}_{it}^{LI}$		0.00			0.02	
		(0.02)			(0.02)	
R^2	0.051	0.050		0.027	0.027	
Selection Correction	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y
Num of Obs	3751	3751	3751	3751	3751	3751
Num of Firms	716	716	716	716	716	716

Notes: Panel A reports the second stage impact of exporting on emission intensity, Panel B reports the corresponding first-stage impacts of demand shocks on log export value, and Panel C reports the reduced form. Columns 1-2, 4-5 report instrumental variable estimates, while columns 3 and 6 report the OLS. All regressions control for the inverse mills ratio of exporting in a given year. Emissions values computed from firm-level energy reports. Sample includes only firms that export some positive value over the period and firms that report outputs in consistent units across products and over time, and only single-product firms. Columns 1-3 report estimates for emission intensity in value (E_i/V_i), while columns 4-6 report estimates for emission intensity in quantity (E_i/Q_i). All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the HS6 category of the core product in the base year (first year the firm appears in the dataset). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 1.11: IV Estimates of Exporting Impacts on Prices and Core Share at the Firm Level

Dependent Variable:	Log Unit Value			Core Share		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Second Stage						
Log Export Value	0.19***	0.14***	0.05***	0.26***	0.20***	-0.00
	(0.05)	(0.04)	(0.01)	(0.05)	(0.03)	(0.00)
Panel B: First Stage						
Log $\widetilde{D}_{it}$	0.17***			0.16***		
	(0.04)			(0.04)		
Log $\widetilde{D}_{it}^{HI}$		0.17***			0.17***	
		(0.03)			(0.04)	
Log $\widetilde{D}_{it}^{LI}$		0.03			0.03	
		(0.03)			(0.03)	
R^2	0.086	0.089		0.116	0.102	
F-stat	8.70	9.96		8.29	8.60	
Panel C: Reduced Form						
Log $\widetilde{D}_{it}$	0.03*			0.04***		
	(0.02)			(0.01)		
Log $\widetilde{D}_{it}^{HI}$		0.01			0.03***	
		(0.02)			(0.01)	
Log $\widetilde{D}_{it}^{LI}$		0.02**			0.01***	
		(0.01)			(0.00)	
R^2	0.072	0.073		0.074	0.076	
Selection Correction	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y
Num of Obs	9008	9008	9008	5257	5257	5257
Num of Firms	1587	1587	1587	871	871	871

Notes: Panel A reports the second stage impact of exporting on emission intensity, Panel B reports the corresponding first-stage impacts of demand shocks on log export value, and Panel C reports the reduced form. Columns 1-2, 4-5 report instrumental variable estimates, while columns 3 and 6 report the OLS. All regressions control for the inverse mills ratio of exporting in a given year. All data aggregated to the firm-year level. Demand shocks are averaged over products produced within the firm using base year product sales shares (first year the firm appears in the dataset). Sample includes only firms that export some positive value over the period and firms that report outputs in consistent units across products and over time. Columns 1-3 report estimates for average unit value (total sales over total production), while columns 4-6 report estimates for core share of production. All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the HS6 category of the core product in the base year (first year the firm appears in the dataset). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

with the case in which core products require lower emission intensity to produce.

In summary, evidence based on the firm-level energy reports confirm the preliminary results from section 3, but are robust to endogeneity in the export choice. Exporting lowers both emission intensity in value and emission intensity in quantity at the firm-level, but not at the firm-product level. Price effects explain about 1/3rd of the reduction in the emission intensity in value, with the rest coming from product-mix. The results imply that export market access induces firms to increase average price and substitute towards producing cleaner goods. The null result on firm-product emission intensity is based on a restricted set of single-product firms. It remains to be seen whether firm-product emission intensity adjusts to foreign demand shocks in the larger product-specific dataset, which includes firm-product observations of multi-product firms.

Product-Level Evidence

Finally, we estimate the impact of foreign trade shocks on emission intensity in quantity of individual firm-products from the product-specific dataset. This dataset includes firm-product observations from both single-product firms and multi-product firms, in contrast to the results in Table 1.10. Additionally, we no longer must restrict the sample to firms with constant units reported across products within the firm. As long as units are reported consistently within the firm-product over time, any remaining influence of differential unit reporting will be captured by the firm-product fixed effect. Lastly, with the product-specific dataset, we no longer have to average product-code demand shocks across products within the firm. Since estimation is at the firm-product level, we can use the product-code shocks $\widetilde{D}_{pt}$ directly.

We estimate

$$\text{Log } EQ_{ipt} = \alpha_{ip} + \alpha_t + \beta_7 \text{Log } \widetilde{D}_{pt} + \gamma_7 W_{it} + \epsilon_{ipt} \tag{1.29}$$

where EQ_{ipt} is the emission intensity in quantity of firm-product ip in year t. β_7 captures the technological impact of foreign demand shocks. The dependent variable is denominated in quantity, so prices effects are excluded. Also, by taking the firm-product as the unit of observation, we have ruled out the product-mix channel. If firms in fact adopt cleaner technology when foreign demand increases, we should have $\beta_7 < 0$.

Results are reported in Table 1.12. In column 1, we find that a 1% increase in foreign demand translates into 0.028% *higher* emission intensity at the firm-product level. The point estimate is statistically significant at the 10% level (p-value = .091). Standard errors have been clustered on the HS6 code as before. The estimate is based on a sample of 2,773 firms and 4,249 firm-products, which is a substantial increase in size over the single-product firm estimates in Table 1.10, thus, statistical power should be much less of an issue. With a p-value of 0.091, we can reject null that $\beta_7 < 0$ with a p-value of $0.091/2 = 0.045$. Thus, at the 5% level, we can reject that foreign demand shocks lowers emission intensity at the firm-product level. In column 4, we break out demand shocks by destination income, and again find that, if anything, foreign demand shocks *increase* emission intensity. The point estimate on high-income countries is 0.018, statistically significant at the 10% level. The p-value for the one-tailed test is 0.0365.

Table 1.12: Product Specific Data

Dep Var:	Log (E/Q)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log $\widetilde{D}_{pt}$	0.028*	0.023*	0.022			
	(0.016)	(0.014)	(0.014)			
Log $\widetilde{D}_{p,t-1}$		0.013	0.011			
		(0.011)	(0.010)			
Log $\widetilde{D}_{p,t-2}$			0.012			
			(0.008)			
Log $\widetilde{D}^{HI}_{pt}$				0.018*	0.014	0.014
				(0.010)	(0.008)	(0.008)
Log $\widetilde{D}^{HI}_{p,t-1}$					0.013	0.008
					(0.008)	(0.008)
Log $\widetilde{D}^{HI}_{p,t-2}$						0.017**
						(0.007)
Log $\widetilde{D}^{LI}_{pt}$				0.006	0.005	0.003
				(0.007)	(0.006)	(0.006)
Log $\widetilde{D}^{LI}_{p,t-1}$					0.008	0.006
					(0.005)	(0.005)
Log $\widetilde{D}^{LI}_{p,t-2}$						0.009
						(0.008)
Selection Correction	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm-Product FE	Y	Y	Y	Y	Y	Y
Num of Obs	27954	27954	27954	27954	27954	27954
Num of Firms	2773	2773	2773	2773	2773	2773
Num of Firm-Products	4239	4239	4239	4239	4239	4239
Mean Dep. Var	6.09	6.09	6.09	6.09	6.09	6.09
R squared	0.01	0.01	0.01	0.01	0.01	0.01

Notes: Emission intensity in quantity computed at firm-product level from product-specific reports. Sample includes only firms that export some positive value over the period. All regressions control for the inverse mills ratio of exporting in a given year. All regressions include years 1997-2011. Years 1995 and 1996 are excluded so that all observations have two years of lagged data. The same restriction on outliers as throughout applies. Standard errors are clustered on the HS6 category corresponding to Prowess product p. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

One explanation why do foreign demand shocks fail to translate into cleaner production techniques (at the firm-product level) in this context could be that year-to-year foreign demand shocks are not strong enough to induce firms to adjust behavior. Hummels et al. (2014) finds that year-to-year changes in foreign demand impacts the hiring and wage offerings of Danish firms, and Bernard, Moxnes, and Ulltveit-Moe (2014) finds that foreign demand impacts several sales margins of Norwegian firms. Additionally, we find that the year-to-year shocks impact exports, prices, and product-mix of the firm, so the evidence suggests year-to-year demand shocks *do* influence firm behavior. Finally, we note that the period captures many episodes of sustained demand shifts, such as the Asian financial crises and the Great Recession in the US and EU; thus, at least some of the variation in our sample is likely based on persistent shocks.

Still, perhaps technological investment requires longer lead times to adjust. Managers may need some time to decide to adjust the capital stock. In addition, securing financing can take time, especially in the developing world where credit markets are often incomplete. To capture this time dimension, we estimate the impact of one and two-year lagged shocks in columns 2-3 and 5-6, in addition to contemporaneous shocks. We find that point estimates on lagged shocks are also positive, though usually not statistically significant. These results confirm the initial findings from Table 1.10: in both sets of evidence it does not appear that emission intensity in quantity at the firm-product level declines with foreign demand shocks. This result support the theory that firms adjust emission intensity through product-mix.

Alternative Channel: Quality Upgrading

An alternative explanation of the evidence in Tables 1.9 and 1.10 is that the export market demands higher quality goods, which in turn commands higher prices (Harrigan, Ma, and Shlychkov, 2011; Manova and Zhang, 2012). If higher quality products require higher energy intensities, then the quality channel would push up emission intensity at the firm-product level, despite any technological upgrading. This could explain why the negative impact on E_i/Q_i does not survive to the firm-product level: quality upgrading works against technological upgrading, so even if firms invest in cleaner technology, the increase in energy intensity due to increases in quality makes the impacts difficult to see.

We investigate this possibility by re-estimating equations (1.27)-(1.28) for industries with low scope for quality differentiation, thus ruling out the quality channel. Khandelwal (2010) computes the degree of quality differentiation within an industry by estimating variety-specific quality scores for different products within an industry and taking the difference between the minimum and the maximum of these scores within the industry. Khandelwal (2010) refers to the resulting measures as quality "ladders." Industries with long quality ladders are industries for which quality differentiation could play an important role, while industries with short quality ladders leave less room for this channel to matter. To distinguish long-ladder industries from short-ladder industries in our sample, we take the HS6-level quality ladder estimates from Khandelwal (2010) and pass them through our Prowess-HS6 mapping from Appendix A.2. We then take the average quality ladder across Prowess products at the industry-level. Resulting ladder lengths are reported in Table 1.13. Short-ladder industries, i.e. industries with comparatively lower scope for quality differentiation, include

Table 1.13: Quality Ladders by Industry

2-Digit Industry	Ladder Size
Mineral Products	1.19
Base Metals	1.36
Wood, Pulp & Paper Products	1.53
Food products, beverages & tobacco	1.66
Agricultural products	1.80
Non metallic mineral products	1.84
Plastics & Rubbers	2.02
Machinery	2.09
Misc. Manufactured Articles	2.18
Chemicals	2.21
Textiles	2.40
Transport equipment	2.66

Notes: These estimates of sector-specific quality ladders are obtained by taking the HS6-level quality ladder estimates from Khandelwal (2010) and passing them through our Prowess-HS6 mapping, and then by computing the sector average.

"Mineral Products," "Base Metals," and "Wood, Pulp, & Paper Products." Industries with greater scope for quality differentiation include "Textiles" and "Transportation Equipment."

Table 1.14 reports the results of re-estimating equations (1.27)-(1.28) for industries with below-median ladder estimates. The firm count drops to 649 for the pooled sample of multi-product and single-product firms (columns 1-6) and 307 for just single product-firms (columns 7-9). We find that restricting to short-ladder industries does not change the qualitative results. E_i/V_i still falls dramatically with an exogenous increase in export value (column 1), while we cannot reject the null hypothesis of no impact on E_{ij}/Q_{ij} (column 7). If quality explained the null result in Table 1.10 column 4 then we should not be able to reproduce the findings in these short-ladder industries. Since the result extends to the restricted sample, we conclude that quality likely does not explain the null result in Table 1.10.

1.6 Conclusion

Many worry that globalization exacerbates environment-related market failures, yet new evidence suggests that exporting encourages firms to increase productivity, which may lower emission intensity of production. This latter possibility is difficult to test empirically because emission intensity is usually measured in value, not quantity, and usually aggregated across products within the firm. Previous work indicates that prices and product-mix could also explain firm-level trade impacts. We model how these alternative channels impact emission intensity theoretically. Then, using a highly detailed dataset on Indian manufacturing firms and an instrumental variable strategy to address endogeneity, we separately estimate the different ways that exporting impacts emission intensity of the firm.

First, we find that prices systematically bias estimates when emission intensity is mea-

Table 1.14: IV Estimates of Exporting Impacts on Emission Intensity in Homogeneous Industries

Dependent Variable:	Log (E/V)			Log (E/Q)			Log (E/Q)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Second Stage									
Log Export Value	-0.55***	-0.37***	-0.08***	-0.24	-0.30***	-0.05**	0.05	0.11	-0.06
	(0.18)	(0.10)	(0.02)	(0.15)	(0.11)	(0.02)	(0.18)	(0.19)	(0.04)
Panel B: First Stage									
Log $\widetilde{D}_{it}$	0.14***			0.14***			0.26		
	(0.05)			(0.05)			(0.17)		
Log $\widetilde{D}_{it}^{HI}$		0.18***			0.18***			0.19*	
		(0.04)			(0.04)			(0.11)	
Log $\widetilde{D}_{it}^{LI}$		-0.02			- 0.02			-0.03	
		(0.04)			(0.04)			(0.12)	
R^2	0.090	0.095		0.090	0.095		0.050	0.050	
F-stat	9.85	11.33		9.85	11.33		4.43	3.95	
Panel C: Reduced Form									
Log $\widetilde{D}_{it}$	-0.08***			-0.03			0.01		
	(0.02)			(0.03)			(0.06)		
Log $\widetilde{D}_{it}^{HI}$		-0.06**			-0.06**			0.02	
		(0.02)			(0.02)			(0.04)	
Log $\widetilde{D}_{it}^{LI}$		-0.01			0.02			0.02	
		(0.02)			(0.02)			(0.03)	
R^2	0.057	0.057		0.035	0.035		0.04	0.04	
Selection Correction	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Energy Prices	Y	Y	Y	Y	Y	Y	Y	Y	Y
Includes MP Firms	Y	Y	Y	Y	Y	Y	N	N	N
Num of Obs	3355	3355	3355	3355	3355	3355	1495	1495	1495
Num of Firms	649	649	649	649	649	649	307	307	307

Notes: Panel A reports the second stage impact of exporting on emission intensity, Panel B reports the corresponding first-stage impacts of demand shocks on log export value, and Panel C reports the reduced form. Columns 1-2, 4-5, 7-8 report instrumental variable estimates, while columns 3, 6 and 9 report the OLS. All data are aggregated to the firm-year level. Columns 1-6 Include multi-product and single-product firms, while columns 7-9 include just single-product firms. Emissions values computed from firm-level energy reports. Sample includes only firms that export some positive value over the period and firms that report outputs in consistent units across products and over time. Columns 1-3 report estimates for emission intensity in value (E_i/V_i), while columns 4-6 report estimates for emission intensity in quantity (E_i/Q_i). All regressions control for the inverse mills ratio of exporting in a given year. Demand shocks are averaged over products produced within the firm using base year product sales shares (first year the firm appears in the dataset). All regressions include years 1995-2011. The same restriction on outliers as throughout applies. Standard errors are clustered on the HS6 category of the core product in the base year (first year the firm appears in the dataset). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

sured in value. This is because firms charge higher prices in the export market, especially when exporting to high-income countries (which proxy for high-regulation countries). Emission intensity in value at the firm level falls 0.57% with an (instrumented) 1% increase in export value, but approximately 1/3rd of the effect is purely from higher prices, confirming the OLS results in section 3. The results caution against interpreting productivity impacts when outputs are only denominated in value, like De Loecker (2011) and De Loecker et al. (2012).

Second, we find that firm-level emission intensity in quantity falls 0.38% with an (instrumented) 1% increase in export value. The sign is the same as in the OLS, but the effect is much stronger. However, disaggregating to the firm-product level, we can reject the technological channel at the 5% level. Thus, all of the firm-level impact can be attributed to product-mix. The result implies that firms *do* respond to changes in foreign market conditions, but not through the technological upgrading mechanism advanced by the literature. Researches and policy-makers should take note of the product-mix effect when measuring the impact of policies on clean technology investments.

A natural question to ask with respect to the results of this paper would be: why did foreign demand shocks fail to induce cleaner technology adoption in this case? One explanation is that technological change is biased against environmental inputs, so firms *do* adopt more-efficient technologies, but these technologies are no cleaner than older processes. Addressing this possibility with the Prowess dataset is difficult because only energy inputs are broken down by output product. Thus, estimating the correlation between total factor productivity and emission intensity would require additional structure. We leave it for future research to investigate biases either towards or against environmental inputs in technology adoption in the developing world.

Another possible explanation for the absence of technological upgrading is pre-existing market failures. Greenstone and Jack (2013) describe how weak environmental regulations and imperfect capital markets inhibit investments in developing countries. With two pre-existing market distortions (regulations and capital markets), perhaps the signal from foreign demand shocks is insufficient on its own to induce clean technology investment. In this case, the results call for a mixture of trade policy and environmental regulation/credit market interventions in order to bring firms in the developing world to the technological frontier. We leave investigation of these market frictions for future research.

Chapter 2

Clean Clothes: Exporting and the Environmental Impact of Textile Production under the MFA

With Hélène Ollivier

2.1 Introduction

A growing body of research finds that trade liberalization leads to better environmental performance within firms over time.[1] While the results support the environmental benefits of globalization, estimates neglect "third-party" effects on countries that *lose* market share due to liberalization. In most standard trade models, when two countries bilaterally reduce tariffs with each other, "third-party" countries – peripheral to the liberalization – lose market share as competition increases. If export market access encourages firms to reduce emission intensity, then firms whose sales are crowded-out of the newly-liberalized markets may *increase* their emission intensity as an indirect result of the liberalization. In this paper, we provide what we believe are the first estimates of this relationship, and thus offer a new channel through which trade influences the environmental consequences of production.

The liberalization episode we study comes from a well-known quasi-natural experiment in global trade in textile and apparel in the late 1990's and early 2000's. Prior to 2005, exports of textile and apparel products from the developing world to the developed world were restricted under a system of bilateral product-by-country-specific quotas known as the Multifiber Arrangement (MFA). Under the MFA, different exporting country-product pairs were subject to different regulations across destination markets, but not all country-product categories were subject to quota, and not all quotas were binding. Thus, with the complete dismantling of the MFA between 1995 and 2005, firms operating in different product codes in different countries experienced differential changes in market access both from reductions of their own quotas (direct effect) and from reductions of quotas on competing firms in other countries (third-party effects). This episode has been used to study the impacts of trade

[1] See Gutiérrez and Teshima (2011), Martin (2012), and Cherniwchan (2013) for evidence that firm-level emission intensity falls with increased export market access.

policy in many contexts (Harrigan and Barrows, 2009; De Loecker, 2011; Bloom, Draca, and Van Reenen, 2011; Khandelwal, Schott, and Wei, 2013; Rotunno, Vézina, and Wang, 2013), though never the environmental effects before.

To test for environmental impacts of MFA liberalization, we map detailed production and emission data for Indian firms to product-by-year export quotas to the US for these firms, along with the weighted-average quota rates for India's competitors (i.e., quota rates for all countries other than India to the US). The production dataset – Prowess, a long panel of large Indian manufacturers – is uniquely well-suited for our purposes because inputs and outputs are reported with enough granularity that we can circumvent many well-known problems that arise when estimating productivity measures in firm-level datasets (De Loecker, 2011; De Loecker et al., 2012). First, Prowess reports outputs in both quantity and value at the product level, so we can distinguish between physical productivity and revenue productivity. Second, Prowess contains energy input information at the *product level*, which enables us to compute CO_2 intensity of production directly from the data. With emission intensity computed in physical quantities of CO_2 per physical quantities of output at the firm-product level, we can isolate technological effects from across-product substitution effects, which previous work has shown to be important (Bernard, Redding, and Schott, 2011; De Loecker et al., 2012; Mayer, Melitz, and Ottaviano, 2014).

The identification strategy is to compare changes in physical emission intensity in quantity within firm-products over time for products that were previously constrained under the MFA (via binding US-India quota) vs products that were not, and products whose competitors were previously constrained (via binding US-competitor-country quotas) vs products whose competitors were not. To illustrate the strategy, consider the case of two goods: "Men's & Boy's shirts," whose exports from India to the US were constrained by binding quota throughout the late 1990's/early 2000's, vs "gloves," whose exports were not. By 2007, when all MFA quotas had expired, Indian firms producing Men's & Boy's shirts saw market access to the US increase compared to producers of gloves, since producers of Men's & Boy's shirts in India were no longer constrained by binding quotas. However, at the same time, throughout the late 1990's/early 2000's, 68% of glove exports (on average) to the US from countries *other* than India were subject to binding quotas, while the corresponding figure for Men's & Boy's shirts was only 25%. By 2007, both shares dropped to 0, but producers of gloves saw greater reductions in competitor export constraints relative to producers of Men's & Boy's shirts, since India's competitors in gloves were initially more constrained. If export market access induces lower emission intensity, then direct trade effects (reductions of US-India quotas) should lead to lower emission intensity in Men's & Boy's shirt producers relative to glove producers, whereas third-party effects (reductions in competitor quota constraints) should lead to higher emission intensity in glove producers relative to Men's & Boy's shirt producers. This is the basic logic of the estimation strategy.

We first demonstrate that MFA liberalization affected the exports of firms in Prowess in the predicted way. Using a fixed-effect estimator that controls for unobserved time-invariant firm-product factors and year-specific macro shocks, we find that firm-level exports correlate positively with our measures of competitor constraints, statistically significant at the 1% level. This result implies that when the US eliminated quotas for India's competitors, firms in Prowess lost market share. We estimate that the average exporter in Prowess lost 14% export sales as a result of MFA quota liberalization. By contrast, we find that binding India-

US quotas did *not* reduce exports to the US, but given that India was relatively unconstrained under the MFA compared to competitor countries, this is perhaps not too surprising. The results are consistent with previous research that documents large increases in exports from India's competitors after the MFA, especially from China (Harrigan and Barrows, 2009). Additionally, the results are consistent with firm responses in European and other East Asia countries documented in previous studies (De Loecker, 2011; Amiti and Khandelwal, 2013; Bloom, Draca, and Van Reenen, 2011), but where previous work infers competition effects from productivity responses, we can see the crowding-out effects of third-party quota constraints on exports directly.

Next, we relate the evolution of emission intensity within firm-product over time to the elimination of quota constraints in the difference-in-difference manner illustrated above. Consistent with the export effects, we find that emission intensity falls with competitor constraints, but not India's own quota constraints. This result implies that third-party effects of liberalization lead to higher emission intensity for Indian firms, on the order of about 9%. We also present results from a placebo test of quota impacts on non-exporters, and find no statistically significant impact on the emission intensity of firms that never export over the period. This result is what one would expect, if there are no spillover effects from exporters to non-exporters. In a final section, we consider several possible channels for the results, including fuel switching, factor-biased productivity gains, and quality adjustments.

This paper relates to a large literature on the impacts of trade on the environment. Early work assumed homogeneous firms (Copeland and Taylor, 1994; Antweiler, Copeland, and Taylor, 2001), finding that trade tends to reduce not only emission intensity, but overall emission levels, through income-induced endogenous environmental regulation. But this literature did not consider third-party effects in general equilibrium. More recently, another line of papers relaxes the homogeneity assumption and highlights selection effects and productivity growth holding regulation fixed (Kreickemeier and Richter, 2014; Holladay, 2010; Forslid, Okubo, and Ulltveit-Moe, 2011; Cui, Lapan, and Moschini, 2012; Jing Cao and Zhou, 2013; Galdeano-Gómez, 2010; Gutiérrez and Teshima, 2011; Martin, 2012; Cherniwchan, 2013). This literature also argues that trade reduces emission intensity in countries that participate to the liberalization, but again there is no discussion of how bilateral liberalization between two countries impacts the production decisions of firms elsewhere in the world.

Beyond the trade and environment literature, this paper also relates to a large literature that investigates the impact of competition on productivity. The emphasis in this literature tends to be on effects for domestic firms that see increased foreign import competition from developing countries (see for recent examples De Loecker (2011) Bloom, Draca, and Van Reenen (2011)). By focusing on third-party effects, we extend the competition results to markets beyond the developed world, where most of the import penetration occurs. Additionally, the previous literature usually requires strong functional form assumptions to compute the outcome variable – total factor productivity. Instead, we compute productivity of a single factor directly from the data without imposing any structure at all.

Finally, the paper also relates to a literature on the causes of high emission intensity of firms in developing countries (Duflo et al., 2013; Greenstone and Jack, 2013; Greenstone and Hanna, 2014). The literature here focuses mainly on market imperfections in the exporting country (e.g., weak environmental regulation, corruption). Our results show that firms' environmental performance also responds to international competition.

2.2 Background and Empirical Strategy

In this section, we present the quasi-experimental setting and the measures of competitor constraints. In the following section, we present the production data and test the parallel trends assumption.

The Multifiber Arrangement

In order to estimate the impact of international competition on the emission intensity of individual firms, we must observe a large change in competition whose cause is unrelated to unobservable determinants of the production process. A common strategy is to relate changes in production to changes in tariff rates across different goods; however, tariffs often change gradually over time, and potentially endogenously to strategic trade interests (Trefler, 1993). In this paper, we exploit the rapid and complete dismantling of a large system of import quotas in textile and apparel between 1995 and 2005 – the end of the MFA – which generates large and abrupt changes in competition in the destination markets.

Prior to 2005, export flows in different textile and apparel products from developing countries to the developed world were restricted through a system of bilateral import quotas known as the Multifiber Arrangement (MFA). The MFA began with US-imposed import quotas on a few Japanese textile products following World War II (under the name "Agreement on Textile and Apparel"), and then expanded to encompass the entire sector and many developing countries. Quotas levels were set in physical quantities and applied to fairly aggregate product groups, e.g., "Dozens of Men's & Boy's shirts." Temporal variation in protection rates is due to the fact that not all country-by-product categories were subject to quotas and not all quotas were binding. For example, in Figure 2.1, we find that China had the most binding quotas by far throughout most of the 1990's and 2000's, with on average 60 out of 167 categories subject to binding quota, while other East Asian countries like Indonesia, Hong Kong, and Bangladesh were also heavily constrained, though at lower rates.[2]

As part of the Uruguay Round of the General Agreement on Tariffs and Trade, the protected countries agreed to eliminate MFA quotas in each of four predetermined phase-out years – 1995, 1998, 2002, and 2005 – lifting all quota constraints by 2005 (see Figure 2.1).[3] Since the agreement mandated the complete removal of all quotas, there was little scope for differential enforcement across product categories, so endogenous lobbying efforts from

[2] Quotas are defined as "binding" if the fill rate for the quota in a given year exceeds 90%, following the majority of the MFA literature. Quota limits and fill rates were tracked by the US Commerce Department's Office of Textile and Apparel (OTEXA), and published online by Brambilla, Khandelwal, and Schott (2010). See Appendix B.1 for details.

[3] Protected countries were required under the agreement to retire at least 16.7% of all quotas in each of the years 1995, 1998, and 2002 and the remainder in 2005. Brambilla, Khandelwal, and Schott (2010) document that the US complied with the agreement, phasing out the required levels in each of the four years. However, note that the onslaught of cheap textile goods into the US triggered a safeguard mechanism whereby the US was permitted under the agreement to re-impose temporary quotas on a few categories that saw particularly high spikes in foreign import activity. The safeguards did not last long, as all quotas were lifted by 2006. In the results we present, we use the actual binding status of quotas to construct the instrument, but for robustness, we also consider setting 2005 quota levels to 0, as was intended under the agreement, and the results are unchanged.

Figure 2.1: Top 10 Constrained Countries under the MFA

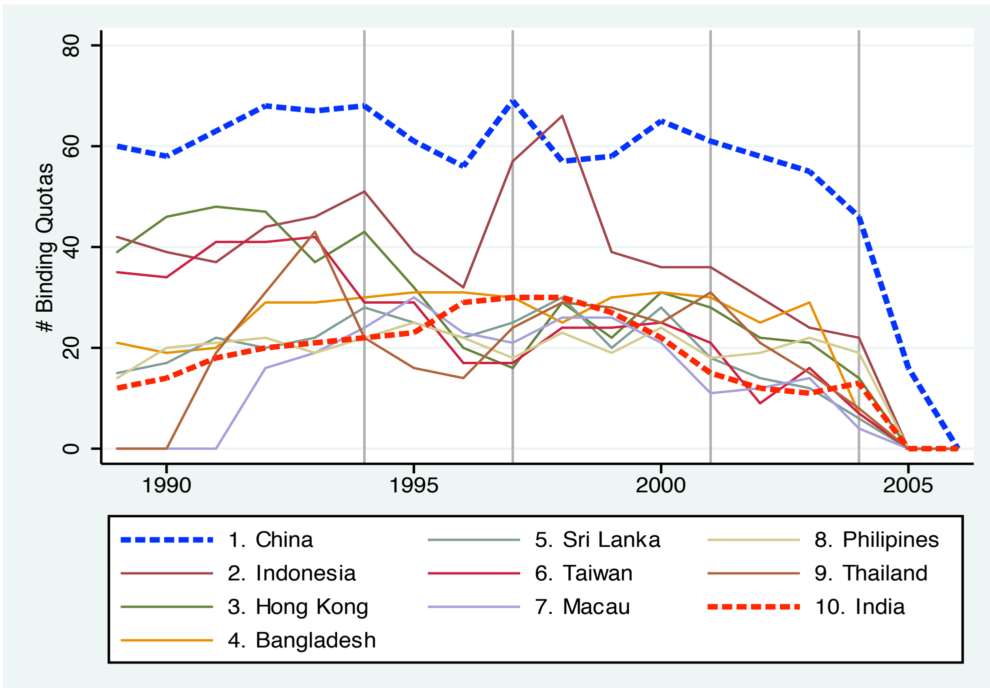

Notes: Figure plots number of binding quotas in place by year for the ten most constrained countries under the MFA based on 1994 ranking. A quota is considered binding if the fill-rate is greater than 90%

specific industries likely did not play a role in the timing or enforcement of the liberalization across different product groups. Furthermore, initial quota levels set in the 1950's and 1960's depended on bargaining power of different sub-industries at that time, and tended to persist strongly through the century (Bloom, Draca, and Van Reenen, 2011). Thus, the level of, and hence the *change* in, quota protection across different goods in the late 1990's/early 2000's can be taken as exogenous to trends in unobservable determinants of firm-level productivity at that time.[4]

If we focus in on two key countries for the analysis, China and India, we can see the liberalization in action more clearly. Figure 2.2 plots the number of categories subject to active quota (whether binding or not) and binding quota for these two countries over the period. The solid lines count the number of binding quotas while the dashed lines count the number of total active quotas. The four phase-out years are again denoted with vertical lines. We can see that between 1997 and 1998, i.e., in the second phase of liberalization, the number of total quotas for China dropped slightly, while staying fairly constant in the 10 years prior. Next, between 2001 and 2002, i.e., the third phase of liberalization, Chinese

[4]De Loecker (2011), Khandelwal, Schott, and Wei (2013), and Rotunno, Vézina, and Wang (2013) similarly assume exogeneity of regulation-era quota levels, and hence interpret firm responses to quota removal as causal.

Figure 2.2: China and India Active and Binding Quotas

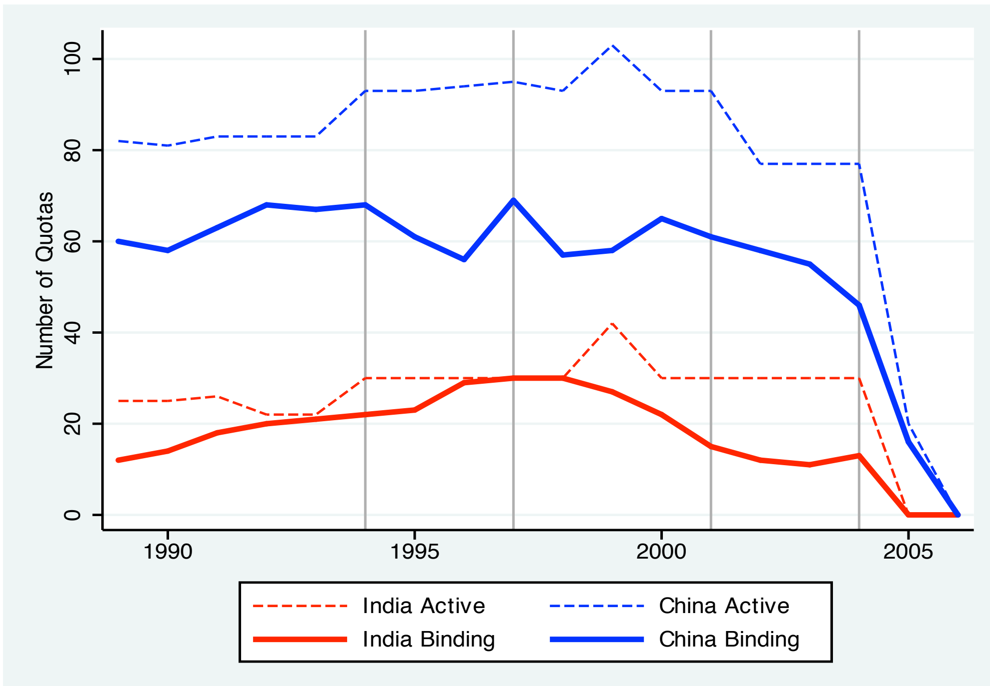

Notes: Figure plots number of active and binding quotas in place by year for India and China. A quota is considered binding if the fill-rate is greater than 90%

quotas drop suddenly, from 90 to 80. India, by contrast, saw no such discrete drops in either of these phases. Finally, both countries saw huge drops after 2004, as all quotas were lifted (subject to the safeguard caveat for China). We exploit the timing of these quota liberalizations to test for production responses to changes in the competitive environment.

Competitor constraint indices

The expiration of MFA quotas should have both a direct effect on India firm exports from the removal of binding US-India quotas, and an indirect effect from the removal of binding US-other country quotas.[5] The latter effect is due to the increase in competition in the destination market: the elimination of India's competitors' quota constraints lowers the cost

[5]Since MFA quotas were specified in quantity, instead of value, the precise impact on sales depends on the assumption of market structure (see Krishna (1989) for discussion of quantity-based quota regulations and exporter response). However, virtually any model predicts that the cost of shipping to the regulated market increases: either a regulatory body imposes a formal apparatus for quota licensing, complete with permit prices, or an informal market arises where firms negotiate for the rights to use a quota permit. Thus, the elimination of binding US-India quotas should comparatively reduce the cost of shipping to the US for Indian firms operating in those product categories that see their quotas reduced, and increase their export sales. Harrigan and Barrows (2009) find that the expiration of the MFA caused export prices for quota-constrained products to drop 30% between 2004 and 2005, suggesting that the regulatory cost associated with quotas was quite large under the MFA.

of shipping for those firms, and hence lowers the price index in the destination market. In the language of Anderson and Wincoop (2003), the destination market sees lower "multilateral resistance" as a result of other-country MFA quota removals, and hence becomes harder to reach for Indian firms.

To compute an aggregate measure of Indian competitor constraints by product category under the MFA, we take a weighted average of imports into the US that arrive under binding quota from countries other than India each year. Let g denote a quota group product category, t denote year, and j the exporting country to the US. We define an index of competitor constraints from the rest of the world (ROW) i.e., besides India

$$ROW_{gt} \equiv \frac{\sum\limits_{i \in \Lambda} V_{gj0} * I_{gjt}}{\sum\limits_{i \in \Lambda} V_{gj0}}, \tag{2.1}$$

where V_{gj0} is the value of exports in good g from j to the US in some pre-liberalization base year 0, I_{gjt} is an indicator for whether good g from country j was subject to binding quota in year t, and Λ is the set of all countries other than India that export to the US. It is important that import shares V_{gj0} are evaluated prior to liberalization because the year-to-year shares may respond endogenously to the liberalization itself.[6] The index ROW_{gt} ranges from 0 to 1, where $ROW_{gt} = 0$ indicates that none of the exports in g from India's competitors to the US were subject to a binding quota in year t (conditional on pre-period weights), while $ROW_{gt} = 1$ indicates that all exports to US (other than those from India) in g were subject to a binding quota in t. As ROW_{gt} increases, Indian firms operating in product group g enjoy greater comparative access to the US, regardless of the constraint status of the US-India quota rates, and thus should export more.

The construction of ROW_{gt} in equation (2.1) follows the same logic as the computation of product-specific instruments in many recent papers. For example, Park et al. (2010) weights exchange rate movements in currencies of importers of Chinese prodcuts during the Asian financial crises by base-year export shares, Hummels et al. (2014) weights foreign demand shocks in importers of Danish products by base-year export shares, and Bloom, Draca, and Van Reenen (2011) weights Chinese export growth in products to the EU by base-year exposure to Chinese exports. The common idea is to interact some macro-economic (or product-specific) shock with pre-period exposure rates to compute idiosyncratic changes in market access across similar firms. An advantage of our approach is that we utilize discrete jumps in policy, so pre-existing trends in macro variables are unlikely to contaminate the results. A disadvantage of our specification relative to say, Hummels et al. (2014), is that we don't know which firms sell which products in which markets. With base-year firm-product-destination trade flows, we could exploit variation in firm-specific exposure to various markets and compute ROW_{gt} at the firm-product level.

Figure 2.3 plots ROW_{gt} against the India fill rate for each year between 2000-2005. To compute the ROW index, we aggregate ten-digit harmonized trading system (HS10)

[6] In the estimation, we take 1994 – the final year before the first phase-out round – as the initial period, though results are not at all sensitive to the choice of base year. In fact, for robustness, we compute the ROW_{gt} using both base-year weights and current-year weights and find that in practice, it does not make much difference which one we use.

trade flows from Feenstra, Romalis, and Schott (2002) to the g-exporter level, merge quota indicators indicators I_{gjt} from the Office of Textile and Apparel (OTEXA), and take the weighted average constraint across OTEXA category (g) as in (2.1). Vertical red lines denote the critical 90% fill rate, hence any mark to the right of this line signifies a binding quota for India. Each blue "X" or red "O" corresponds to a product category g. Red "O's" indicate that quota coverage for the product will fall below 1% in the following year: red "O's" are about to see ROW liberalization. The X's denote otherwise. For example, in 2001, we see several categories for which the ROW index will fall to near 0 in 2002, which corresponds to the third phase of liberalization. Similarly, in 2004, all categories other than those for which safeguards will be re-imposed on China are marked with O's.

Figure 2.3: Competition Index and India Fill Rates by Year

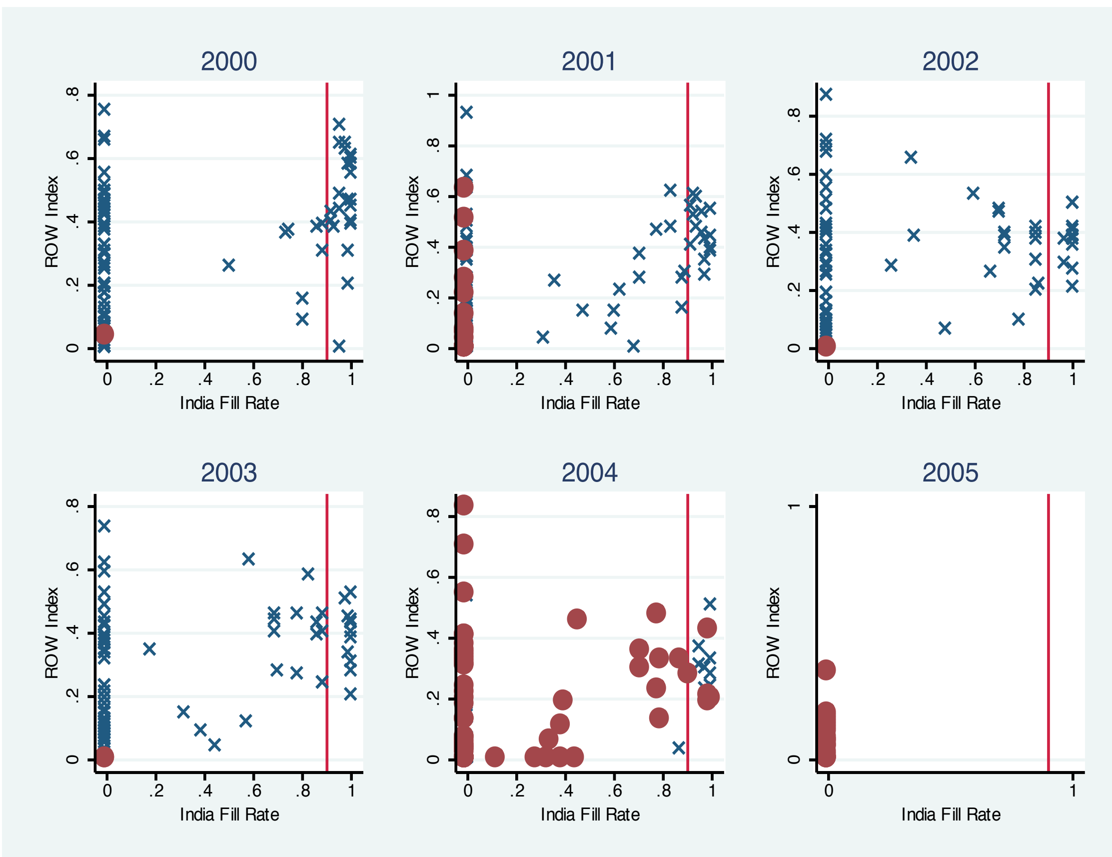

Notes: Each year plots the ROW index against the India fill rates for each 3-digit MFA category. The red line indicates the 90% threshold for binding status. Red dots indicate a category whose ROW quota coverage will fall to below 1% in the next year. Blue X's indicate otherwise.

Two important facts can be gleaned from Figure 2.3. First, in each year prior to 2005, there were several categories in which India's competitors were constrained, but India was not. This can be seen by the mass of X's and O's along the 0 India fill rate mark. This implies that India's competitors will see increased market access in these categories, while India itself will not. In these categories, the prediction for changes in India exports is unambiguous – Indian exports should fall. Second, in each year prior to 2005, we see that categories for which Indian fill rates were high tended also to have high values in the competitor index. Regressing India fill rates on the competitor index for the years 1994-2004 yields a coefficient of 0.88 and a t-stat of 23. Most of the categories in which India was constrained in under the MFA, India's competitors were constrained as well. When these quotas are liberalized, the direct impact fights the crowding-out effect, and the overall effect on Indian exports is ambiguous.

2.3 Prowess Data and Merging to Quota Data

This section presents the output and emissions data contained in the Prowess database. The dataset is described in detail in Chapter 1, so we keep the presentation brief here and refer to Chapter 1 for further discussion. Next, we discuss how to merge the Prowess data to the quota constraint indices.

Production Data

As part of the Indian Companies Act of 1956, Indian firms are required to issue annual reports detailing a wide array of economic activity, including quantity and value of output by product and energy inputs by product. The Center for Monitoring the Indian Economy (CMIE) collected a large subset of these reports, standardized and digitized the information, and published the resulting database, Prowess, for use by investors and researchers. Most inputs such as labor or capital are reported at the firm-level, but concerns about energy security lead the Indian government to mandate greater detail in energy-use reporting. Firms still report total energy consumption (by energy type) for the year, as they would other inputs, but they also report product-specific energy intensity (also by energy type), though not total energy consumption per product. Additionally, firms report the export share in total firm-year revenue, which we use to check that quota constraints impact exporting behavior in the predicted way. Previous researchers have exploited the multi-product dimension of the data (Goldberg et al., 2010; De Loecker et al., 2012), but to our knowledge, no other paper (beyond Chapter 1) have utilized the product-specific energy reports yet.

The product-specific energy reports are quite important for our purposes because they allow us to compute CO_2 intensities at the firm-product level directly from the data.[7] To generate these values, we multiply product-specific energy-type intensities by CO_2 emission intensities and sum over energy types at the product level (see Chapter 1 for details). One caveat to mention with regards to the emission intensity data is that they are reported in a separate module of the dataset, so connecting emission intensity to sales and prices at the firm-product level requires merging between the two modules. Unfortunately, this process is not straight-forward because neither product name nor product classification are entirely

[7] An alternative approach would be to estimate a structural model of production for single-product firms and then impute product-specific energy intensity based on the estimated coefficients (as in De Loecker et al. (2012). Our approach requires no functional form assumptions at all.

consistent between the two modules. In Chapter 1, we present an automated process for merging the two modules, but here, since we are only using the textile and apparel sector, we chose to map by hand to ensure a tighter link between the emissions and the output data. See appendix B.1 for details.

Table 2.1: Summary Statistics

Variable	Exporters	Non-Exporters	Diff
Sales Value (Mill of Rs)	872.8	217.6	***
	(2285.0)	(423.9)	
CO_2 Emissions (MT)	0.022	0.006	***
	(0.043)	(0.028)	
CO_2 Intensity (MT/ Real Units Output)	2.94	2.71	***
	(4.05)	(3.30)	
Export Share (f.o.b.)	20.92	0	-
	(26.87)		
N	5751	1259	
# Firm-Products	814	280	
# Firms	523	226	

Notes: Annual production data from 1994-2007. An observation corresponds to a firm-product-year, except for Export share, which is computed at the firm-year level. "Exporters" earn positive export revenue for at least one year over the period.

Descriptive statistics broken down by exporters and non-exporters are reported for the years 1994-2007 in Table 2.1. Exporters are defined as those firms that earn positive export sales for some year over the period. We have an unbalanced panel of 1,094 firm-products manufactured by 749 distinct firms. The mean sales generated by a firm-product line in a given year is 872.8 million Rs (roughly 19 million USD) for exporters and 217.6 million Rs (roughly 5 million USD) for non-exporters. With the mean number of products sold by a firm in a given year at 1.46, the average firm in the dataset generates 17 million USD in revenue. By comparison, the average textile firm in the comprehensive ASI dataset generates only 4 million USD in revenue, so firms in our sample tend to be larger than the average textile firm in India.[8] However, since exporters tend to be larger in general, the sample is appropriate for studying trade impacts.

Since exporters are bigger than non-exporters, they also generate more pollution. We calculate that the average exporter product-line generates 0.022 MT of CO_2 in a given year, with the corresponding figure for non-exporters being 0.006 MT of CO_2. In terms of emission intensity, the average product-line generates 0.0256 kg CO_2 per Indian rupee. For comparison, the average textile firm in the US generates 0.005 CO_2 per Rs, so the mean firm in the data is about 5 times less efficient than the mean US textile and apparel plant.[9]

[8] In the 2010 ASI, Table 5-A reports output for textile and apparel sector (NIC codes 13-14) as 2720 billion RS and 766 billion RS, while total establishments equals 13,159 and 2,052. Summing output and dividing by number of establishments yields 243 Million RS per factory. Converting the USD with an exchange rate of 62 RS/USD yields 4 million USD of output per plant

[9] From 2006 MECS data, we compute the average US textile plant use 0.427 KhW of electricity per dollar

Finally, the average exporter earns a little over 20% of revenue on the export market in a given year. We observe export share over time within the firm, and thus can check that quota constraints impact exporting behavior in the predicted ways.

Figure 2.4: Changes in emission intensity

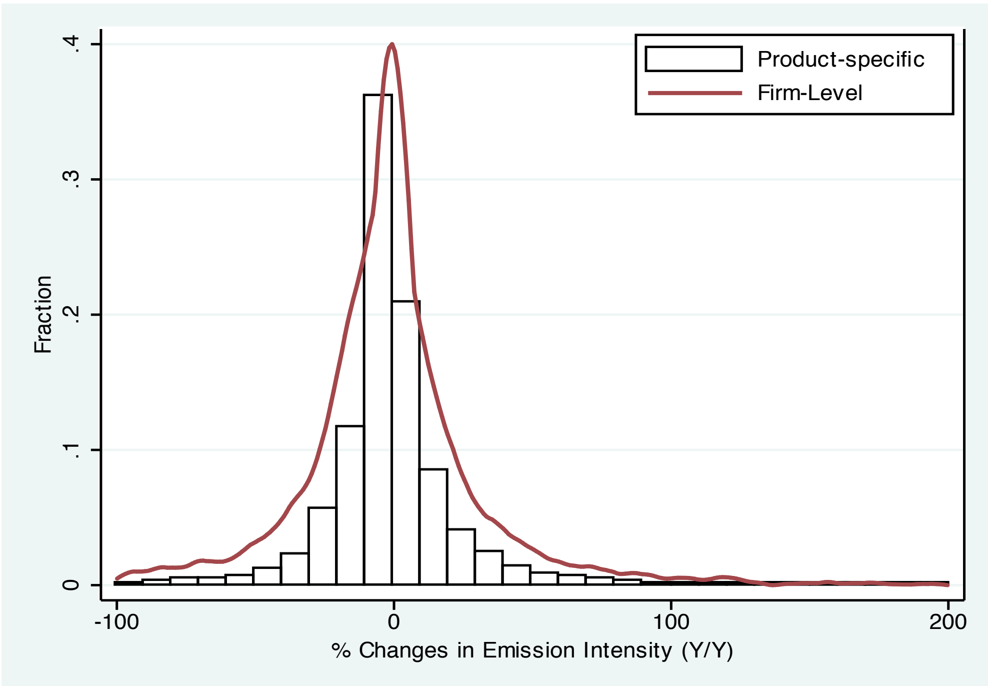

Notes: Figure plots year-on-year percent changes in emission intensity for the years 1994-2007 at different levels of aggregation.

To get a preliminary sense of the temporal variation in emission intensity, we plot the distribution of year-to-year percentage changes within the firm-product in Figure 2.4. Denoted by transparent bars, we find significant year-to-year variation at the firm-product level: 10% of the observations register a year-on-year increase of greater than 26%, while 10% of the observations see a decline of more than 21%. Certainly, these changes reflect some measurement error as well, but the unconditional distributions reflect sufficient variation that we cannot rule out trade impacts *prima facie*. Additionally, we plot the kernel density of changes in emission intensity computed from firm-level energy reports and find a similar degree of within-firm variation.

Merging Quota Data to Prowess

The last task to perform before estimating the impact of the MFA on exports and emission intensities is to merge the quota constraints to the Prowess data. This merge presents a

shipped. Multiplying by a constant emission factor of 1.34 lbs of CO_2 to KwH and converting to kg per RS, we estimate average emission intensity of 0.005 kg CO_2 per Rs.

significant challenge in that Prowess organizes products according to an internal classification system that does not map directly to other known classifications, while OTEXA uses an idiosyncratic classification system as well (the 167 3-digit categories discussed in section 2.2). We proceed by first mapping Prowess production codes to HS trade data classifications, and then mapping to the OTEXA data. Note that competitor constraints ROW_{gt} are defined as in section 2.2. US-India constraints are denoted $India_{gt}$, and take the value 1 when product g is subject to a binding US-India quota in year t and 0 otherwise.

The Prowess-HS map was discussed in Chapter 1, so we refer the reader there for details. But to re-iterate, both Prowess and the HS classifications hew closely to the ISIC system, so while names and numbering are different between the two, it is fairly straightforward to map between them by hand. Alternative methods of mapping between Prowess and HS data exist (De Loecker et al., 2012), but they involve mapping first through the India National Industry Classification system (NIC), which is fairly aggregate compared to the 6-digit HS trade classifications.

Next, with our Prowess-HS correspondence, we merge to OTEXA codes via an HS10 mapping published on OTEXA's website.[10] The correspondence assigns each HS10 product to one of the 167 OTEXA categories (e.g., "Men's and Boy's shirts"); however, mapping from HS4 or HS6 through the OTEXA-HS10 mapping requires a few steps. First, in the Prowess-HS map, Prowess codes do not always map to unique HS4 or HS6 categories. In the case that these HS4 or HS6 represent multiple OTEXA categories, then there will not be a unique OTEXA category for each Prowess code. Additionally, even if the Prowess code matches to a single HS4 or HS6, these four or six digit classifications may span multiple OTEXA categories. For example, we map Prowess product "woven fabric of carded wool" to HS4 code 5111, but this 4-digit HS code contains HS10 codes that map to both OTEXA 410 and 414 ("wool yarn" and "woven fabric of wool"). To compute a single index value for each Prowess product code, we take a simple average of $India_{gt}$ and ROW_{gt} over possible OTEXA codes (in example above, we would average $India_{gt}$ and ROW_{gt} for codes 410 and 414). In a slight abuse of notation, we continue to refer to product regulation categories as g, though when we merge to Prowess, they are really averages over multiple OTEXA categories. See appendix B.1 for details.

Tables 2.2 and 2.3 report the top ten most-constrained and least-constrained Prowess products under the MFA in terms of $India_{gt}$ and ROW_{gt}, respectively. In each table, we list the products with the ten highest average constraints during the regulation period (1994-2004) on the left, and the ten products with the lowest constraints on the right. Looking first at Table 2.2, we find that "Sacks/Bags (Cotton)" had the highest average $India_{gt}$ over the regulation period ($India_{gt}$=1.00), followed by "Textile labels, badges, etc." ($India_{gt}$=0.50), and then "Carpets, etc." ($India_{gt}$=0.33), "Curtains, blinds, etc." ($India_{gt}$=0.33) and "Cotton fabrics grey (Mill sector)" ($India_{gt}$=0.33). Products initially unconstrained in exports to the US include "Bedspreads," "Cotton & blended yarn, texturised," and "Fishing nets."[11] No clear pattern emerges in terms of types of products that are likely to have have high vs low

[10]http://otexa.trade.gov/corr.htm

[11]Note that in fact, 69 of the 95 Prowess product codes take the value $India_{gt} = 0$ throughout the pre-period, so these products all tie for the lowest $India_{gt}$ (0). We just choose ten of these products for illustration.

Table 2.2: High vs Low constrained products under MFA from India

High-constraint products		Low-constraint products	
Product	India index	Product	India index
Sacks/Bags (Cotton)	1.00	Cotton & blended yarn, texturised	0
Textile labels, badges, etc.	.50	Textile products for technical uses	0
Carpets, etc.	.33	Wool yarn	0
Curtains, blinds, etc.	.33	Yarn of artificial staple fibres	0
Cotton fabrics grey (Mill sector)	.33	Woven fabrics of made filaments	0
Towels including turkish towels	.31	Other synthetic filament yarns	0
Cotton fabrics (Handloom sector)	.27	Polypropylene filament yarn	0
Other clothing accessories, knitted or crocheted	.23	High tenacity yarn of viscose rayon	0
Millmade fabric	.22	Bedspreads	0
Blankets & travelling rugs	.20	Fishing nets	0

Notes: This table reports average values for the India constraint Index for the years 1994-2004 by prowess product classification. The left panel reports the top-10 highest constraint values, while the right panel reports a selection of the lowest constraint values.

Table 2.3: High vs Low constrained products under MFA from ROW

High-constraint products		Low-constraint products	
Product	ROW index	Product	ROW index
Sacks/Bags (Cotton)	.63	Cotton & Polypropylene fibre	0
Textile labels, badges, etc.	.32	Hessian	0
Curtains, blinds, etc.	.28	Wool yarn	0
Other clothing accessories, knitted or crocheted	.25	Other jute products	0
Cotton fabrics (Handloom sector)	.23	Rubberised textile fabrics	0
Cotton fabrics (Powerloom sector)	.22	Textured yarn of syn filament	0
Cotton fabrics grey (Mill sector)	.22	Fishing nets	0
Apparels - knitted / crocheted	.21	Polypropylene filament yarn	0
Carpets, etc.	.21	Knitted fabrics	0
Woven fabrics, of man-made filaments	.18	Felt	0

Notes: This table reports average values for the ROW constraint Index for the years 1994-2004 by prowess product classification. The left panel reports the top-10 highest constraint values, while the right panel reports a selection of the lowest constraint values.

constraints, which supports the orthogonality of the regulation to unobservable determinants of emission intensity.

In Table 2.3, we find that Indian product initially heavily shielded from international competition in the US market (high ROW_{gt}) include the same three heavily-regulated products as in Table 2.2, in addition to "Other clothing accessories," and "Cotton fabrics (Handloom sector)." Products categories in which India's competitors were already unconstrained in their exports to the US under the MFA include "Cotton & Polypropylene fibre," "Hessian," and "Wool yarn."[12] Again, there appears to be no clear pattern of difference between constrained or unconstrained products in Table 2.3.

In Table 2.4, we test formally for endogenous regulation following a strategy from Topalova and Khandelwal (2011) that relates industry characteristics to the strength of pre-shock regulation. We regress product-category characteristics averaged over the pre-liberalization period (1988-1994) on the category's constraint level prior to the start of the sample (1994) to see if regulation systematically targets certain kinds of industries. Characteristics we consider are the average and standard deviation of emission intensity in terms of value and energy within product codes. To compute these summary statistics, we first regress log emission intensity (either in value or energy) on a full set of year dummies, compute the residual, exponentiation, and then compute the mean and standard deviations. We restrict the sample to product categories with substantial coverage (hence the number of product categories drops to 53). We find in Table 2.4 that we cannot reject the null hypothesis of no correlation for any of the product-specific characteristics with either of the constraint indices, which further supports the exogeneity of the trade shock to firm-level trends.

2.4 Results

In this section, we relate the time-varying MFA regulation variables $India_{gt}$ and ROW_{gt} to firm-level exports and firm-product-level emission intensity in the Prowess dataset. We then investigate possible channels to explain the results.

Trade Liberalization Impacts on Exports

We begin by estimating the impact of the quota constraint indices on exports of firms in Prowess. Exports are reported at the firm-level (in contrast to the production data, which we have at the firm-product level), so we first aggregate $India_{gt}$ and ROW_{gt} over all products produced by the firm. For each firm-year, we weight $India_{gt}$ and ROW_{gt} for each product by base-year output shares within the firm, where the base year corresponds to the first year the firm appears in the dataset. Denoting base-year-weighted average constraints by $India_{ft}$ and ROW_{ft}, we estimate

$$Export\ share_{ft} = \alpha_f + \alpha_t + \beta^X * India_{ft} + \delta^X * ROW_{ft} + W'_{ft}\Gamma + \epsilon_{ft} \tag{2.2}$$

where $Export\ share_{ft}$ represents the share of revenue earned from exporting by firm f in year t (across all products and all destinations), W_{ft} represents firm-year controls such as

[12]For ROW_{gt}, 47 of the 95 products are unconstrained throughout the period 1994-2004. Again, we just take a sample for Table 2.3

Table 2.4: Correlation of constraint indices with product-code characteristics

	Kg CO_2 per Rs		Kg CO_2 per mmBTU	
	(1) Mean	(2) Sd Dev	(3) Mean	(4) St Dev
ROW	-4.2	7.7	0.5	-0.3
	(8.5)	(46.7)	(0.4)	(0.3)
India	-3.2	-26.3	0.2	-0.1
	(5.7)	(31.3)	(0.3)	(0.2)
Obs	53	53	53	53
R2	0.03	0.02	0.11	0.07

Notes: Dependent variable is the average and standard deviation of emission intensity in terms of value (kg CO_2 per Rs) by prowess product code in columns 1 and 2, and average and standard deviation of emission intensity in terms of energy ((kg CO_2 per mmBTU) in columns 3 and 4. Firm-year-product values are the residuals from regressions with year fixed effects. Constraint indices are 1994 values. Product-code statistics are averaged over the period 1988-1994. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

energy prices and total sales, ϵ_{ft} represents a time-varying unobserved idiosyncratic shock to exports, and α_f and α_t denotes firm and year fixed effects. The direct trade impact (β^X) and the indirect competition effect (δ^X) are identified from year-to-year variation in $India_{ft}$ and ROW_{ft} induced by quota expiration under the MFA. This specification is similar to the one used by De Loecker (2011) to estimate the MFA's impact on productivity of Belgian producers, except that we separately estimate β^X from δ^X. The specification is also reminiscent of Amiti and Khandelwal (2013) and Bloom, Draca, and Van Reenen (2011), in which changes in MFA quota restrictions on China were used to instrument competition faced by firms in other countries. A key conceptual difference is that we allow quota constraints from all countries to influence the competition levels in the US, while both Amiti and Khandelwal (2013) and Bloom, Draca, and Van Reenen (2011) implicitly assign zero weight to all countries other than China.

We estimate equation (2.2) via OLS for all 523 textile and apparel exporters in the dataset and report results in Table **??**. Column 1 estimates the model without firm-level controls, while column 2 includes controls for firm-specific energy prices and overall firm-year sales. In both specifications, we find that higher ROW index values are associated with higher export shares. Thus, when India's competitors were more constrained with respect to exporting to the US, Indian firms' exports as a share of total revenue increased. Conversely, with the end of the MFA, Indian firms operating in industries initially protected by third-party quota constraints saw their privileged access to the US erode. This is exactly what one would expect if multilateral resistance matters for exports. Point estimates are statistically

Table 2.5: The Impact of MFA Quotas on Exports of Indian Firms

	(1)	(2)
ROW_{ft}	25.15***	25.31***
	(8.22)	(7.98)
$India_{ft}$	7.13	4.78
	(15.16)	(14.77)
Controls	N	Y
Year Fixed Effects	Y	Y
Firm Fixed Effects	Y	Y
Obs	4312	4312
# Firms	523	523
R^2 (within)	0.02	0.04
Mean Dep. Var	20.92	20.92

Notes: Dependent variable is export share of revenue (0 to 100). Controls in column 2 include sales and firm-specific energy prices. Standard errors that allow for clustering at firm level are reported in parentheses. Asterisks indicate statistical significance at the 1% *** level.

significant at the 1% level with standard errors clustered on the firm.[13]

Surprisingly, we find no evidence that US-India quotas restrained Indian exports (no statistically significant impact on $India_{ft}$). There are at least two interpretations of this null result. First, it is possible that while some US-India quotas had fill rates over 90% under the MFA, the quotas did not in fact constrain export sales. I.e., the quota limits could have been set precisely at equilibrium export supply, so the removal of US-Indian quotas did not relax a constraint at all. Alternatively, the result may be a statistical artifact stemming from the small sample of product categories subject to US-India quotas. To see this point, note that the US tended to coordinate policy across countries within quota categories (Figure 2.3), so goods that were subject to binding US-India quotas were also subject to binding quotas for other countries as well. Thus, when the US eliminated all quotas, Indian firms producing these goods would have been subject to both the direct effect (from lower $India_{ft}$) and the indirect effect (from lower ROW_{ft}), which pull in opposite directions. The simultaneity of the two effects makes the direct effect difficult to pick up because there are few goods for which India was constrained while the rest of the world was not.

Overall, the estimates imply that the end of the MFA delivered a *positive competition* shock to Indian firms, on net. In terms of magnitude, the point estimates imply that if the ROW index increases from 0 to 1, then export shares of Indian firms increase 25 percentage

[13]While we would like to control for correlations across firms operating in the same regulation code, regulations differ at the firm level due to averaging over multiple categories, which makes it difficult to cluster at the regulation level. One robustness check we make is to cluster on the regulation of the highest-sales product within the firm. Results are robust to this procedure, but we report standard errors clustered at the firm level throughout the paper.

points, or 120% on a baseline average of 21% percent export share.[14] Alternatively, evaluated at the mean ROW value over the period 1994-2004 ($ROW_{ft} = 0.1$), we calculate the average firm lost 2.5 percentage points on export sales share as a result of the end of the MFA, or 14% of total export sales.[15] With such large export impacts, we should see firms adjusting emission intensity in response to the elimination of MFA constraints, if exporting matters for environmental performance.

Trade Liberalization Impacts on Emission Intensity

Next, we estimate the impacts of the quota constraints on emission intensity in physical quantity (i.e., kg of CO_2 per unit of production) at the firm-product level. We estimate

$$log\ EQ_{pgft} = \alpha_{pgf} + \alpha_t + \beta^E * India_{gt} + \delta^E * ROW_{gt} + W'_{ft}\Gamma + \epsilon_{pgft} \tag{2.3}$$

where $log\ EQ_{pgft}$ represents log emission intensity for product p belonging to product group g produced by firm f in year t, W_{ft} represents firm or firm-product controls for scale and energy prices, α_{pgf} and α_t denotes firm-product and year fixed effects, and ϵ_{pgft} represents a time-varying unobserved idiosyncratic shock. We expect $\delta^E \neq 0$, since we found strong export impacts from the ROW index in the previous section. In particular, if exporting encourages firms to reduce emission intensity, we expect $\delta^E \leq 0$. However, with no statistically significant effect of US-India quotas on Indian firm exports, we expect $\beta^E = 0$: if US-India quotas do not impact Indian exports in a statistically significant way, then they should not impact emission intensity either.

In Table 2.6, we estimate equation (2.3) via OLS for both exporters (columns 3-4, 7-8) and non-exporters (column 1-2, 5-6) over the period 1994-2007. Estimates are reported based on firm-product regulations (columns 1-4) and firm-average regulations (columns 5-8), where the latter are computed from averaging product-specific regulation over products generated by the firm as in section 2.4. Estimates from non-exporters represent placebo checks, as foreign competition should not affect domestic-oriented Indian firms, absent spillover effects.

We find in columns 3-4 that higher ROW values are associated with lower emission intensities for exporters, statistically significant at the 1%, whether we control for output scale and energy prices (column 4) or not (column 3). There is no corresponding impact from US-India quotas. Additionally, point estimates for neither ROW nor $India$ are statistically significant for non-exporters (columns 1-2). The fact that non-exporters do not respond in the same way is a confirming check that firms operating in product categories that happened to see quotas eliminated post 2004 were not generally trending differently from other firms. A qualitatively similar pattern is found in columns 5-8, with statistically significant negative impacts of ROW on firm-product emission intensity of exporters (columns 7-8), but not non-exporters (columns 5-6).

[14] The point estimates are virtually identical between columns 1 and 2, so it doesn't matter which one we choose to assess economic significance

[15] If export share equals 21% before liberalization, then normalized domestic sales equals 79 (unitless). The change in export sales corresponding to a fall in export share of 2.5 percentage points can be calculated as $\frac{x}{79+x} = (.21 - .025) \longrightarrow x = 18$. Thus, normalized export sales fall from 21 to 18, or by 14%.

Table 2.6: MFA Quota Impact on CO_2 Intensity in Output

	Product-level Regulation				Firm-Average Regulation			
	Non-Exporters		Exporters		Non-Exporters		Exporters	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ROW	0.049	0.140	-0.857**	-0.860**	-0.225	-0.057	-0.463*	-0.463*
	(0.763)	(0.748)	(0.435)	(0.435)	(0.862)	(0.880)	(0.270)	(0.269)
India	0.773	0.349	0.140	0.119	0.737	0.323	-0.016	-0.036
	(0.786)	(0.836)	(0.379)	(0.384)	(0.824)	(0.863)	(0.475)	(0.476)
controls	N	Y	N	Y	N	Y	N	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm-product FE	Y	Y	Y	Y	Y	Y	Y	Y
Obs	1243	1243	5621	5621	1243	1243	5621	5621
# Firm-prds.	280	280	802	802	280	280	802	802
# Firms	226	226	517	517	226	226	517	517
R^2 (within)	0.02	0.04	0.02	0.02	0.02	0.04	0.01	0.02

Notes: Dependent variable is log CO_2 emissions intensity per physical unit of output. Sample includes years 1994-2007. In columns 1-4, ROW and India constraint indices are product-specific, while in columns 5-8 we have aggregated to the firm-year level. "Exporters" includes all firms that export some value in some year. Standard errors that allow for clustering at the firm level are reported in parentheses. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

The results imply that the loss in exports found in section 2.4 translated into higher emission intensity for Indian firms. This is the central result of the paper: trade liberalization between the US and other countries negatively impacted the environmental performance of firms in "third-party" countries. In terms of magnitude, the point estimates in columns 3-4 imply that the average exporter, who saw ROW fall from 0.1 to 0 as a result of MFA quota expiration, *increased* emission intensity at the firm-product level by $e^{(0.1*0.86)} = 9.0\%$.

Possible Channels

There are several possible explanations for the environmental results found in the previous section. First, it may be that increased competition on the export market induced firms to search for cheaper energy sources, which, for some reason, may have been more intensive in CO_2 emissions (Cicala, 2015). We investigate this possibility in Table 2.7, in which we re-estimate equation (2.3) replacing $log\ EQ_{pgft}$ with log CO_2 emissions per mmBTU of energy. If firms purchase cheaper and dirtier forms of energy to produce the same level of physical output, then log CO_2 emissions per mmBTU should increase with lower ROW index values. We find in Table 2.7, that neither for exporters nor non-exporters does this appear to be the case.

Table 2.7: MFA Quota Impact on CO_2 Intensity in Energy

	Product-level Regulation				Firm-Average Regulation			
	Non-Exporters		Exporters		Non-Exporters		Exporters	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ROW	-0.190	-0.251	-0.016	-0.007	-0.166	-0.247	0.071	0.086
	(0.314)	(0.275)	(0.087)	(0.086)	(0.300)	(0.262)	(0.092)	(0.092)
India	0.724*	0.723*	-0.035	-0.024	0.776*	0.773*	0.039	0.050
	(0.425)	(0.425)	(0.084)	(0.083)	(0.448)	(0.450)	(0.118)	(0.116)
controls	N	Y	N	Y	N	Y	N	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm-product	Y	Y	Y	Y	Y	Y	Y	Y
Obs	1205	1205	5599	5599	1205	1205	5599	5599
# Firm-prds.	279	279	790	790	279	279	790	790
# Firms	225	225	513	513	225	225	513	513
R^2 (within)	0.06	0.09	0.05	0.06	0.06	0.09	0.05	0.06

Notes: Dependent variable is log CO_2 emissions intensity per mmBTU of Energy. Sample includes years 1994-2007. In columns 1-4, ROW and India constraint indices are product-specific, while in columns 5-8 we have aggregated to the firm-year level. "Exporters" includes all firms that export some value in some year. Standard errors that allow for clustering at the firm level are reported in parentheses. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Next, several models in the broader trade literature predict that increased competition induces productivity improvements (Bloom, Draca, and Van Reenen, 2011; Amiti and Khandelwal, 2013). If these improvements are Hicks-neutral, then emission intensity should *fall* with export market competition. The environmental impacts in section 2.4 go against this prediction, i.e., we see CO_2 intensity *increasing* with MFA liberalization, which is inconsistent with Hicks-neutral productivity gains. Thus, we conclude that Hicks-neutral productivity enhancements could not explain the results.

Still, there is no theoretical reason that productivity growth must be Hicks-neutral. Perhaps firms adopt technologies that lower variable cost, but increase CO_2 intensity of output. Energy is the only input we have at the firm-product level, so we cannot estimate productivity without putting a lot more structure on the estimation (as in De Loecker et al. (2012)). However, with readily available data, we can assess the possibility of factor-biased technological change by estimating the impact on firm-level capital stock, a decent proxy for technology in manufacturing sectors. If firms adopt new technology to combat increased competition on the export market, we should see it in larger capital stocks.[16] In fact, we find support for this hypothesis in Table 2.8, in which we find that total value of the capital

[16] A caveat here is that capital stock is not the same things as investment. Also, capital is denominated in value, so there could be unobserved changes to the quality or price of capital. Finally, since capital is only reported at the firm level, output has to be denominated in value, so the estimates are inclusive of

Table 2.8: Capital Investments

Dep Var:	Log Capital		Log Capital/Sales	
	(1)	(2)	(3)	(4)
	All Firms	Exporters	All Firms	Exporters
ROW	-0.8**	-0.9**	-1.1**	-1.4***
	(0.4)	(0.4)	(0.5)	(0.5)
India	0.4	0.6	-0.7	-0.3
	(0.4)	(0.4)	(0.6)	(0.6)
Obs	4702	3831	4702	3831
R2	0.08	0.08	0.02	0.02

Notes: Dependent variable is log capital value in columns 1 and 2 and log capital per sales value in columns 3 and 4. Estimates in columns 1 and 3 are based on the entire sample of firms, while estimates in columns 2 and 4 are only for firms that export positive value in some year through the period. Sample includes years 1996-2007. ROW and India constraint indices are weighted averages at the firm-year level. All regressions include firm and year fixed-effects. Standard errors that allow for clustering at the firm level are reported in parentheses. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

stock (columns 1-2) and log capital/sales ratio (columns 3-4) both increase with competition (lower ROW values). If productivity gains are factor-biased, the increase in capital stock could increase CO_2 intensity of output.

A final potential explanation for the increased emission intensity is that firms sell higher quality varieties on the export market, and these high-quality varieties generate lower emission intensities. Thus, when Indian firms are crowded out of the US market, they shift the variety mix (within product code) towards lower quality, higher-emission intensity outputs. Verhoogen (2008) finds evidence of a similar mechanism with respect to labor inputs in Mexico, whereby Mexican firms sell relatively higher quality/higher labor-intensity varieties for the export market relative to the domestic market. We hypothesize that a similar mechanism could generate the environmental impacts documented in section 2.3, if quality is decreasing in emission intensity.

To assess this possibility, we estimate the impact of MFA quotas on two different measures of the quality of firm-product outputs in Prowess. First, we measure quality simply as the unit value of sales, assuming that higher unit value goods are higher "quality." This is a common measure of quality in the literature (Baldwin and Harrigan, 2011). We report estimates in columns 1-2, 5-6 in Table 2.9.

Second, we implement a procedure from Khandelwal, Schott, and Wei (2013) where

output price effects. With these qualifications in mind though, it appears that capital stocks increase when the ROW index falls, which is consistent with competition-induced investments.

Table 2.9: MFA Quota Impact on Quality Measures

	Product-level Regulation				Firm-Average Regulation			
	Log Unit Value		Log Quality		Log Unit Value		Log Quality	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ROW	0.360*	0.346	0.500	0.461	0.420*	0.441*	0.690**	0.589*
	(0.212)	(0.215)	(0.317)	(0.287)	(0.232)	(0.231)	(0.329)	(0.308)
India	-0.418	-0.375	-0.057	-0.499	-0.182	-0.108	0.348	-0.145
	(0.295)	(0.283)	(0.338)	(0.377)	(0.485)	(0.456)	(0.532)	(0.608)
controls	N	Y	N	Y	N	Y	N	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm-product	Y	Y	Y	Y	Y	Y	Y	Y
Obs	6308	6308	6308	6308	6308	6308	6308	6308
# Firm-prds.	1015	1015	1015	1015	1015	1015	1015	1015
# Firms	712	712	712	712	712	712	712	712
R^2 (within)	0.02	0.02	0.00	0.12	0.02	0.02	0.00	0.12

Notes: Dependent variable is log CO_2 emissions intensity per mmBTU of Energy. Sample includes years 1994-2007. In columns 1-4, ROW and India constraint indices are product-specific, while in columns 5-8 we have aggregated to the firm-year level. "Exporters" includes all firms that export some value in some year. Standard errors that allow for clustering at the firm level are reported in parentheses. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

quality is computed as the residual of the regression

$$\ln x_{pgft} + \sigma \ln p_{pgft} = \alpha_{pgf} + \alpha_t + \epsilon_{pgft} \tag{2.4}$$

where x_{pgft} is quantity, p_{pgft} is price, and σ is the elasticity of substitution between products. Khandelwal, Schott, and Wei (2013) show how this specification results from CES preference structure. With enough data, σ could be estimated; however, we follow Khandelwal, Schott, and Wei (2013) and adopt the mean σ calculated by Broda and Weinstein (2006) for the textile and apparel sector ($\sigma = 4$). Our results are robust to varying $\sigma \in [2, 10]$. Intuitively, the procedure assigns higher quality to goods with higher demand, conditional on price. Hence, netting out price effects, if a product generates more sales, then it has higher quality. We refer to this second measure as "quality" and report estimates in columns 3-4 and 7-8 of Table 2.9

Starting in columns 1-4, we estimate MFA impacts on unit values and the Khandelwal, Schott, and Wei (2013) quality measure using the product-level regulation variables $India_{gt}$ and ROW_{gt}, controlling for scale and energy prices in columns 2 and 4. The point estimate on the ROW index is positive in columns 1-4, which is consistent with the case in which higher export share shifts production towards higher quality, though only marginally significant. Moving to columns 5-8, we re-estimate the model using the firm-average regulation as the input variables ($India_{ft}$ and ROW_{ft}). These estimates account for spillovers across product groups within the firm, as they take the overall constraints of the firm as the relevant measure. Here, we find that the point estimate on the ROW index is positive and statistically significant at the 10% level. Additionally, the results hold both for unit values and the Khandelwal, Schott, and Wei (2013) quality measure.

Overall, the evidence suggests that the end of the MFA, as measured by a decline in ROW_{ft}, induced quality-downgrading by Indian firms. Evaluating the impact implied by our preferred specification (column 8) at the average pre-2005 regulation in the sample ($ROW = 0.1$), we estimate the average firm lowered quality $e^{(0.1*0.589)} = 11\%$ due to MFA liberalization. If higher quality products generate lower emissions per unit of output, than this reduction in quality could explain the results in section 2.3.

2.5 Concluding Remarks

Previous work has argued that trade liberalization reduces emission intensity of production in participating countries, either through endogenous regulation or productivity growth. Neither strand of literature considers "third-party" effects on countries peripheral to the liberalization. If competition matters for exports, and exporting matters for emission intensity, then emission intensity reductions in countries that gain market share may be offset by emission intensity increases in countries that are crowded-out of the newly liberalized markets.

We present the first estimates of this effect in the literature and find that it can be quite large. Exploiting quasi-natural variation arising from the elimination of quota constraints under the MFA, we find that Indian exporters in Prowess lost on average 14% export sales as a result of liberalized trade between the US and India's competitors. This loss of export sales was accompanied by an increase in CO_2 intensity of 9%. The results do not appear to be due to fuel-switching, but there is suggestive evidence that capital investments and/or increased output shares devoted to low-quality/high-emission-intensity varieties may have played a role.

Chapter 3

The Impact of Agricultural Biotechnology on Supply and Land-Use

With Steven Sexton and David Zilberman

3.1 Introduction

Meeting growing agricultural demand despite severe resource constraints is among the greatest challenges of the 21st century. New evidence on the environmental cost of land-use change has raised the stakes, suggesting that externalities associated with cropland expansion are more costly than previously understood (Fargione et al., 2008). Stagnating crop yield and increasing demand from growing populations, rising meat demand in transition economies, and increasing biofuel production create tradeoffs between environmentally costly land conversion and higher food prices (Rajagopal et al., 2007). Like manna from heaven, any technology that boosts yield per hectare helps navigate this neo-Malthusian dilemma by increasing supply without converting lands to agriculture.

In this paper, we provide new evidence of how genetically engineered (GE) seeds have increased aggregate supply, reduced the agricultural footprint, eased pressure on prices, and abated greenhouse gas (GHG) emissions. The principle GE traits of first-generation biotechnology were intended to improve pest control. First commercialized in 1996, insect resistant and herbicide tolerant seeds allow farmers to better control pests at lower cost, generating higher yields. To the extent the GE gene increases yield per hectare, the technology not only increases supply and lowers prices, but also reduces demand for new cropland: without GE technology, greater agricultural land-base would be needed to meet demand. In this sense, GE can be said to have preserved lands and "saved" GHG emissions associated with land-use change.

There exists a large literature estimating GE impacts based on small-scale farm trials or survey data, which finds moderate yield and pesticide impacts, on average.[1] But to our

[1] In two recent surveys, Carpenter (2010) and Finger et al. (2011) review 49 and 203 studies of GE yield impacts, respectively, all based on farm-level data. Both reviews find substantial yield impacts across all GE

knowledge, there are no macro-level econometric studies of yield or price effects.[2] Two factors motivate our interest in a macro (i.e., country-level) analysis. First, while micro-level studies hold other inputs constant in order to estimate internally valid impacts of the GE gene itself, the full impact of GE adoption also includes induced changes to complimentary inputs, including variable inputs like fertilizer, water, pesticides etc, and fixed inputs like farmer education and land quality. For an aggregate assessment of the technology, we want to include such complimentary inputs effects. Panel analysis controls for time-invariant differences between adopters and non-adopters, but allows for endogenous changes to complimentary inputs.

Second, while micro-level analyses focus on the impact of switching from traditional technology to GE (i.e., the intensive margin), aggregate data allow us to analyze an extensive margin owing to expansion of production onto previously unprofitable lands. The extensive margin has important implications for supply and commodity prices. If GE technology enables production on extensive margin lands, then the change in supply caused by GE includes not only the yield gain on the intensive margin, but all of the production on the extensive margin as well. Thus, taking the extensive margin into account, the supply and price of GE technology are larger than previously realized. It is beyond the scope of the paper to attribute increased land-use to the GE technology empirically, but we can decompose overall GE adoption into intensive margin lands that switch from traditional technology, and extensive margin lands that were previously devoted to some other purpose. This decomposition allows us to assess the potential importance of this previously unrecognized source of supply gains.

In terms of land-use, though the extensive margin means that more (marginal) lands enter production, overall land inputs would still have to increase to meet a fixed demand if the intensive margin yield boost from GE were not available. Abstracting from the extensive margin effect, we can compute "saved" lands as the difference between hectarage needed to meet observed demand without the GE yield boost and the observed hectarage. The extensive margin has implications for land-use as well, though the predictions are not unambiguous. If extensive margin lands come from previously marginal lands, then land-use saving impacts are larger with the extensive margin effects. Our reasoning here is that without these marginal lands, extensive margin production would have had to come from converting nonmarginal lands, which generate more GHG emissions when converted to cropland. If, however, extensive margin lands come from deforestation, then GHG emissions savings might be smaller with the extensive margin, or even negative, as converting forest to agriculture releases considerable GHG emissions.

We explain both the intensive and extensive margins with a simple adoption model that yields clear predictions and guides estimation. As in Lichtenberg and Zilberman (1986), damage control agents (here, GE) raise the marginal product of complementary inputs and reduce risk, thus increasing yield per hectare. Additionally, marginal lands – on which pest pressure is initially too high to farm profitably without the GE technology – will be brought into production once GE becomes available.

We first estimate the yield effect of GE using a cross-country panel of annual hectarage

crops, for the most part.

[2] Other than a related paper from two of the same authors (Sexton and Zilberman, 2011).

and output that takes into account complimentary input effects. Our approach builds on the work of Sexton and Zilberman (2011), which also estimates yield, price, and land-use-saving effects of different GE crops from a country-level panel. The novel features of our work here is that we use a longer panel and estimate a different specification from Sexton and Zilberman (2011) that controls for inter-temporal variation in crop area and land devoted to GE technology. Endogenous input-use and technology adoption at the farmer level does not threaten identification as long as the timing of changes in national regulation does not correlate with trends in variables that correlate with both yield and adoption – education, risk preferences, etc – and so long as individual farmer adoption decisions are not correlated with time-specific idiosyncratic deviations in yield-effecting characteristics, e.g., rainfall. We argue that access to the technology is exogenous, since the licensing of GE technology is largely driven by political concerns (Just, Alston, and Zilberman, 2006). Furthermore, based on results from Imbens and Wooldridge (2007), we argue that farmer-level adoption is unlikely correlated with time-varying shocks, since adoption for the most part monotonically increases.

Next, we derive an algorithm for quantifying the extensive margin based on the adoption model and decompose total GE hectarage into intensive margin and extensive margin lands. We then compute supply and price effects with and without the extensive margin, and land-use and GHG effects from just the intensive margin. We estimate that in 2010, GE technology increased the world supply of corn between 5-12%, cotton 15-20%, and soybeans 2-40%. Given a range of estimated elasticities of demand and supply in the literature, these supply impacts translate into 5-19% lower corn prices, 19-33% lower cotton prices, and 3-66% lower soybean prices. We also compute the same effects based on other yield estimates from the literature and find that our estimates imply somewhat higher impacts, which is to be expected since we take complementary-input and extensive margin effects into account. Furthermore, we find that absent the intensive margin yield effects, farmers would have needed to convert another 5 million hectares, 6 million hectares, and 2 million hectares to corn, cotton, and soybeans, respectively, to match observed 2010 output. Employing the oft-cited Searchinger et al. (2008) land-use-change GHG release figure, these hectarage conversions translate into 0.15 Gt of averted GHG emissions, which is, for comparison, equivalent to about 1/8th the annual emissions from automobiles in the US.

Together, these results suggest that the first generation of GE technology significantly increased crop production, lowered crop prices, and preserved natural land. These effects imply the poor likely disproportionately benefit from GE technology since they spend a relatively large share of their incomes on food. Additionally, as Schelling (1992) famously argued, the poor are predicted to suffer the most from climate change because they live in exposed areas and lack the means to adapt. To the extent GE technology lowers GHG emissions, it also benefits the poor by averting costs related to climate change.

3.2 Model

The first generation of agricultural biotechnology introduced insect resistant (IR) and herbicide tolerant (HT) traits into 3 principle row crops in order to mitigate crop damage from insects and weeds, respectively. There have been several applications of the IR trait thus far, having been inserted into corn, cotton and soybean.[3] The most notable trait causes plants to produce the naturally occurring chemical Bacillus thuringiensis (Bt), which is toxic to common agricultural pests, like the European corn borer, but harmless to humans and relatively environmentally benign. In producing the toxin, which has been applied to plants for nearly a century and is employed in modern organic farming, GE crop plants fend off pests without chemical applications by farmers. HT crops express tolerance to glyphosates, a class of broad-spectrum, low toxicity herbicides that includes Round-Up, a Monsanto product employed also in residential settings. Such tolerance, introduced into corn, soybeans and canola, allows farmers to more easily control weeds. Absent HT varieties, farmers must rely more heavily on pre-emergence weed control, like tilling operations, and on more toxic and narrow spectrum chemicals that can target weeds without impacting the crop plant.

The IR and HT traits can be modeled as damage control agents that reduce the fraction of crop lost to pests. The framework was first introduced by Lichtenberg and Zilberman (1986) to model pesticide adoption, and subsequently applied to GE by Qaim and Zilberman (2003). A wide range of applications followed and are reviewed by Qaim et al. (2009) and Benjamin, Sithole-Niang, et al. (2013). We apply the generalized framework from Sexton and Zilberman (2011) to show how adoption boosts supply on the intensive margin through both a gene effect and complementary-input effects, and along the extensive margin by expanding the range of land that can be profitably farmed.

Production occurs on lands that differ only with respect to pest pressure, denoted by n. The pest pressure at a location may be measured by the average number of pests absent the use of any mitigating agents (like pesticide). Farmers have access to two seed technologies indexed by i with $i = 0$ denoting traditional seed varieties, and $i = 1$ denoting GE varieties. GE varieties are considered damage control inputs that affect yields only indirectly by reducing the fraction of crops lost to pests, which affects both the mean and risk of production. We assume a Just and Pope (1978) and constant return to scale production function following Just and Zilberman (1988). Thus production per unit of land with technology i is denoted by

$$y_i = f(z_i)\, g(i; x_i, n) + h(z_i, i; x_i, n)\, \epsilon \tag{3.1}$$

The deterministic part of the production function is the product of expected potential output $f(z_i)$, the average output without pest damage, and the expected efficacy of production $g(i; x_i, n)$, the share not lost to pest damage. Expected potential output is a concave function of inputs like fertilizer per unit of land z_i that increases output directly.[4] Expected efficacy $g(i; x_i, n)$, lies between 0 (complete crop destruction) and 1 (no pest damage) and is increasing in pesticide use and decreasing in pest pressure.[5] It is assumed that on average pest damage is lower under the GE technology than under the traditional technology $g(0; x, n) < g(1; x, n)$. The stochastic element of the per unit of land production function, $h(z_i, i; x_i, n)\,\epsilon$, is multiplicative in ϵ a random variable with zero mean and variance σ^2. It

[3]Another substantial feed crop that has adopted GE is rapeseed, but the hectarage and impact is much less substantial and therefore it is not addressed in this paper

[4]We assume that when the farmers evaluate the GE technology they consider the number of seeds to be constant under both technologies. This seems reasonable because seeding density tends to depend heavily on cultural practices exogenous to the seed technology choice.

[5]x_i is a measure of the pesticides that are substitute to the GE technology, the pesticides that are part of the GE technology (roundup in the case of herbicides tolerant varieties) are considered part of it.

is assumed that risk is increasing with fertilizer use, i.e, $\frac{\partial h}{\partial z_i} \geq 0$, and the pest infestation $\frac{\partial h}{\partial n} \geq 0$ and declining with the level of the pesticides use $\frac{\partial h}{\partial x_i} \leq 0$. It is also reasonable to assume that the risk is smaller after adoption of GE technology (Bennett et al., 2013) so that $h(z, 0; x, n) > h(z, 1; x, n)$.

Let p, w and v denote exogenous prices of outputs, pesticides, and productive inputs, respectively.[6] When farmers initially evaluate the GE technology, they assume seed density per hectare to be constant under both technologies, because changes in seeding density is associated with exogenous cultural practices. The seed cost per unit for technology i is denoted by k_i and it is assumed that $k_1 > k_0$, since seed companies assess technology fees for access to proprietary GE varieties.[7]

Because of the concern for segregation between GE and non-GE products we assume that a small producer makes a choice between the two. As in Just and Zilberman (1988) we assume constant absolute risk aversion, with a coefficient denoted by r, and a normal distribution of risk. Under these conditions, the maximization of expected utility is equivalent to maximization of expected profit adjusted for expected risk cost:

$$V_i = \max_{z_i, x_i} \; pf(z_i)\, g(i; x_i, n) - wx_i - vz_i - k_i - .5rp^2 \left[h(z_i, i; x_i, n)\right]^2 \sigma^2 \tag{3.2}$$

The first term of V_i is expected revenue, followed by the costs of pesticides, fertilizers and seeds, minus the expected costs of risk – the product of variance of profits multiplied by the marginal cost of risk per unit of variance of profits. Solving this optimization problem, where the optimal outcomes are denoted by z_i^* and x_i^* respectively suggests that the optimal fertilizer level under the ith technology is determined by

$$\frac{\partial f(z_i)}{\partial z_i} g(i; x_i, n) = v + rp^2 \frac{\partial h}{\partial z_i} h(z_i, i; x_i, n)\, \sigma^2 \tag{3.3}$$

Equation (3.3) states that optimal level of the fertilizer z_i^* is where marginal contribution of fertilizers to expected revenue is equal to the price of fertilizers plus the marginal increase in risk because of added fertilizer. Comparative static analysis suggests that more fertilizer will be applied with GE than without $z_1^* \geq z_0^*$. Also, as the pest pressure (n) increases, the expected marginal gain from fertilizers declines, and the marginal risk costs of fertilizers increases, so that application of fertilizers declines $(dz_i/dn \leq 0)$. This impact on fertilizer contributes to the decline of expected profits adjusted for risk with the pest population. Since pest damage increases less under GE, the reduction in fertilizer use as pest pressure increases is lower with GE than without it $(dz_0/dn \leq dz_1/dn \leq 0)$.

The optimal level of the pesticide x_i^* is where the marginal contribution of pesticides to expected revenue is equal to their price less the marginal increase in the cost of risk because of pesticides use, i.e.,

$$\frac{\partial g(i; x_i, n)}{\partial x_i} f(z_i) = w + rp^2 \frac{\partial h}{\partial x_i} h(z_i, i; x_i, n)\, \sigma^2. \tag{3.4}$$

[6] Results would be qualitatively the same if we allowed output prices to differ across the two technologies, so we assume identical prices for simplicity.

[7] Pesticides that are complementary to the GE variety are assumed to be applied at a fixed quantity and are considered part of the technology, hence adding to its costs.

Because the GE combats pests, the marginal productivity of the pesticides it substitutes is smaller for every level of pest damage, and, thus, less of these pesticides are applied with GE, i.e., $x_0^* \geq x_1^*$. Because the expected marginal gain from pesticides increases and the marginal risks costs of pesticides declines, application of fertilizers increases with pest pressure ($dx_i^*/dn \geq 0$). This increase is likely to be greater with the non GE ($dx_0/dn \geq dx_1/dn$) which will further increase its relative costs.

Introduction of GE technology has several effects in comparison to the traditional technology. First, there is a "gene" effect of reducing pest damage. Second, there is a related effect of increasing seed cost. Third, GE induces a substitution effect, reducing pesticides use (altogether we assume that "gene" effect dominate the substitution effect on pest damage). Finally, there is an indirect (complementarity) effect of increasing fertilizer use. The "gene" effect (net of substitution) is likely to increase (expected) yield and reduce risk.[8] The indirect effect on fertilizers is likely to increase expected yield and risk further. Furthermore, our analysis suggests that since the GE technology is more capable of addressing the pest pressure, the reduction in expected profits adjusted for risk with GE is smaller than without it ($dV_0^*/dn \leq dV_1^*/dn \leq 0$).

The model enables assessment of the impact of risk and risk aversion on various outcome measures. Without risk ($\sigma^2 = 0$) or under risk neutrality ($r = 0$), conditions (3.3) and (3.4) do not include the risk component. That means that under these conditions fertilizers use will increase and use of pesticides will decline.[9] This means higher (expected) yield and higher variance of profits under risk neutrality. This also means that the net gain from adoption of the GE technology with ($d\left(V_1^* - V_0^*\right)/dn \geq 0$) increases with pest pressure.

Producers adopt the technology that yields highest expected profits (adjusted for risk). Their problem is solved recursively. First, conditional on seed technology choice and pest pressure, they choose variable inputs (pesticides and fertilizer). Then they choose the seed that yields highest expected profits, conditional on optimal input use and provided expected profits are non-negative. Given heterogeneity in pest conditions, adoption follows the threshold model (David, 1969; Feder, Just, and Zilberman, 1985), in which more vulnerable farmers who gain most from a new technology adopt first and aggregate adoption increases over time as the technology improves or costs of adoption fall.

For a given period, the adoption pattern predicted by the model can be depicted as in Figure 1. Line segment AB depicts profit per hectare as a function of initial pest pressure under the traditional technology and line segment CD depicts the same for the new technology. At locations with low pest pressure, it is profitable (from now on we will use the term profitable to mean expected profitability adjusted for risk) to farm under either seed technology, but the conventional technology yields higher profits because crop losses are too small to compensate for the technology fee. Thus, below a threshold n_l, farmers produce using the traditional technology. For pest pressure greater than n_l and less than a threshold

[8] The gene effect need not be positive. Adoption of GE tends to reduce damage of pests targeted by the GE trait. On the other hand if the trait is not introduced in the best local variety there is a yield loss. For example, Benbrook (1999) found that adoption of HT yield tolerance variety may result in "yield drag". We expect that if adoption occurs the damage reduction effect is greater than the variety effect (Qaim and Zilberman, 2003).

[9] This is consistent with the Just and Pope (1978) prediction that use of risk increasing (decreasing) inputs declines (increases) under risk neutrality and when risk is increasing.

Figure 3.1: Adoption of GE Technology

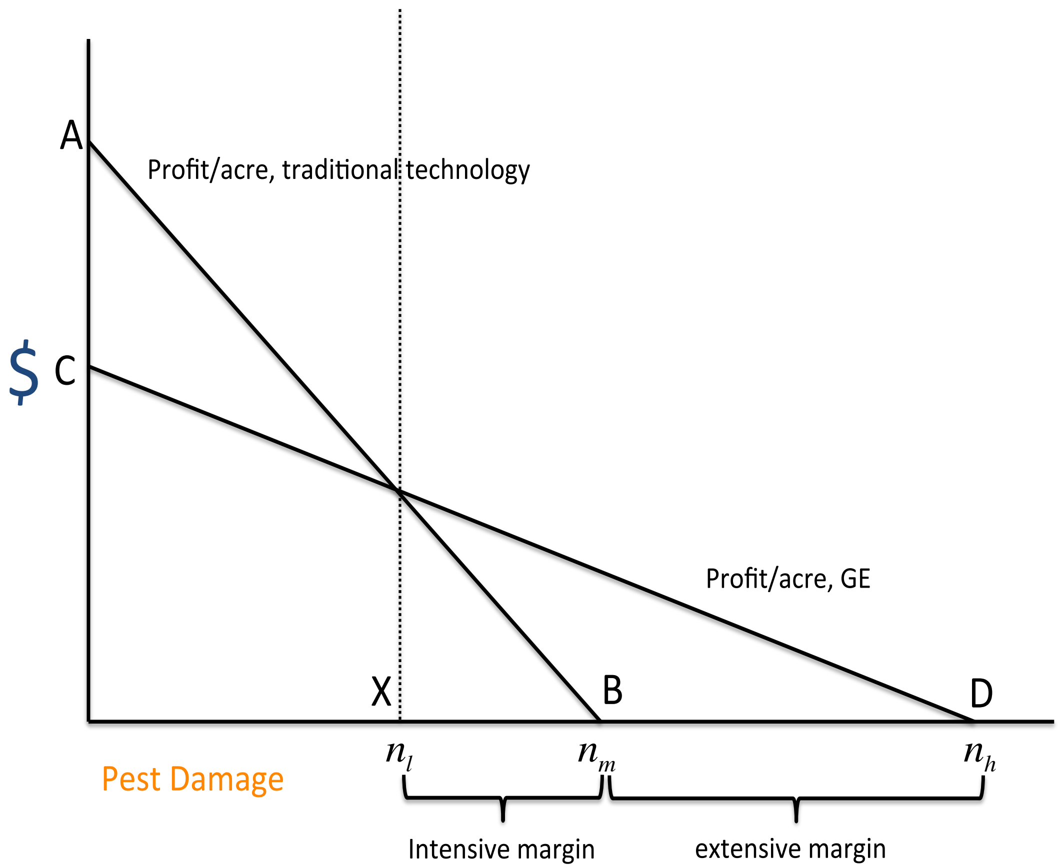

Notes: The figure plots optimized profits on the y-axis against initial pest pressure on the x-axis.

n_m, it is profitable to use either technology, but higher crop losses from greater pest pressure make damage abatement more valuable so that the new technology yields higher profits. Above n_m and below a high threshold of pest pressure, n_h, it is not profitable to produce under the conventional technology, but it is profitable to produce under the new technology. Above n_h, it is not profitable to produce under any technology; such land is unfarmed.

The area between n_m and n_h is where farmers adopt the new technology and recruit into production land that was too marginal to be profitably farmed under the old technology. This area represents the extensive margin, which we are particularly interested in quantifying. The pest pressure levels n_l, n_m, n_h determine the adoption decision, but the overall magnitude of adoption depends on the amount of land associated with each level of pest damage. If, for example, there is a small amount of hectarage between n_l and n_m and large amount of land between n_m and n_h, then the intensive margin is small in magnitude while the extensive margin is large. On the other hand, if there is no land with pest damage below n_m, then there is no extensive margin and all the impact is intensive. We return to this issue when we decompose adoption into intensive and extensive margin lands.

3.3 Estimation

To compute supply, price, land-use, and GHG impacts of GE technology, we first compute intensive margin yield impacts from real-world (i.e. non-experimental) production data. While a vast literature estimates yield parameters from farm-level survey or experimental data, we use a cross-country time series of adoption and yield to assess the overall impact of the technology. Our estimates can be interpreted as the average treatment on the treated effects and are inclusive of all complimentary input effects. One reason that previous work has focused on farm-level data is that until recently, there has not been enough annual observations with positive adoption levels at the country level to consistently estimate macro yield parameters; however, with initial commercialization in 1996, we now have sufficient "treatment" exposure to estimate impacts from within-country time series variation in GE adoption.

In Figure 3.2, we plot world aggregate adoption by crop over time. Area planted to GE varieties was provided by Graham Brookes, who compiled the data from the International Service for the Acquisition of Agri-Biotech Applications. We can see that GE adoption scaled incredibly fast. In 2010, 15 years after the commercialization of GE technology, GE corn accounted for 42 million ha worldwide across 14 countries, representing 25% of world corn hectarage, GE cotton accounted for 19 million ha worldwide across 10 countries, representing 60% of total cotton hectarage, and GE soybeans accounted for 72 million ha worldwide across 9 countries, representing 70% of total soybean hectarage.

Though GE adoption was rapid, enthusiasm for GE was not shared equally across crops and countries. In Table A.3, we report adoption by crop and country for the years 2000 and 2010. The table includes all 26 countries that planted any GE seed from 1996 to 2010. Cumulative hecatres over the entire adoption period planted to GE seed in any of the three crops is reported in the last column. The US is by far the largest adopter, accounting for 58% of total cumulative adoption, with Argentina and Brazil following with a combined 30% of total adoption, mostly in soybeans. India and China follow next, with a combined 7% of total adoption, mostly in cotton. Outside of these 5 major adopters, 21 other countries had planted some GE seeds in at least one of the 3 crops by 2010, though at much lower levels. Our strategy is to exploit this variation in adoption rates over time to estimate yield effects.

While country-level data are well-suited to estimating global effects – where "global" means inclusive of all corresponding changes induced by adoption – estimating production function parameters at the aggregate level is problematic. Felipe and Fisher (2003) show aggregate production function are weighted sum of micro level production function that may change over time and reflect spatial and dynamic variability. Thus, estimating parametric coefficients with aggregate data is feasible only under restrictive assumptions. In agriculture, however, a large literature uses macro data to identify nonparametric productivity coefficients (Huffman and Evenson, 1992). Following this tradition, we develop a simple method to recover the yield effects of adoption from aggregate national data by decomposing aggregate

Table 3.1: Area Planted to GE seeds (Millions of Ha) by Country

	Cotton		Corn		Soybeans		All Crops
	2000	2010	2000	2010	2000	2010	1996-2010
United States	3.83	4.11	8.05	28.19	18.21	29.35	579.51
Argentina	0.03	0.44	0.58	2.75	8.64	18.02	188.78
Brazil	0.00	0.37	0.00	7.51	1.30	18.36	110.63
India	0.00	9.40	0.00	0.00	0.00	0.00	36.30
China	1.22	3.45	0.00	0.00	0.00	0.00	34.38
Canada	0.00	0.00	0.45	1.13	0.21	1.03	17.53
Paraguay	0.00	0.00	0.00	0.00	0.09	2.67	16.47
South Africa	0.02	0.01	0.08	1.88	0.00	0.31	11.73
Uruguay	0.00	0.00	0.00	0.10	0.00	0.86	3.61
Bolivia	0.00	0.00	0.00	0.00	0.00	0.78	3.06
Australia	0.17	0.21	0.00	0.00	0.00	0.00	2.04
Philippines	0.00	0.00	0.00	0.54	0.00	0.00	1.81
Mexico	0.03	0.05	0.00	0.00	0.00	0.02	0.68
Spain	0.00	0.00	0.03	0.08	0.00	0.00	0.61
Romania	0.00	0.00	0.00	0.00	0.04	0.00	0.44
Burkina Faso	0.00	0.26	0.00	0.00	0.00	0.00	0.38
Colombia	0.00	0.04	0.00	0.04	0.00	0.00	0.25
Honduras	0.00	0.00	0.00	0.01	0.00	0.00	0.05
France	0.00	0.00	0.00	0.00	0.00	0.00	0.03
Czech Republic	0.00	0.00	0.00	0.00	0.00	0.00	0.03
Portugal	0.00	0.00	0.00	0.00	0.00	0.00	0.02
Bulgaria	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Germany	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Slovakia	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Egypt	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Poland	0.00	0.00	0.00	0.00	0.00	0.00	0.00
# Adopting Countries	6	10	6	14	7	9	26

Notes: Table values represent millions of Ha planted to GE technology in the given crop-country-year. Countries are sorted on cumulative adoption over the entire period, which is reported in the last column. Data comes from Graham Brookes.

Figure 3.2: World Area Planted to GE Seeds by Crop

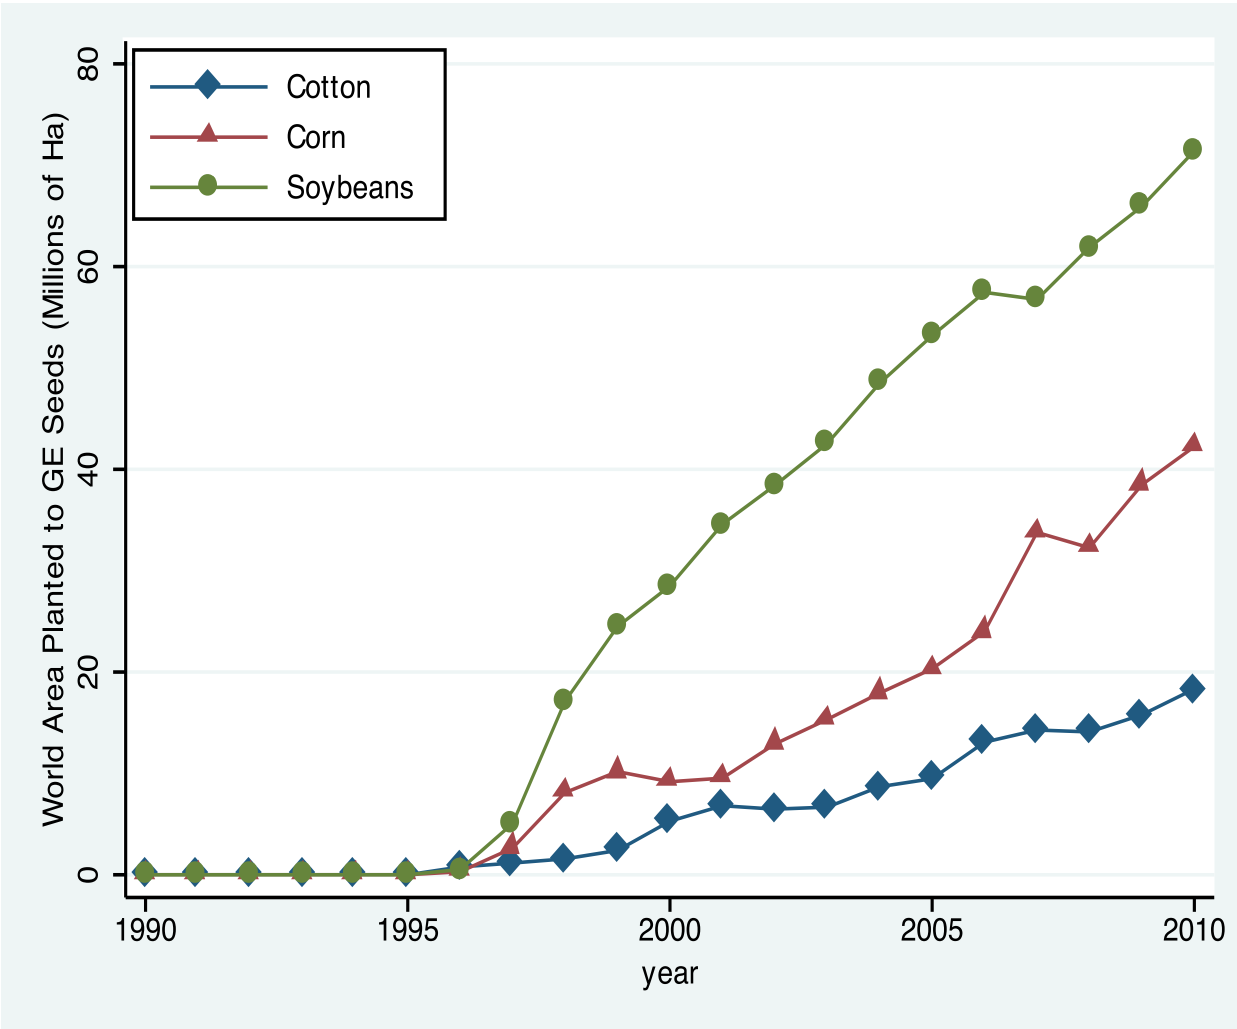

Notes: Authors' own calculation derived from data from Graham Brookes

output into sources of variability in the data. This decomposition approach is similar to the one used in Just et al. (1990). We use the land share of GE as a measure of adoption, while dummy variables allow for variation over crops, countries, and time (Feder, Just, and Zilberman, 1985). Statistical power constrains our ability to estimate parameters beyond average effects, but with more data one would be able to further investigate the impact of interaction between factors.

According to equation (3.1), yield per hectare depends on prices, pest pressure, and the technology choice, with no scale effects. The assumption of perfect input and output markets implies prices effects are captured by year dummy variables. The time-invariant component of pest pressure is likewise absorbed by country dummies, along wth all other time-invariant unobservable determinants of adoption. Remaining agnostic on the precise functional form in (3.1), the model implies we can write the deterministic component of yield per hecatre y_{cti} in country c, in year t with technology i as the sum of a technology effect β_i country effect α_c and time effect γ_t:

$$y_{cti} = \beta_i + \alpha_c + \gamma_t \tag{3.5}$$

where γ_t absorbs the intertemporal impact of prices common to all countries and α_c absorbs time-invariant determinants of yield, such as land quality and pest pressure. The structural technology parameters β_i are the parameters of interest.

Total output for a given crop per country-year (Q_{ct}) can be written as the sum of output produced with technology $i \in 0, ...I$:

$$Q_{ct} = \sum_{i=0}^{I} Q_{cti} = \sum_{i=0}^{I} y_{cti} L_{cti} \tag{3.6}$$

where L_{cti} is land planted to technology i in country c in time t. Substituting for y_{cti}, we have:

$$Q_{ct} = \sum_{i=0}^{I} [\beta_i + \alpha_c + \gamma_t] L_{cti} \tag{3.7}$$

Sexton and Zilberman (2011) estimate the technology parameters with a fixed effect model:

$$Q_{ct} = \delta_0 L_{ct} + \delta_1 L_{ct1} + \gamma_t D_t + \alpha_c D_c + \nu_{ct} \tag{3.8}$$

with $i = 0$ denoting traditional seed technology $i = 1$ denoting GE technology and D_t and D_c representing dummies for years and countries. While the fixed-effect model (3.8) controls for country and time specific unobservables that correlate with adoption, it subsumes country-time specific hectarage deviations in the error term, which correlate with the adoption decision.[10] This correlation generates bias in the δ_i's, as they pick up some of the country and time specific effects, which multiply the deviations in ν_{ct}. The direction of the bias is ambiguous, but since bigger countries adopted GE more heavily, it is likely that Sexton and Zilberman (2011) overestimate the technology effect.

Departing from Sexton and Zilberman (2011), we divide (3.7) through by total hectarage and simplify:

$$y_{ct} = \beta_0 s_{ct0} + \beta_1 s_{ct1} + \gamma_t + \phi_c + \epsilon_{ct} \tag{3.9}$$

[10] To see this, note that time and country dummies can be rescaled with time and country averages

$$Q_{ct} = \delta_0 L_{ct} + \delta_1 L_{ct1} + \gamma_t D_t \bar{L}_t + \alpha_c D_c \bar{L}_c + \nu_{ct}$$

but that a direct derivation from (3.7) delivers

$$Q_{ct} = \delta_0 L_{ct0} + \delta_1 L_{ct1} + \gamma_t D_t L_{ct} + \alpha_c D_c L_{ct} + \epsilon_{ct}$$

The use of L_{ct} instead of L_{ct0} does not matter, since it just alters the definition of the excluded category. But multiplying the time and country effects $\gamma_t D_t$ and $\alpha_c D_c$ by $\bar{L}_t$ instead of L_{ct} means that country-time specific deviations appear in the error term multiplied by the time and country effects γ_t and α_c:

$$Q_{ct} = \delta L_{ct0} + \delta_1 L_{ct1} + \gamma_t D_t \bar{L}_t + + \phi_c D_c \bar{L}_c + \underbrace{\gamma_t D_t \left(L_{ct} - \bar{L}_t\right) + \phi_c D_c \left(L_{ct} - \bar{L}_c\right) + \epsilon_{ct}}_{=\nu_{ct}}$$

The country-time deviations from averages $L_{ct} - \bar{L}_t$ and $L_{ct} - \bar{L}_c$ in ν_{ct} are obviously correlated with L_{ct0} and L_{ct1}.

where y_{ct} is yield per hectare, s_{ct0} and s_{ct1} represent shares of hectarage devoted to traditional and GE technology, respectively (again, $i = 0$ denoting traditional seed technology $i = 1$ denoting GE technology), and ϵ_{ct} represents an idiosyncratic shock to country-level yield per hecatre. Coefficients β_0 and β_1 correspond directly with the structural parameters in (3.5) and are recovered by estimating (3.9) via OLS.

In estimating (3.9), as in Sexton and Zilberman (2011), the key identification assumption is that country-level adoption shares s_{ct0} and s_{ct1} are orthogonal to the time-varying shocks ϵ_{ct}. Though adoption is not randomly assigned, we argue that unconfoundedness is likely to hold. There are two components of country-level adoption. First, governments have to approve the technology crop by crop. Just, Alston, and Zilberman (2006) argue that this process is driven largely by political concerns, and hence can be taken as exogenous to unobservable determinants of yield. Second, conditional on government approval, farmers adopt. While it has been shown that GE-adopting farmers are more educated (Crost et al., 2007), and less risk-averse (Liu, 2013), and so should have systematically higher yields, to the extent that unobservable farm or farmer characteristics are time-invariant, they are absorbed by the country dummies. Given endogenous selection at the farm-level, our aggregate estimates should be interpreted as average treatment on the treated (ATT) measures, where the "treated" here refer to adopting farmers within adopting countries.

One remaining concern is that time-varying shocks to prices or pest pressure – possibly through weather – could bias estimates in equation (3.9). But because adoption tends to monotonically increase over the sample, time-varying shocks likely do not influence the adoption decision too much. As Imbens and Wooldridge (2007) noted, if $s_{ct1} > s_{cr1}$ for $r < t$, then strict exogeneity is a reasonable assumption. Intuitively, if farm-level adoption in period r were induced by a stochastic positive shock to an underlying characteristic, like weather, then a stochastic negative shock to the same characteristic at a time $t > r$ should induce switching back to traditional technology in period t. Since reductions in s_{ct1} are rare (only 79 instances out of 4989 possibly country-crop-year observations since 1996), we conclude idiosyncratic shocks to farm or farmer characteristics are unlikely to bias estimates in (3.9).[11]

For each GE crop $\in \{corn, cotton, soybeans\}$, we estimate equation (3.9) with the same data sources as in Sexton and Zilberman (2011), though we extend the panel to include more years. Output and area by crop-country-year for 1990-2010 come from FAO Stat. Descriptive statistics are reported by crop and adopting vs non-adopting countries in Table 3.2, where "adopting countries" have positive GE area for some year for the given crop. The panel for each crop includes all GE adopters and all other 100 top-producing countries.[12] We drop observations with 0 output (134 for cotton, 43 for corn, 151 for soybeans) generating an unbalanced panel for each crop. We find in Table 3.2 that adopting countries have higher yields per hectare and higher harvested area in all three crops. Some of the difference is attributable to selection bias into adoption at the country-level, while some (potentially) is

[11] A similar justification based on sequential exogeneity was made in Sexton ("Automatic Bill Payment, Price Salience, and Consumption: Evidence from Residential Electricity Consumption")

[12] There are only 93 soybean producers and 95 cotton producers in the FAO data, so for these crops, we keep the entire sample. There are 173 corn producers, so we censor to only keep the top 100, which includes all 19 GE adopters

due to adoption of GE. We control for country fixed-effects in order to distinguish between the two.

In Table 3.3 we report estimates of equation (3.9) by crop. The regression coefficients for traditional and GE technology correspond directly with the structural yield parameters β_0 and β_1. In Panel A, we begin in columns 1, 4, and 7 by estimating (3.9) via OLS for cotton, corn, and soybeans, respectively. All regressions include year and country fixed-effects. Standard errors are clustered at the country level, so estimates are robust to serial correlation in the error term. For all crops, the coefficients for both traditional and GE yield are individually significant, jointly significant, and statistically different, all at the 1% level. The yield effect can be computed as $\frac{\beta_1-\beta_0}{\beta_0} * 100$ and is reported in Panel B.[13] We find that the yield effect for cotton is 34%, corn is 12%, and soybeans is 3%. By contrast, the yield effects from Sexton and Zilberman (2011) are 65% for cotton, 45% for corn, and 13% for soybeans, again all significant at the 1% level. Our estimates here are smaller than those from the Sexton and Zilberman (2011) specification, but still mostly larger then other studies in the literature (see Qaim et al. (2009)). This is likley due to the fact that our estimates are based on real-world outcomes, and hence are inclusive of all complimentary input effects. The one exception is that we find almost no yield impact of GE soybeans. The small impact on soybeans are possibly explained by "yield drag" resulting from the HT trait. Indeed much of the benefits from planting HT crops is due not to any inherent yield advantage but the lower cost of managing weeds. Considerable evidence exists that the presence of HT traits in a crop actually lowers yield (Benbrook, 1999). However as we show in latter sections, even if the yield gain from GE soybeans is low, GE may still boost soybean supply via the extensive margin.

Next, in columns 2, 5, and 8 of Table 3.3, we estimated weighted regressions to account for difference in country size. Weights correspond to total agricultural area of the country. Point estimates are still all significant, jointly significant, and different from each other. Implied yields increase substantially for all crops, reflecting the correlation between country-level adoption and country size: larger countries adopted more heavily (eg US, China, India, Brazil), and their yields increased due to GE, so the average GE effect seems larger when we weight by size. Finally, in columns 3, 6, and 9, we log the dependent variable and weight by size. While the model does not call for logs, it is a common specification in the literature, so we include it for comparison (Lobell, Schlenker, and Costa-Roberts, 2011).

All three specifications deliver significant GE yield impacts. The specifications reported in columns 1, 4, and 7 provides yield effects estimates that are below those reported in columns 2, 5, and 8 and somewhat above those in columns 3, 6, and 9. As our baseline specification corresponds directly with the statistical model and the results lie between the two other specifications, we prefer the baseline and use the implied yield impacts in columns 1, 4, and 7 to derive supply, price, and land-use effects, though we note that other specifications common in the literature (i.e., columns 2-3, 5-6, 8-9) also deliver positive GE yield impacts, though at significantly different magnitudes.

[13] In the case of logged dependent variable, the yield impact is $\frac{\exp^{\beta_1} - \exp^{\beta_0}}{\exp^{\beta_0}} * 100$

Table 3.2: Summary Statistics of Adopters and Non-adopters by Crop (1990-2010)

	Cotton		Corn		Soybeans	
	Non-Adopters	Adopters	Non-Adopters	Adopters	Non-Adopters	Adopters
Yield (Lbs/Ha)	1.37	2.11	3.09	5.13	1.44	2.10
	(0.91)	*(1.06)*	*(2.67)*	*(2.62)*	*(0.76)*	*(0.53)*
Total Area (Mil Ha)	0.15	2.12	1.02	3.21	0.26	5.69
	(0.38)	*(2.93)*	*(3.03)*	*(6.85)*	*(1.24)*	*(9.18)*
GE Area (Mil Ha)	0.00	0.59	0.00	0.71	0.00	2.89
	(0.00)	*(1.50)*	*(0.00)*	*(3.34)*	*(0.00)*	*(6.73)*
# Countries	85	10	81	19	83	10

Notes: Top value in each cell reports the mean over the period, 1990-2010, for the specified country-category. "Non-Adopters" are the countries that report 0 GE planted area throughout the period, while "Adopters" report positive GE planted area in some year over the period. Standard deviations are reported below in parenthesis.

Table 3.3: Yield Effects of GE Technology

Crop:	Cotton			Corn			Soybeans		
Dep Variable:	yield (1)	yield (2)	log(yield) (3)	yield (4)	yield (5)	log(yield) (6)	yield (7)	yield (8)	log(yield) (9)
Panel A: Estimated Coefficients									
β_0	2.13***	0.31***	-0.79***	10.61***	5.93***	1.65***	1.95***	0.71***	-0.25***
	(0.00)	(0.12)	(0.07)	(0.15)	(0.10)	(0.04)	(0.02)	(0.06)	(0.04)
β_1	2.86***	0.78**	-0.63***	11.93***	7.39***	1.73***	2.01***	0.88***	-0.24***
	(0.25)	(0.35)	(0.18)	(0.28)	(0.23)	(0.12)	(0.15)	(0.12)	(0.09)
Panel B: Implied Yield Impact									
	34%	152%	17%	12%	25%	8%	3%	24%	1%
Number of Obs.	1864	1864	1864	1955	1955	1955	1805	1805	1805
R squared	0.96	0.97	0.93	0.98	0.99	0.98	0.96	0.98	0.92
Mean of Dep. Var.	1.46	2.03	0.53	3.60	4.11	1.19	1.52	1.74	0.46
Analytical Weights	none	ag. area	ag. area	none	ag. area	ag. area	none	ag. area	ag. area

*Notes: OLS coefficients are reported by crop for cotton (columns 1-3), corn (column 4-6) and soybeans (columns 7-9) in Panel A. In columns 1-2, 4-5, 7-8, dependents variable is yield per hectare, and in columns 3, 6,and 9, log of yield per hecatre. Standard errors are clustered at the country level and reported in parentheses below. Asterisks indicate statistical significance at the 1% ***, 5% ***, and 10% * levels. All regressions are estimated for 1990-2010 and include country fixed-effects and year fixed effects. In Panel B, yield impact is computed as* $\frac{\beta_1-\beta_0}{\beta_0} * 100$ *when the dependent variable is yield, and* $\frac{\exp^{\beta_1} - \exp^{\beta_0}}{\exp^{\beta_0}} * 100$ *when the dependent variable is log yield. Columns 2-3, 5-6, and 8-9 weight by agricultural area.*

3.4 Extensive Margin

The previous section estimates the increase in yields associated with switching from traditional technology to GE (intensive margin switching), though the adoption model also predicts that GE brings more land into production by extending the range of land that can be profitably farmed. If this relationship is causal, as the model predicts, then output on the extensive margin should be credited to the GE technology, thereby increasing the supply effect of GE seeds. As mentioned in the introduction, we estimate the quantity of new land converted to a given crop since GE was introduced, though we cannot attribute this extensification to the GE technology. We will compute supply and price effects under the bounding assumptions that none (all) of the production on the extensive margin is attributable to GE to understand how important the effect could be.

With plot-level data, the task of decomposing the supply effect into intensive and extensive margins is a simple matter of separating the plots that switched from traditional to GE from those newly planted to GE and summing over the yield increases in each group. Since our data is country-level, additional structure is needed to guide the calculation. We appeal to the adoption model from Section 2 again to generate the necessary structure.

To illustrate the strategy, consider again Figure A.3. In some base year – prior to GE entry – the profit curve with traditional technology is given by line segment AB. In a future period, GE becomes available and generates profit curve CD. Suppressing time and country indices, let ΔL_0 and ΔL_1 denote the change in traditional and GE hectarage between two periods, with the total change in area $\Delta L = \Delta L_0 + \Delta L_1$. Total hectareage expands in the figure by $\Delta L = D - B$, and GE hectareage expands by $\Delta L_1 = D - X$, where X indicates the break-even point on GE technology. As described in Section 2, the extensive margin, denoted ext, is given by $n_h - n_m$, or

$$ext = \Delta L = D - B \tag{3.10}$$

Furthermore, the intensive margin, denoted int, is given by $n_m - n_l$, or

$$int = \Delta L_0 = B - X \tag{3.11}$$

Thus, in this case, all we need to compute the intensive and extensive margins are the change in total hectareage ΔL and the change in traditional technology hectareage ΔL_0, which are figures readily computed from the data.

While this simple example illustrates how the adoption model generates enough structure to calculate intensive and extensive hectareage from the aggregate data, the example is not sufficiently general to handle all cases. In particular, we have assumed that the traditional technology profit curve does not change over time. In this case, we have $\Delta L_1 > \Delta L > 0$, and thus equations (3.10) and (3.11) yield the intensive and extensive margins. However, in reality, prices, growing conditions, and policy all change from year to year, which shifts the traditional technology profit curve. If this profit curve shifts concurrently with the entry of GE technology, we could observe $\Delta L > \Delta L_1 > 0$, for example. In this case, the switchover point X would exceed the x-intercept of the original traditional technology profit curve, implying that all traditional technology hectares from the base year remain traditional technology hectares in the future year. I.e., in such a case there is no intensive margin switching. All GE hectares should be counted as extensive margin. Furthermore, it's possible that the traditional technology curve shifts in such that $\Delta L < 0$. In this case, no new lands enter production in the future year, so there can be no extensive margin. In this case, all GE lands should be considered intensive margin.

The general structure for these three cases are presented in Table 3.4. The three cases are distinguished by the ordering of ΔL_1, ΔL, and 0. In the first case (which includes the first example above), $\Delta L_1 > \Delta L > 0$, and there is adoption on both the extensive and intensive margins.[14] For this ordering to occur, it is possible that the traditional profit technology curve shifts in or out, but it must be that X', the observed break-even point in the future period, lies to the left of the initial x-intercept, B. That is, in order for the change in GE hectarage to exceed the change in total hectarage, there must be some intensive margin switching, which implies the break-even point exceeds the initial marginal hectare. In Table 3.4, we illustrate this case in the first row with a small outward shift of the line segment AB to $A'B'$. The column labeled "Ordering" describes the case, and the columns "*int*" and "ext_1" give the calculation of the intensive and extensive margins ("ext_1" indicates extensification with GE technology, while "ext_0" indicates extensification with traditional technology, and $ext = ext_1 + ext_0$). We find that the intensive margin is computed as the negative of the change in traditional technology hectarage $(B - X')$, while the extensive margin is computed as the change in total hectarage $(D - B)$. In the second case, illustrated in row 2, we have a large outward shift in AB such that $X' > B$. In this case, the total extensive margin is given by the change in total hectarage $D - B$, but these hectares are divided between extensive margin traditional hectares, $ext_0 = X' - B$, and extensive margin GE hectares $ext_1 = D - X'$. There are no intensive margin hectares.[15] In the final case, AB shifts in such that the total hectarage decreases. With no new hectares entering production, $ext = 0$, and any GE hectares come from the intensive margin $int = \Delta L_1 = D - X'$.

The three cases in Table 3.4 exhaust the possible outcomes when comparing any post-adoption year to the pre-adoption base year.[16] Using the data described in Section 3, for every country c and year t, we compute the change in total area and GE area (for each crop) between year t and some pre-adoption base year b as $\Delta L_{ct} = L_{ct} - L_{cb}$ and $\Delta L_{ct1} = L_{ct1} - L_{cb1}$, where the base year is defined as the year immediately prior to the first positive value for GE hectarage for the given country-crop observation. Given ΔL_{ct} and ΔL_{ct1}, we classify every country-crop-year as belonging to one of the three cases in Table 3.4 and compute the corresponding intensive and extensive margins according to the formulas in columns 3 and 4.[17] We then sum over the given year to generate world hectarage by crop, divided be-

[14] This will only occur if there is land available on the extensive margin and the land with the highest pest damage has more pest damage than n_h in Figure 1. See Section 2.

[15] Of course, extensification only occurs if there exists lands with pest damage that is greater than point B

[16] A fourth case corresponds to the possibility that AB shifts out so much that traditional technology profits dominate GE profits for any initial pest pressure. In this case, GE hectarage equals 0, so trivially $ext = int = 0$. We leave this case out of Table 3.4 to reduce clutter, but we allow it in the empirical exercise.

[17] The model predicts extensification onto marginal lands that presumably were not used for anything before the introduction of GE. In this sense, the extensive margin is extensive to agriculture overall. However, we want to quantify the extensive margin *to a given crop so that we can compute supply effects by crop. Defining the extensive margin as crop-specific means that extensive margin lands might be coming from any previous employment other than the production of the given crop, including the production of other crops.*

Table 3.4: Computation of the Intensive vs Extensive Margins

(1) Case	(2) Ordering	(3) Intensive Margin (int)	(4) Extensive Margin (ext_1)
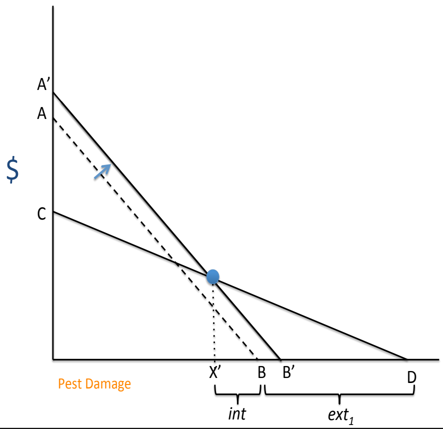	$\Delta L_1 > \Delta L > 0$	$-\Delta L_0 = B - X'$	$\Delta L = D - B$
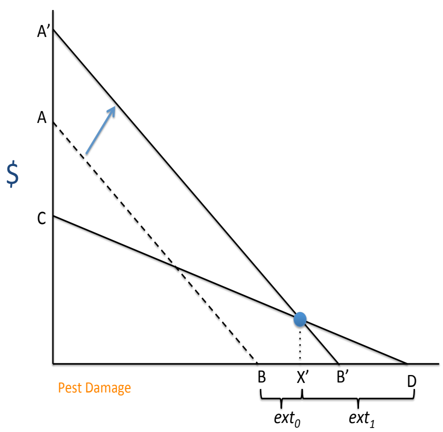	$\Delta L > \Delta L_1 > 0$	0	$\Delta L_1 = D - X'$
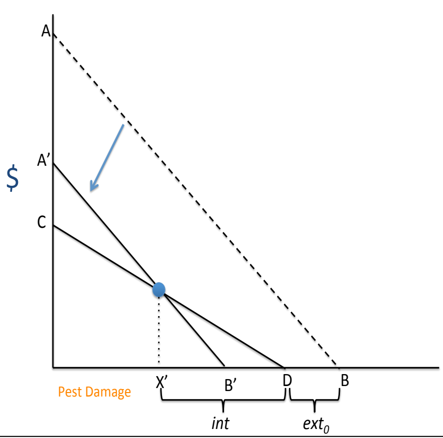	$\Delta L_1 > 0 > \Delta L$	$\Delta L_1 = D - X'$	0

Notes: Intensive and extensive margin lands are computed based on the ordering of $\Delta L_1, \Delta L, 0$. The three cases are summarized in the different rows here. The first column depicts the case graphically. Column 2 reports the ordering. Column 3 reports the formula for the intensive margin. Column 4 reports the formula for the extensive margin.

tween traditional seed technology, GE intensive margin hectarage, and GE extensive margin hectarage. We present results for corn, cotton and soybeans in Figure 3.3.

In Figure 3.3, we find that for corn, most adoption of GE occurred on the intensive margin, with the extensive margin only accounting for 16% of total GE hectarage in 2010. The share of GE corn hectarage in total corn hectarage is not very large (26%), but because total corn hectarage is so large (the largest world hectarage of all crops) absolute extensification is still substantial. In cotton, overall GE cotton adoption rates are much higher (57%), though mostly still on the intensive margin (only 12% extensive margin). By contrast, adoption of GE soybeans has been high (70%) and more concentrated on the extensive margin than the other crops (49%). As a result, soybean hectarage grew more than 50% since the introduction of the GE seed.

The data shows that much of the potential of GE has been realized in cotton and soybean. In the case of cotton, there is a relatively small extensive margin effect, and the adoption of most of the GE has occurred on land previously in production. However, GE cotton is the only GE crop that has been adopted globally, as it did not suffer from bans that apply to corn and soybean. In the case of soybean, the high rate of adoption of GE is attributable to an expansion of the hectarage of the crop (thus the large extensive margin effect), and virtually all of the adoption of GE soybean occurs in the U.S., Brazil, and Argentina. In the case of corn, a majority of corn in the world is located in countries in Europe and Africa that have banned the adoption of GE corn, and therefore overall adoption is below 30% of global hectarage. Because yield per hectare of adopters is higher than that of non-adopters, the GE share of corn is about 43%. Nevertheless, there is a large potential for increased adoption of GE corn if practical bans on the technology are removed.

Breaking down GE area by country, we find that our estimates of extensification are broadly in line with country-specific aggregates. For example, in the case of cotton, 70% of adoption occurred in India and China in 2010, where total agricultural area has actually *declined* slightly since the introduction of GE seeds.[18] However, we estimate that only 20% and 4% of GE cotton adoption respectively for India and China occurred on the extensive margin. With such low GE cotton extensification, it is entirely possible that declining footprints in other crops contributes to overall lower agricultural land base. By contrast, 72% of the extensification effect in soybeans in 2010 is estimated to have come from Brazil and Argentina, where total agricultural land base has increased since 1995. In fact, 86% of extensification across all crops is estimated to have originated in Brazil, Argentina and the US, where total agricultural area has increase by a combined 2 million hectares since 1995. Thus, while total agricultural area has not increased in all GE adopting countries, it has increased where our model predicts large extensifications from GE. While a complete analyses of these transition dynamics would require a deeper structural model, it seems the aggregate flows of country-level agricultural area are consistent with our estimates of GE extensification.

This definition of the extensive margin is broader than the one proposed by the model, however, absent plot-level time-series data, it is impossible to know from where the extensive margin is recruited. Thus, defining the extensive margin as all hectarage not previously devoted to a specific crop is as precise as we can be given data constraints.

[18] We thank an anonymous referee for pointing this out to us.

Figure 3.3: World hectarage of GE Crops by Technology and Intensive/extensive Margins

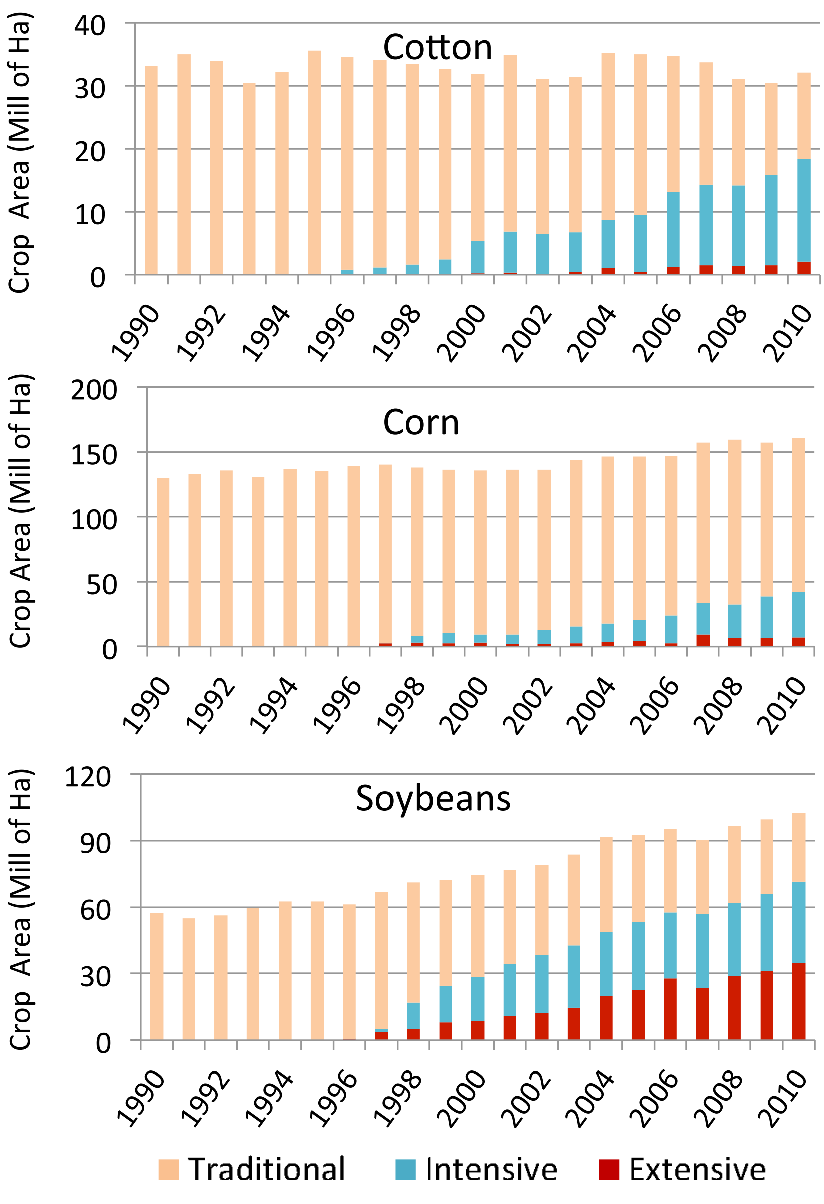

Notes: For each crop cotton, corn, and soybeans, we plot total world area broken down by traditional technology, GE adopted on the intensive margin, and GE adopted on the extensive margin. Area is millions of Ha harvested. Intensive margin indicates lands that switched from traditional technology to GE in the same crop, while extensive margin indicates lands that switched from some other crop or purpose into producing the given crop with GE seeds. GE area is divided between intensive margin and extensive margin by the algorithm described in Section 4.

As discussed above, the model predicts that the extensive margin draws from lands that were previously too low-quality (high pest pressure) to farm profitably. In the data, extensive margin lands could draw from marginal land or any other land that was not previously employed in the given crop. Without plot-level data, it is difficult to determine from which uses extensive margin lands are recruited. Given that most of the extensification occurred in Brazil and Argentina, there is some concern that extensive margin lands result from deforestation, rather than the incorporation of marginal lands. In fact, the stock of forested lands has declined since the introduction of GE seeds, but recent research fails to establish a causal link between soybean expansion and deforestation (Hausman, 2012).

An alternative explanation of the extensive margin owes to particularities of the data collection process. Total harvested area from FAO generally counts physical plots, but the GE data counts plot-seasons, i.e. the number of plot-seasons using GE seeds over the year. Thus, if farmers plant multiple seasons within the year using GE technology, it would look like extensification in the data, though really the physical footprint has not increased. Trigo and Cap (2006) attribute some of the 9.9 million-hectare expansion of soybean area in Argentina to such "double-cropping" activity.[19] It is consistent with our model that GE technology permits double-cropping by extending the range of initial pest pressures accommodated by profitable farm operations. In one example, HT varieties permit control of weeds after the crop plant has emerged from the ground, which speeds up production and allows time for follow-on crops to maturate. A fuller description of land-use change dynamics resulting from GE seeds is the subject of ongoing research.

3.5 Estimated Impacts

What do the intensive and extensive margin mean for supply, prices, land-use, and GHG emissions? While it is beyond the scope of this paper to conduct a full general equilibrium analysis, we can assess magnitudes of these impacts in partial equilibrium in two thought experiments. First, we construct a counterfactual supply curve for each crop assuming that GE technology had not been available in 2010. The impact is the horizontal shift in the supply curve caused by the technology. Then, conditional on assumptions of the elasticity of supply and demand, counterfactual equilibrium quantities and prices are computed. The price impact is the percentage difference between counterfactual price and observed price. Next, to assess the land-use impacts, we calculate how much more land would have been needed to meet observed 2010 demand if GE technology were not available. Finally, we multiply the land-saving figure by a constant GHG per hectare emission rate to compute averted GHG emissions. All impacts are estimated country by country (for each crop) and

[19] Quoting from Trigo and Cap (2006) p. 24: "The second source of benefits has its source in the expansion of the area planted with soybeans, above the trend pre-existing before 1996. This occurred through two mechanisms: the first one was the increase in double-cropping, especially through the combination no-till farming – GE soybeans. This implies that this segment of the area expansion took place without substitution for other crops. The second one is the widening of the 'agricultural frontier' of soybeans towards non-Pampean regions where it substituted for other crops, especially cotton and also ventured into areas considered, until then, 'marginal' for agriculture, where it substituted for livestock production, resulting in an increase in the stock of arable land, induced by a technological innovation."

then aggregated to global figures. We relegate most of the details to the online appendix and just present the results here.

We define the supply effect as the horizontal shift in the supply curve due to GE technology. Only considering the intensive margin yield impact, the supply effect is the percentage difference between observed supply and the counterfactual supply that would have obtained from planting all harvested lands with the lower-yielding traditional technology. However, if one attributes production on the extensive margin to GE technology as well, then the counterfactual supply should subtract production on extensive margin lands as well. We calculate supply effects under these two bounding cases using both our own estimated yield impacts of from section 3, and a range of other yield impacts from the literature.

As shown in Figures C.1 in the online appendix, we find that GE technology increased the supply of corn in 2010 between 5-12% based on our preferred yield effects from column 4 of Table 3.3, depending on how much of the extensive margin is attributed to GE. Thus, even though extensive margin lands represent a small share of total GE corn hectarage, accounting for the extensive margin can potentially make a large difference for the supply effect. We also estimate supply effects based on Sexton and Zilberman (2011) along with all the studies reviewed in Qaim et al. (2009) and find that our supply effects are usually larger, since our estimated yield effects were larger, but the supply effects computed from other yield estimates still generate significant impacts. Estimates range from 2-14% without the extensive margin, and 9-19% with the extensive margin. The notable exceptions are Sexton and Zilberman (2011) and Yorobe and Quicoy (2006), which generates slightly larger supply estimates than ours.

Our estimates imply that GE technology increased the supply of cotton between 15-20% in 2010, depending on the extensive margin. These results are shown in Figure C.2 in the online appendix. Again, these estimates are larger than what would be implied from the yield effects in the Qaim et al. (2009) review. Finally, for soybeans, we find that because the estimated yield effect is small and the estimated extensive margin effect is large, almost all of the supply effect comes from the extensive margin. We estimate that the supply effect was only 2% without the extensive margin, but as large as 40% with the full extensive margin.

The supply effect from GE technology can be translated into price effects using a methodology from De Gorter and Zilberman (1990) and Alston, Norton, Pardey, et al. (1995), where the percentage change in equilibrium price is equal to the supply effect divided by the difference between price elasticity of supply and price elasticity of demand (see online Appendix). In our estimates we assume an elasticity of supply to be 0.3, while we allow elasticity of demand could be either low (-0.3), or high (-0.5).[20] For each elasticity scenario, we also vary the assumption on the extensive margin as before. For each of these 4 scenarios {low elasticity, no extensive margin ; low elasticity with extensive margin; high elasticity, no extensive margin; high elasticity, with extensive margin} price effects are computed conditional on yield estimates from section 3, Sexton and Zilberman (2011), and all the studies reviewed in Qaim et al. (2009).

[20] Roberts and Schlenker (2010) suggest that supply elasticities vary between 0.08 and 0.13 for supply of grain calories and demand elasticities vary between -0.05 and -0.08. Thus, the magnitude of the price effect should be greater than five times the magnitude of the supply effect, which are greater than the impacts estimated here.

In Figure C.3 in the online Appendix, we find that corn prices would have been between 5-19% higher, depending on the assumption of elasticity and extensive margin effect (using our own yield estimates). The price effects in Yorobe and Quicoy (2006) and Sexton and Zilberman (2011) are higher than our estimates, while other studies are roughly 5 percentage points lower. We find that in all cases, the estimates are more sensitive to the inclusion of the extensive margin than the assumption of demand elasticity. In Figure C.4, we find cotton prices would have been 19-33% higher without GE technology. Again, the estimates are higher using our yield impacts instead of others in the literature, but even low yield estimates as in Traxler and Godoy-Avila (2004) and Falck-Zepeda, Traxler, and Nelson (2000) predict that cotton prices would have been 7-19% higher.[21] Finally, for soybeans, the price effect depends heavily on the extensive margin assumption. Without the extensive margin, the price effect is between 3-4%. Including the extensive margin, the price effect is between 50-66%.

Lastly, we estimate land-use saving effects and the corresponding GHG emissions savings due to GE technology as the difference between observed hectarage in 2010 and counterfactual hectarage that would be needed to produce the same output without the GE supply effects. In this thought experiment, the impact of the intensive margin is unambiguous: without the yield boost from GE, more lands would have been recruited to meet observed demand. However, the impact of the extensive margin is not clear. If extensive margins lands come from truelly marginal lands that could not have been used for anything else, then the extensive margin contributes to GHG savings, because without the GE technology, supply from those marginal lands wold have been unavailable and more productive lands would have had to have been converted to agriculture. However, if extensive margin lands come from forest, for example, then their conversion (owing to the GE technology) increases GHG emissions because converting forest to agriculture is highly damaging to the environment. Here, we remain agnostic on the source, and hence GHG impacts of extensive margin lands, and only consider the intensive margin impact. In Table C.1 of the online Appendix, we estimate the land use savings associated with GE cotton are 6 million Ha, or roughly 18% of observed 2010 cotton hectarage. Corn land-use savings equals 5 million Ha, or 3% of observed corn hectarage. Finally, soybean land-use savings are small, at less than 2% of total soybean hectarage.

In the last column of Table C.1, we translate land-use savings into Gt of averted GHG emissions by multiplying the hectares saved by GHG emissions per hectare of land-use change per year. Searchinger et al. (2008) estimates that converting land to cropland generates on average (across the world) 11.7 metric tonnes (t) of GHG per Ha per year.[22] The US EPA estimates a similar figure and other studies have also applied the Searchinger et al. (2008) figure to estimate GHG impacts of various activities (Avery and Avery, 2008; EPA, 2012). Multiplying hectares saved by this GHG impact, we find that across all three crops, GE technology saved 0.15 Gt of GHG emissions in 2010. To put this figure in perspective, the total emissions from all passenger cars in the US in 2010 was roughly 1.28 Gt of GHG[23],

[21] Of course, the estimates from Fitt (2003) yield even lower price effects, but that is because Fitt (2003) estimate no yield impact

[22] The total stock of GHG released is $351t$, amortized over 30 years, for an annual figure of $11.7t$ per year.

[23] The EPA calculates that the average passenger vehicle in the US generates 5.1 metric tons of GHG per year, and the National Transportation Statistics table 1-11 reports there were 250,272,812 passenger

which means the land-use savings effects of GE technology was roughly 1/8th the size of all emissions from automobiles in the US.

3.6 Conclusion

Growing demand for food, feed, fiber and energy means that without new sources of yield gains, new lands must be recruited into production, or else prices must rise to equilibrate the market. Rising prices disproportionally hurt the poor, while clearing lands generates harmful environmental emissions. Agricultural biotechnology can potentially increase yields per hectare, thus boosting supply and preserving lands. In this paper, we generate new estimates of the yield effect that takes account of complementary input use and find larger impacts than most studies in the literature. We also develop a methodology for decomposing observed hectarage into intensive and extensive margin. While we cannot say if GE technology has *caused* the increase in the range of lands that can profitably be farmed, we have found that hectarages have increased since the introduction of GE technology, and counterfactual supply scenarios suggest that the extensive margin effect could make a large difference in computing supply, price, and land-use saving effects. Future research using experimental variation to identify the causal link between GE adoption and the extensive margin would constitute a significant contribution.

We find that adoption of GE has significant impact on the price of cotton, corn, and soybeans. As corn and soybeans are used extensively in the production of food, these price effects likely translate into lower food prices, benefiting the poor (Hochman et al., 2011). The analysis suggests that while high adoption rates of GE cotton and soybean has contributed to a significant price reduction in these commodities, bans and other regulations limited the adoption of GE corn to less than 30% of total corn hectarage, reducing its total price effect. If adoption of corn is expanded globally, we expect much larger increases in supply both because of reduction in pest damage as well the complementary input effect, resulting in further corn price reductions. The use of GE is practically banned everywhere for major food grains like wheat and rice, even though existing traits could reduce pest damage in these two crops. Our analysis suggests that developing new GE varieties in these crops has the potential to reduce their prices as well as the environmental side effects from producing these crops.

Finally, we find that GE has had significant environmental benefits, even considering just the intensive margin. We estimate that GE technology slowed land-use change and prevented GHG emissions on the order of 1/8 the annual GHG emissions caused by driving in the US. As the poor are expected to suffer the most from climate change, these environmental gains also mean distributional gains for the poor.

vehicles in the US in 2010, which implies that total GHG emissions from passenger vehicles equaled $5.1 * 250,272,812 * \frac{1}{1,000,000,000} = 1.28$Gt.

Bibliography

Alston, Julian M, George W Norton, Philip G Pardey, et al. (1995). *Science under scarcity: principles and practice for agricultural research evaluation and priority setting.* Cornell University Press.

Amiti, Mary and Amit K Khandelwal (2013). “Import competition and quality upgrading”. In: *Review of Economics and Statistics* 95.2, pp. 476–490.

Anderson, JE and E van Wincoop (2003). “Gravity with Gravitas: A Solution to the Border Puzzle”. In: *The American Economic Review* 93.1, pp. 170–192.

Antweiler, Werner, Brian R Copeland, and M Scott Taylor (2001). “Is Free Trade Good for the Environment?” In: *The American Economic Review* 91.4, pp. 877–908.

Avery, Alex and Dennis Avery (2008). “Beef production and greenhouse gas emissions”. In: *Environmental health perspectives* 116.9, A374.

Baldwin, Richard and James Harrigan (2011). “Zeros, Quality, and Space: Trade Theory and Trade Evidence”. In: *American Economic Journal: Microeconomics*, pp. 60–88.

Barrows, Geoffrey and Helene Ollivier (2014). “Product mix, trade, and the environment: theory and evidence from Indian manufacturing”. In: *mimeo*.

Bastos, Paulo and Joana Silva (2010). “The quality of a firm’s exports: Where you export to matters”. In: *Journal of International Economics* 82.2, pp. 99–111.

Bastos, Paulo, Joana Silva, and Eric Verhoogen (2014). *Export Destinations and Input Prices.* Tech. rep. National Bureau of Economic Research.

Batrakova, Svetlana and Ronald B Davies (2012). “Is there an environmental benefit to being an exporter? Evidence from firm-level data”. In: *Review of World Economics* 148.3, pp. 449–474.

Benbrook, Charles (1999). “Evidence of the magnitude and consequences of the Roundup Ready soybean yield drag from university-based varietal trials in 1998”. In: *AgBiotech Infonet technical paper* 1.

Benjamin, Falck-Zepeda, Idah Sithole-Niang, et al. (2013). “Genetically modified crops in Africa: Economic and policy lessons from countries south of the Sahara”. In:

Bennett, Alan B et al. (2013). “Agricultural biotechnology: economics, environment, ethics, and the future”. In: *Annual Review of Environment and Resources* 38.1.

Bernard, A., S. Redding, and P. Schott (2013). “Forthcoming,“Multi-Product Firms and Product Switching,”” in: *American Economic Review*.

Bernard, A.B., S.J. Redding, and P.K. Schott (2011). “Multiproduct firms and trade liberalization”. In: *The Quarterly Journal of Economics* 126.3, pp. 1271–1318.

Bernard, Andrew B and J Bradford Jensen (1999). “Exceptional exporter performance: cause, effect, or both?” In: *Journal of international economics* 47.1, pp. 1–25.

Bernard, Andrew B, Andreas Moxnes, and Karen Helene Ulltveit-Moe (2014). *Two-Sided Heterogeneity and Trade*. Tech. rep. National Bureau of Economic Research.

Bloom, N., M. Draca, and J. Van Reenen (2011). *Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity*. Tech. rep. National Bureau of Economic Research.

Brambilla, I., A.K. Khandelwal, and P.K. Schott (2010). "China's experience under the multi-fiber arrangement (mfa) and the agreement on textiles and clothing (atc)". In: *China's Growing Role in World Trade*. University of Chicago Press, pp. 345–387.

Brambilla, Irene, Daniel Lederman, and Guido Porto (2012). "Exports, Export Destinations, and Skills". In: *American Economic Review* 102.7, pp. 3406–3438.

Broda, Christian and David E Weinstein (2006). "Globalization and the Gains from Variety". In: *The Quarterly Journal of Economics* 121.2, pp. 541–585.

Bustos, P. (2011). "Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms". In: *The American economic review* 101.1, pp. 304–340.

Carpenter, Janet E (2010). "Peer-reviewed surveys indicate positive impact of commercialized GM crops". In: *Nature biotechnology* 28.4, pp. 319–321.

Chakraborty, Pavel (2012). "The Great Trade Collapse and Indian Firms". In: *PhD Student, Graduate Institute of International and Development Studies, Geneva*.

Cherniwchan, JEVAN (2013). *Trade liberalization and the environment: Evidence from nafta and us manufacturing*. Tech. rep. mimeo.

Cicala, Steve (2015). "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation". In: *American Economic Review* 105.1, pp. 411–44.

Copeland, B.A. and M.S. Taylor (1994). "North-South trade and the environment". In: *Quarterly Journal of Economics* 109.3, pp. 755–787.

Copeland, Brian R and M Scott Taylor (2004). "Trade, Growth, and the Environment". In: *Journal of Economic Literature* 42, pp. 7–71.

Copeland, Brian Richard and Michael Scott Taylor (2003). *Trade and the Environment*. Princeton University Press Princeton.

Crost, Benjamin et al. (2007). "Bias from Farmer Self-Selection in Genetically Modified Crop Productivity Estimates: Evidence from Indian Data". In: *Journal of Agricultural Economics* 58.1, pp. 24–36.

Cui, Jingbo, Harvey Lapan, and GianCarlo Moschini (2012). *Are Exporters More Environmentally Friendly than Non-Exporters? Theory and Evidence*. Tech. rep.

David, Paul A (1969). *A Contribution to the Theory of Diffusion*. Research Center in Economic Growth, Stanford University.

De Gorter, Harry and David Zilberman (1990). "On the political economy of public good inputs in agriculture". In: *American Journal of Agricultural Economics* 72.1, pp. 131–137.

De Loecker, Jan (2011). "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity". In: *Econometrica* 79.5, pp. 1407–1451.

De Loecker, Jan et al. (2012). *Prices, markups and trade reform*. Tech. rep. National Bureau of Economic Research.

Debroy, B and A Santhanam (1993). "Matching Trade Codes with Industrial Codes." In: *Foreign Trade Bulletin* 24.1.

Duflo, Esther et al. (2013). “Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India*”. In: *The Quarterly Journal of Economics*, qjt024.

Durceylan, Esra (2009). “Managerial Effort, Agency, and Industrial Evolution”. In:

Eckel, Carsten and J Peter Neary (2010). “Multi-product firms and flexible manufacturing in the global economy”. In: *The Review of Economic Studies* 77.1, pp. 188–217.

EPA (2012). ““Inventory of U.S. Greenhouse Gas Emissions and Sinks, 1990-2010.””. In: EPA 430-R-12-001.

Faber, Benjamin (2012). “Trade liberalization, the price of quality, and inequality: Evidence from mexican store prices”. In: *London School of Economics*.

Falck-Zepeda, Jose Benjamin, Greg Traxler, and Robert G Nelson (2000). “Surplus distribution from the introduction of a biotechnology innovation”. In: *American Journal of Agricultural Economics* 82.2, pp. 360–369.

Fargione, Joseph et al. (2008). “Land clearing and the biofuel carbon debt”. In: *Science* 319.5867, pp. 1235–1238.

Feder, Gershon, Richard E Just, and David Zilberman (1985). “Adoption of agricultural innovations in developing countries: A survey”. In: *Economic development and cultural change* 33.2, pp. 255–298.

Feenstra, Robert C, John Romalis, and Peter K Schott (2002). *US imports, exports, and tariff data, 1989-2001*. Tech. rep. National Bureau of Economic Research.

Felipe, Jesus and Franklin M Fisher (2003). “Aggregation in production functions: what applied economists should know”. In: *Metroeconomica* 54.2-3, pp. 208–262.

Finger, Robert et al. (2011). “A meta analysis on farm-level costs and benefits of GM crops”. In: *Sustainability* 3.5, pp. 743–762.

Fitt, Gary P (2003). “Implementation and impact of transgenic Bt cottons in Australia”. In: *Proceedings of the World Cotton Research Conference*. Vol. 3, pp. 371–382.

Forslid, Rikard, Toshihiro Okubo, and Karen Helene Ulltveit-Moe (2011). “Why are firms that export cleaner? International trade and CO2 emissions”. In:

Frankel, Jeffrey A and Andrew K Rose (2005). “Is trade good or bad for the environment? Sorting out the causality”. In: *Review of Economics and Statistics* 87.1, pp. 85–91.

Galdeano-Gómez, Emilio (2010). “Exporting and Environmental Performance: A Firm-level Productivity Analysis”. In: *The World Economy* 33.1, pp. 60–88.

Goldberg, Pinelopi K et al. (2010). “Multiproduct firms and product turnover in the developing world: Evidence from india”. In: *The Review of Economics and Statistics* 92.4, pp. 1042–1049.

Greenstone, Michael and Rema Hanna (2014). “Environmental Regulations, Air and Water Pollution, and Infant Mortality in India”. In: *American Economic Review* 104.10, pp. 3038–72.

Greenstone, Michael and B Kelsey Jack (2013). *Envirodevonomics: A research agenda for a young field*. Tech. rep. National Bureau of Economic Research.

Gutiérrez, E. and K. Teshima (2011). “Import Competition and Environmental Performance: Evidence from Mexico”. In:

Harrigan, J. and G. Barrows (2009). “Testing the theory of trade policy: Evidence from the Abrupt End of the Multifiber Arrangement”. In: *The Review of Economics and Statistics* 91.2, pp. 282–294.

Harrigan, James, Xiangjun Ma, and Victor Shlychkov (2011). *Export prices of US firms*. Tech. rep. National Bureau of Economic Research.

Hausman, Catherine (2012). "Biofuels and Land Use Change: Sugarcane and Soybean Acreage Response in Brazil". In: *Environmental and Resource Economics* 51.2, pp. 163–187.

Hochman, Gal et al. (2011). "The role of inventory adjustments in quantifying factors causing food price inflation". In: *World Bank Policy Research Working Paper Series, Vol*.

Holladay, J Scott (2010). "Are Exporters Mother Nature's Best Friends?" In:

Huffman, Wallace E and Robert E Evenson (1992). "Contributions of public and private science and technology to US agricultural productivity". In: *American Journal of Agricultural Economics* 74.3, pp. 751–756.

Hummels, David et al. (2014). "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data". In: *American Economic Review* 104.6, pp. 1597–1629.

Iacovone, Leonardo and Beata S Javorcik (2010). "Multi-Product Exporters: Product Churning, Uncertainty and Export Discoveries*". In: *The Economic Journal* 120.544, pp. 481–499.

Imbens, Guido and Jeffrey M Wooldridge (2007). *What's new in econometrics?* NBER.

Jing Cao, Larry D. Qiu and Mohan Zhou (2013). "The Inverted-U Shaped Investment in Abatement The Inverted-U Shaped Investment in Abatement Technology: Theory and Evidence". In:

Just, Richard E, Julian M Alston, and David Zilberman (2006). *Regulating agricultural biotechnology [electronic resource]: economics and policy*. Vol. 30. Springer.

Just, Richard E and Rulon D Pope (1978). "Stochastic specification of production functions and economic implications". In: *Journal of Econometrics* 7.1, pp. 67–86.

Just, Richard E and David Zilberman (1988). "The effects of agricultural development policies on income distribution and technological change in agriculture". In: *Journal of Development Economics* 28.2, pp. 193–216.

Just, Richard E et al. (1990). "Input allocation in multicrop systems". In: *American Journal of Agricultural Economics* 72.1, pp. 200–209.

Kellenberg, Derek K (2009). "An empirical investigation of the pollution haven effect with strategic environment and trade policy". In: *Journal of International Economics* 78.2, pp. 242–255.

Khandelwal, Amit (2010). "The long and short (of) quality ladders". In: *The Review of Economic Studies* 77.4, pp. 1450–1476.

Khandelwal, Amit K, Peter K Schott, and Shang-Jin Wei (2013). "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters". In: *American Economic Review* 103.6, pp. 2169–95.

Kreickemeier, Udo and Philipp M Richter (2014). "Trade and the environment: The role of firm heterogeneity". In: *Review of International Economics* 22.2, pp. 209–225.

Krishna, Kala (1989). "Trade restrictions as facilitating practices". In: *Journal of International Economics* 26.3, pp. 251–270.

Kugler, Maurice and Eric Verhoogen (2012). "Prices, plant size, and product quality". In: *The Review of Economic Studies* 79.1, pp. 307–339.

Levinson, Arik (2009). "Technology, International Trade, and Pollution from US Manufacturing". In: *The American economic review* 99.5, pp. 2177–2192.

Lichtenberg, Erik and David Zilberman (1986). "The econometrics of damage control: why specification matters". In: *American Journal of Agricultural Economics* 68.2, pp. 261–273.

Lileeva, Alla and Daniel Trefler (2010). "Improved access to foreign markets raises plant-level productivity... for some plants". In: *The Quarterly Journal of Economics* 125.3, pp. 1051–1099.

Lipscomb, Molly (2008). "The Effect of Environmental Enforcement on Product Choice and Competition: Theory and evidence from India". In: *Center for Economic Analysis Department of Economics, University of Colorado at Boulder, Working Paper No. 08* 13.

Liu, Elaine M (2013). "Time to change what to sow: Risk preferences and technology adoption decisions of cotton farmers in China". In: *Review of Economics and Statistics* 95.4, pp. 1386–1403.

Lobell, David B, Wolfram Schlenker, and Justin Costa-Roberts (2011). "Climate trends and global crop production since 1980". In: *Science* 333.6042, pp. 616–620.

Manova, Kalina and Zhiwei Zhang (2012). "Export Prices Across Firms and Destinations*." In: *Quarterly Journal of Economics* 127.1.

Martin, L.A. (2012). "Energy efficiency gains from trade: greenhouse gas emissions and india's manufacturing firms.'" In: *Department of Agricultural and Resource Economics, University of California Berkley*.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I. P. Ottaviano (2014). "Market Size, Competition, and the Product Mix of Exporters". In: *American Economic Review* 104.2, pp. 495–536.

Melitz, M. (2003). "The impact of trade on aggregate industry productivity and intra-industry reallocations". In: *Econometrica* 71.6, pp. 1695–1725.

Melitz, M.J. and G.I.P. Ottaviano (2008). "Market size, trade, and productivity". In: *The review of economic studies* 75.1, pp. 295–316.

Park, Albert et al. (2010). "Exporting and firm performance: Chinese exporters and the Asian financial crisis". In: *The Review of Economics and Statistics* 92.4, pp. 822–842.

Perla, Jesse, Christopher Tonetti, and Michael E Waugh (2012). "Equilibrium Technology Diffusion, Trade, and Growth". In:

Qaim, M et al. (2009). "The economics of genetically modified crops." In: *Annual Review of Resource Economics* 1, pp. 665–694.

Qaim, Matin and David Zilberman (2003). "Yield effects of genetically modified crops in developing countries". In: *Science* 299.5608, pp. 900–902.

Rajagopal, Deepak et al. (2007). "Challenge of biofuel: filling the tank without emptying the stomach?" In: *Environmental Research Letters* 2.4, p. 044004.

Roberts, Michael J and Wolfram Schlenker (2010). *Is agricultural production becoming more or less sensitive to extreme heat? Evidence from US corn and soybean yields*. Tech. rep. National Bureau of Economic Research.

Rotunno, Lorenzo, Pierre-Louis Vézina, and Zheng Wang (2013). "The rise and fall of (Chinese) African apparel exports". In: *Journal of Development Economics* 105, pp. 152–163.

Schelling, Thomas C (1992). "Some economics of global warming". In: *The American Economic Review* 82.1, pp. 1–14.

Searchinger, T. et al. (2008). "Use of US croplands for biofuels increases greenhouse gases through emissions from land-use change". In: *Science* 319.5867, pp. 1238–1240.

Sexton, Steven and David Zilberman (2011). *How agricultural biotechnology boosts food supply and accomodates biofuels*. Tech. rep. National Bureau of Economic Research.

Sexton, Steven E. "Automatic Bill Payment, Price Salience, and Consumption: Evidence from Residential Electricity Consumption". In: *Review of Economics and Statistics, forthcoming*.

Topalova, Petia and Amit Khandelwal (2011). "Trade liberalization and firm productivity: The case of india". In: *Review of economics and statistics* 93.3, pp. 995–1009.

Traxler, Greg and Salvador Godoy-Avila (2004). "Transgenic Cotton in Mexico". In: *AgBio-Forum* 7.1&2, pp. 57–62.

Trefler, Daniel (1993). "Trade liberalization and the theory of endogenous protection: an econometric study of US import policy". In: *Journal of Political Economy*, pp. 138–160.

Trigo, EJ and EJ Cap (2006). "Ten years of genetically modified crops in Argentine agriculture". In: *Argentine Council for Information and Development of Biotechnology (Argen-Bio), Buenos Aires, Argentina*.

Verhoogen, Eric A (2008). "Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector". In: *The Quarterly Journal of Economics* 123.2, pp. 489–530.

Yorobe, JM and Cesar B Quicoy (2006). "Economic impact of Bt corn in the Philippines". In: *Philippine Agricultural Scientist* 89.3, p. 258.

Appendix A

Appendix to Chapter 1

A.1 Theory

Proof of Selection into Exporting

Consider for simplicity a world with two countries: l and h. As in Melitz and Ottaviano (2008), the price distributions in country l of domestic firms producing in l, $p_{ll}(\varphi, m)$, and exporters producing in h, $p_{hl}(\varphi, m)$, are identical. Thus, as in the closed economy, the threshold price condition in country l (1.4), along with the resulting Pareto distribution of all prices for varieties sold in l, yield a zero-cutoff profit condition linking the domestic cost cutoff to the mass of varieties consumed in country l:

$$M_l = \frac{2\gamma(k+1)(\alpha - \Phi_{ll})}{\eta\Phi_{ll}}. \tag{A.1}$$

Given a positive mass of entrants $N_{E,l}$ in country l, there will be $N_{E,l}[1 - G(\varphi_{ll})]$ firms producing in country l, and $N_{E,l}[1-G(\varphi_{lh})]$ firms exporting $\rho_{lh}\Phi_{ll}^k\Omega_l N_{E,l}$ varieties to country h, where $\rho_{lh} \equiv \theta_{lh}^{-k} < 1$ is a measure of 'freeness' of trade from country l to country h. Summing over all varieties from countries l and h sold in country l, we get

$$M_l = N_{E,l}\Phi_{ll}^k\Omega_l + \rho_{hl}N_{E,h}\Phi_{hh}^k\Omega_h. \tag{A.2}$$

Combining the two expressions for M_l, and similarly for M_h, gives the number of entrants in each country (by symmetry):

$$N_{E,l} = \frac{2\gamma(k+1)}{\eta(1-\rho_{lh}\rho_{hl})\Omega_l}\left[\frac{\alpha - \Phi_{ll}}{\Phi_{ll}^{k+1}} - \rho_{hl}\frac{\alpha - \Phi_{hh}}{\Phi_{hh}^{k+1}}\right]. \tag{A.3}$$

Assuming a non-specialized equilibrium where both countries produce the differentiated good ($N_{E,l} > 0$) implies that only a subset of relatively more productive firms choose to export in either country, since $N_{E,l} > 0$ is equivalent to

$$\frac{\alpha - \Phi_{ll}}{\Phi_{ll}^{k+1}} > \rho_{hl}\frac{\alpha - \Phi_{hh}}{\Phi_{hh}^{k+1}} \Leftrightarrow \frac{\alpha/\theta_{hl} - \Phi_{hl}}{\alpha - \Phi_{hh}}\left(\frac{\Phi_{hh}}{\Phi_{hl}}\right)^{k+1} > 1, \tag{A.4}$$

which is incompatible with $\Phi_{hl} \geq \Phi_{hh}$. Therefore, $\Phi_{hl} < \Phi_{hh}$.

Predictions for the Product-Mix and Price Effects

The comparative statics analysis focuses on foreign demand shocks to match our empirical exercise. We thus consider a foreign demand shock such that $dL_h > 0$. This shock makes market h less competitive for its domestic firms: given (1.17),

$$\frac{d\Phi_{hh}}{dL_h} = -\frac{\Phi_{hh}}{(k+2)L_h} < 0. \tag{A.5}$$

The demand shock also impacts country l's firms through the export cost cutoff, which is $\Phi_{lh} = \Phi_{hh}/\theta_{lh}$. We derive the endogenous adjustments made by country l's firms in this section.

Proof of Prediction 1

Let $R_{lh}(\varphi)$ denote the export revenue of firm φ located in country l exporting to country h. We have

$$R_{lh}(\varphi) = \sum_{m=0}^{M(\varphi)-1} r_{lh}(\varphi, m) \tag{A.6}$$

$$= \sum_{m=0}^{M(\varphi)-1} \frac{L_h\theta_{lh}^2}{4\gamma}[\Phi_{lh}^2 - \Phi(\varphi, m)^2]. \tag{A.7}$$

For a fixed product scope M with $1 < M \leq M(\varphi)$, this can be written as

$$R_{lh}(\varphi) = \frac{L_h\theta_{lh}^2}{4\gamma} M\Phi_{lh}^2 - \frac{L_h\theta_{lh}^2}{4\gamma} \sum_{m=0}^{M-1} \Phi(\varphi, m)^2, \tag{A.8}$$

subject to φ being in the range of TFP that allows firms to produce optimally M products. The impact of a foreign demand shock is such that

$$\frac{dR_{lh}(\varphi)}{dL_h} = \frac{Mk\Phi_{hh}^2}{4\gamma(k+2)} - \frac{\theta_{lh}^2}{4\gamma} \sum_{m=0}^{M-1} \Phi(\varphi, m)^2. \tag{A.9}$$

Because $M\Phi_{lh}^2 \geq \sum_{m=0}^{M-1} \Phi(\varphi, m)^2$, we find that the most efficient firms (with higher φ) are affected positively ($dR_{lh}(\varphi)/dL_h > 0$) whereas the less efficient firms are affected negatively ($dR_{lh}(\varphi)/dL_h < 0$) by the demand shock. This can be explained by the fact that only the most profitable products see their revenue increase whereas the less profitable products either see their revenue decrease or are dropped by the firm (see prediction 2). Allowing for a variable product scope implies that products with high marginal costs are no longer sold by the firm on market h. Only the less profitable products are dropped; and by continuity, it should not modify the result.

Proof of Prediction 2

Using (1.12), we get

$$\frac{dr_{lh}(\varphi,m)}{dL_h} = \frac{\theta_{lh}^2}{4\gamma}\left[\frac{k}{k+2}\Phi_{lh}^2 - \Phi(\varphi,m)^2\right], \tag{A.10}$$

which is positive for firm-product cost $\Phi(\varphi,m) \leq \sqrt{k/(k+2)}\Phi_{lh}$, and strictly negative for firm-product cost $\Phi(\varphi,m) > \sqrt{k/(k+2)}\Phi_{lh}$.

Proof of Prediction 3

Let $\overline{EQ}_{lh}(\varphi)$ denote the average emission intensity in quantity for a firm φ producing $M(\varphi)$ varieties in country l that are exported to country h. We have

$$\overline{EQ}_{lh}(\varphi) = \frac{\sum_{m=0}^{M(\varphi)-1} EQ(\varphi,m)q_{lh}(\Phi(\varphi,m))}{\sum_{m=0}^{M(\varphi)-1} q_{lh}(\Phi(\varphi,m))}. \tag{A.11}$$

For a fixed product scope M with $1 < M \leq M(\varphi)$, this can be written as

$$\overline{EQ}_{lh}(\varphi) = \frac{\Phi_{lh}\sum_{m=0}^{M-1} EQ(\varphi,m) - \sum_{m=0}^{M-1} EQ(\varphi,m)\Phi(\varphi,m)}{M\Phi_{lh} - \sum_{m=0}^{M-1}\Phi(\varphi,m)}, \tag{A.12}$$

subject to φ being in the range of TFP that allows firms to produce optimally M products. The only impact that a foreign demand shock has on $\overline{EQ}_{lh}(\varphi)$ comes from its impact on the export cost cutoff $\Phi_{lh} = \Phi_{hh}/\theta_{lh}$ described in (A.5). Thus,

$$\frac{d\overline{EQ}_{lh}(\varphi)}{d\Phi_{lh}} = \frac{M\sum_{m=0}^{M-1} EQ(\varphi,m)\Phi(\varphi,m) - \sum_{m=0}^{M-1} EQ(\varphi,m)\sum_{m=0}^{M-1}\Phi(\varphi,m)}{\left[M\Phi_{lh} - \sum_{m=0}^{M-1}\Phi(\varphi,m)\right]^2} \tag{A.13}$$

For all $M \in N^*$, denote the numerator as

$$D_{QM} \equiv M\sum_{m=0}^{M-1} EQ(\varphi,m)\Phi(\varphi,m) - \sum_{m=0}^{M-1} EQ(\varphi,m)\sum_{m=0}^{M-1}\Phi(\varphi,m), \tag{A.14}$$

where $(\Phi(\varphi,m))_{m\in N}$ and $(EQ(\varphi,m))_{m\in N}$ are real positive sequences. First, consider some examples such as $M = 2$. Simplifying notations using EQ_m and Φ_m, we have

$$D_{Q2} = 2(EQ_0\Phi_0 + EQ_1\Phi_1) - (EQ_0 + EQ_1)(\Phi_0 + \Phi_1) = (EQ_1 - EQ_0)(\Phi_1 - \Phi_0). \tag{A.15}$$

Now for $M = 3$,

$$D_{Q3} = 3(EQ_0\Phi_0 + EQ_1\Phi_1 + EQ_2\Phi_2) - (EQ_0 + EQ_1 + EQ_2)(\Phi_0 + \Phi_1 + \Phi_2) \tag{A.16}$$

$$= (EQ_1 - EQ_0)(\Phi_1 - \Phi_0) + (EQ_2 - EQ_0)(\Phi_2 - \Phi_0) + (EQ_2 - EQ_1)(\Phi_2 - \Phi_1). \tag{A.17}$$

Using a recursive argument, we thus deduce that

$$D_{QM} = \sum_{0 \leq i < j \leq M-1} (EQ_j - EQ_i)(\Phi_j - \Phi_i) \tag{A.18}$$

By assumption $(\Phi_m)_{m \in N}$ is strictly increasing, thus for all $0 \leq i < j \leq M-1$ we have $\Phi_j > \Phi_i$. Thus D_{QM} is positive, hence $\overline{EQ}_{lh}(\varphi)$ increases in Φ_{lh} if and only if $(EQ_m)_{m \in N}$ is also strictly increasing whereas it decreases in Φ_{lh} if and only if $(EQ_m)_{m \in N}$ is strictly decreasing.

Even when product scope M drops due to the decrease in Φ_{lh}, the average emission intensity must still have the same variation as $EQ(\varphi, m)$ due to the continuity of $\overline{EQ}_{lh}(\varphi)$ with respect to Φ_{lh} (both total emissions and $Q(\varphi)$ are continuous in Φ_{lh} as the firm produces zero units of a variety right before it is dropped when competition gets tougher).

Product-Mix Effects

Let $\overline{EQ}(\varphi)$ denote the firm emission intensity in quantity for a firm φ producing $M_{ll}(\varphi)$ varieties in country l, and exporting $M_{lh}(\varphi)$ varieties to each country h. We first get

$$\overline{EQ}(\varphi) = \frac{\sum_{h=1}^{H} \sum_{m=0}^{M_{lh}(\varphi)-1} EQ(\varphi, m) q_{lh}(\Phi(\varphi, m))}{\sum_{h=1}^{H} \sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\Phi(\varphi, m))}. \tag{A.19}$$

Suppose, for simplicity, that $H = 2$: a firm located in country l can either sell its products to consumers from country l or to consumers from country h. The aggregate firm emission intensity in quantity is

$$\overline{EQ}(\varphi) = [1 - x_{lh}(\varphi)]\overline{EQ}_{ll}(\varphi) + x_{lh}(\varphi)\overline{EQ}_{lh}(\varphi), \tag{A.20}$$

where

$$x_{lh}(\varphi) \equiv \frac{\sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\varphi, m)}{\sum_{m=0}^{M_{ll}(\varphi)-1} q_{ll}(\varphi, m) + \sum_{m=0}^{M_{lh}(\varphi)-1} q_{lh}(\varphi, m)}. \tag{A.21}$$

$x_{lh}(\varphi)$ denotes the share of exports to country h in total outputs. From Prediction 3, we know the impact of an import demand shock on $\overline{EQ}_{lh}(\varphi)$, whereas $\overline{EQ}_{ll}(\varphi)$ remains unaffected. For a fixed domestic product scope M_{ll} with $1 < M_{ll} \leq M_{ll}(\varphi)$, and a fixed exported product scope M_{lh} with $1 < M_{lh} \leq M_{lh}(\varphi)$, $x_{lh}(\varphi)$ can be written as

$$x_{lh}(\varphi) = \frac{M_{lh} L_h \theta_{lh} \Phi_{lh} - L_h \theta_{lh} \sum_{m=0}^{M_{lh}-1} \Phi(\varphi, m)}{M_{ll} L_l \Phi_{ll} - L_l \sum_{m=0}^{M_{ll}-1} \Phi(\varphi, m) + M_{lh} L_h \theta_{lh} \Phi_{lh} - L_h \theta_{lh} \sum_{m=0}^{M_{lh}-1} \Phi(\varphi, m)}$$

subject to φ being in the range of TFP that allows firms to produce optimally M_{ll} products. Given the equilibrium export cost cutoff, we have

$$\frac{dL_h \Phi_{lh}}{dL_h} = \Phi_{lh} \frac{k+1}{k+2} > 0. \tag{A.22}$$

Thus,

$$\frac{dx_{lh}(\varphi)}{dL_h} = \frac{\left[M_{lh}\theta_{lh}\Phi_{lh}\frac{k+1}{k+2} - \theta_{lh}\sum_{m=0}^{M_{lh}-1}\Phi(\varphi,m)\right]\left[M_{ll}L_l\Phi_{ll} - L_l\sum_{m=0}^{M_{ll}-1}\Phi(\varphi,m)\right]}{\left[M_{ll}L_l\Phi_{ll} - L_l\sum_{m=0}^{M_{ll}-1}\Phi(\varphi,m) + M_{lh}L_h\theta_{lh}\Phi_{lh} - L_h\theta_{lh}\sum_{m=0}^{M_{lh}-1}\Phi(\varphi,m)\right]^2},$$

whose sign only depends on $M_{lh}\Phi_{lh}\frac{k+1}{k+2} - \sum_{m=0}^{M_{lh}-1}\Phi(\varphi,m)$. If $\frac{k+1}{k+2}$ were close to 1 (k tend to infinity), then $dx_{lh}(\varphi)/dL_h$ would be positive. In general, $dx_{lh}(\varphi)/dL_h$ is positive for efficient firms, but negative for less efficient firms. This result can be derived from Prediction 1. Changing the product scope of exported goods would not modify this result.

We further need to compare $\overline{EQ}_{lh}(\varphi)$ with $\overline{EQ}_{ll}(\varphi)$ to assess the impact of a demand shock on $\overline{EQ}(\varphi)$. Selection into exporting implies that $\Phi_{lh} < \Phi_{ll}$: only the most efficient firms export and only the most profitable products are exported. Therefore, $M_{lh}(\varphi) < M_{ll}(\varphi)$: the export basket is skewed toward the core products even more so than the domestic basket. The relative emission intensity of these baskets thus depends on whether the core products are cleaner or dirtier than higher-m products. Hence,

i/ If $EQ(\varphi,m)$ is increasing in m (i.e., core products are cleaner), then $\overline{EQ}_{lh}(\varphi) < \overline{EQ}_{ll}(\varphi)$ and $\overline{EQ}_{lh}(\varphi)$ is decreasing in L_h. Thus, the average firm emission intensity $\overline{EQ}(\varphi)$ of efficient firms decreases with a demand shock whereas the impact is ambiguous for less efficient firms.

ii/ If $EQ(\varphi,m)$ is decreasing in m (i.e., core products are dirtier), then $\overline{EQ}_{lh}(\varphi) > \overline{EQ}_{ll}(\varphi)$ and $\overline{EQ}_{lh}(\varphi)$ is increasing in L_h. Thus, the average firm emission intensity $\overline{EQ}(\varphi)$ of efficient firms increases with a demand shock whereas the impact is ambiguous for less efficient firms.

Generalizing this result requires to compare the average emission intensity of the basket of goods exported to country h where the demand shock occurs with the average emission intensity of the other baskets. If $\overline{EQ}_{lh}(\varphi)$ remains lower (higher) than the average emission intensity of all other baskets of products for firm φ, even after the demand shock modifies it, then $\overline{EQ}(\varphi)$ of efficient firms decreases (increases) with the demand shock. The impact of the demand shock on less efficient firms is ambiguous.

Proof of Prediction 4

Consider first a mono-product firm. The average emission intensity in value for a firm-product with unit cost $\Phi(\varphi,m)$ is

$$\overline{EV}(\varphi,m) = \sum_{h=1}^{H}\left(\frac{r_{lh}(\varphi,m)}{\sum_{h=1}^{H} r_{lh}(\varphi,m)}\right)EV_{lh}(\varphi,m), \tag{A.23}$$

where $EV_{lh}(\varphi,m)$ is the emission intensity in value in market h as defined by (1.8), and $r_{lh}(\varphi,m)/\sum_{h=1}^{H} r_{lh}(\varphi,m)$ corresponds to the share of revenues made in country h. From Prediction 2, we know that a demand shock in country h would raise $r_{lh}(\varphi,m)$ for the most profitable products, and decrease $r_{lh}(\varphi,m)$ for more expensive products, while leaving

$r_{lj}(\varphi, m)$, $j \neq h$ unaffected. A positive shock to L_h also lowers $p_{lh}(\varphi, m)$, which increases $EV_{lh}(\varphi, m)$.

We observe that $EV_{lh}(\varphi, m) < EV_{ll}(\varphi, m)$ is equivalent to $p_{lh}(\varphi, m) > p_{ll}(\varphi, m)$. Using (1.10), it is also equivalent to $\Phi(\varphi, m)\,(\theta_{lh} - 1) > \Phi_{ll} - \Phi_{hh}$. Hence, if country h's market is less competitive than country l's market ($\Phi_{ll} < \Phi_{hh}$), then $EV_{lh}(\varphi, m) < EV_{ll}(\varphi, m)$. Given

$$\Omega_h = \sum_{m=0}^{\infty} e^{-\sigma mk} \left[1 + \tau_h^{\frac{\epsilon}{\epsilon-1}} e^{\frac{m\epsilon(\nu-\sigma)}{\epsilon-1}}\right]^{\frac{(1-\epsilon)k}{\epsilon}} \tag{A.24}$$

we have

$$\frac{d\Omega_h}{d\tau_h} = -k\tau_h^{\frac{1}{\epsilon-1}} \sum_{m=0}^{\infty} e^{\frac{m[(\nu-\sigma)\epsilon+\sigma k(1-\epsilon)]}{(\epsilon-1)}} \left[1 + \tau_h^{\frac{\epsilon}{\epsilon-1}} e^{\frac{m\epsilon(\nu-\sigma)}{\epsilon-1}}\right]^{\frac{k(1-\epsilon)-\epsilon}{\epsilon}} < 0. \tag{A.25}$$

This implies, given (1.17),

$$\frac{d\Phi_{hh}}{d\tau_h} > 0. \tag{A.26}$$

Thus, countries with stringent environmental regulations are less competitive. If firms export to a country with higher regulations, the export price will be higher, and $EV_{lh}(\varphi, m)$ will be lower than $EV_{ll}(\varphi, m)$.

In a multiple destination setting, we must compare $EV_{lh}(\varphi, m)$ with the average emission intensity in value across all other destinations. If the export price to country h is higher (lower) than the average price over other destinations, then $EV_{lh}(\varphi, m)$ is lower (higher) than the average emission intensity in value across other destinations. If $EV_{lh}(\varphi, m)$ is lower, an increase in exporting (for profitable products) would reduce $\overline{EV}(\varphi, m)$ as long as the compositional shift (toward country h's import basket) outweighs the export price decrease. For less profitable products, there is no ambiguity; hence, a demand shock from country h increases $\overline{EV}(\varphi, m)$.

Consider next a multi-product firm operating in several destination markets. The firm average emission intensity in value is

$$\overline{EV}(\varphi) = \frac{\sum_{h=1}^{H} \sum_{m=0}^{M_{lh}(\varphi)-1} EQ(\varphi, m) q_{lh}(\Phi(\varphi, m))}{\sum_{h=1}^{H} \sum_{m=0}^{M_{lh}(\varphi)-1} r_{lh}(\Phi(\varphi, m))}. \tag{A.27}$$

We must combine the impacts of a demand shock in country h on both the product-mix and the price channels. When both channels reinforce each other – when core products are cleaner (dirtier) and the export price to country h is higher (lower) – we conclude that exporting more (for efficient firms) tend to reduce the average emission intensity in value. When both channels counteract each other, however, the impacts of a demand shock on country l's firms are ambiguous.

A.2 Data Appendix

Supporting Materials Relating to Product-Specific Energy Reports

Figure A.1: Amendments to Section 217(1)(e) of the Indian Companies Act, 1988

Companies (Disclosure of Particulars in the Report of Board of Directors) Rules, 1988

In exercise of the powers conferred by section 642 read with clause (e) of sub-section (1) of section 217 of the Companies Act, 1956 (1 of 1956), the Central Government hereby makes the following rules, namely :-

1. (1) These rules, may be called the Companies (Disclosure of Particulars in the Report of Board of Directors) Rules, 1988.

   (2) They shall come into force on the 1st day of April, 1989.

2. Every company shall, in the report of its board of directors, disclose particulars with respect to the following matters namely :-

   A. Conservation of energy :

   (a) energy conservation measures taken ;
   (b) additional investments and proposals, if any, being implemented for reduction
   of consumption of energy ;
   (c) impact of the measures at (a) and (b) above for reduction of energy consumption and consequent impact on the cost of production of goods ;
   (d) total energy consumption and energy consumption per unit of production as per Form A of the Annexure in respect of industries specified in the Schedule thereto.

Notes: Figure presents the 1988 amendments section 217 (1)(e) to the Indian Companies Act of 1956 relating to the disclosure of energy-use reporting. Emphasis added by the authors to highlight the language specific to product-specific energy-use. Source is Ministry of Corporate Affairs, Government of India http://www.mca.gov.in/Ministry/actsbills/rules/CDoPitRoBoDR1988.pdf

Figure A.2: Sample Product-Specific Energy Intensity Report, Form-A Part B

B. Consumption per Unit of Production

	ELECTRICITY (KWH/TONNE)		COAL (M.T./ TONNE)		FURNACE OIL (K.L./ TONNE)		OTHERS/INTERNAL GENERATION (M.T./TONNE)	
	CURRENT YEAR 2000-2001	PREVIOUS YEAR 1999-2000	CURRENT YEAR 2000-2001	PREVIOUS YEAR 1999-2000	CURRENT YEAR 2000-2001	PREVIOUS YEAR 1999-2000	CURRENT YEAR 2000-2001	PREVIOUS YEAR 1999-2000
PAPER	**1,421**	1,436	**1.549**	1.492	**0.003**	0.006	**0.880**	0.881
CAUSTIC SODA	**3,081**	3,063	**0.155**	0.190	**0.000**	0.000	**–**	–
RAYON GRADE PULP	**1,026**	999	**1.018**	1.071	**0.095**	0.093	**–**	–

Notes: Figure presents sample Annexure to Directors' Report Form-A, Part B for Ballarpur Industries Limited fiscal year 2000-2001. Report is publically available at http://www.bilt.com/annual/photo/img/pic70.pdf

Computing Emission from Energy-Use Data

Firms in Prowess report energy-use in two ways. First, firms report the total quantity consumed each year by energy source (e.g., liters of diesel, KwH of electricity, etc.). Second, firms report energy intensity of production by output product. That is, for each product sold, firms report the amount of each energy source used to generate a single unit of the good. We refer to the first report as the "firm-level" energy data, while the second we call the "product-specific" energy data. For both reports, we translate physical quantities of energy consumed into physical quantities of CO_2 emissions, which we take as our measure of pollution.

For each energy source reported by a firm (in either the firm-level or the product specific data), we assign a CO_2 intensity measure based on emission factors from the US EPA 2012 Climate Registry Default Emissions Factors. CO_2 intensities are reported per unit of energy source (e.g., short ton of Lignite), and per mmBTU of energy. The list of energy types and CO_2 emissions factors are listed in Table A.1. There are 25 energy sources described by the EPA report, to which we add electricity generation, for a total of 26 emission intensities. We take the emission intensity of electricity purchased from the grid from X. We assign by hand each of the 140 energy sources reported in Prowess to one of the 26 energy types in Table A.1.

Energy-use is reported in physical quantities of the energy source, while CO_2 emissions factors are also reported in physical quantities of CO_2 per physical quantity of energy source. In order to translate the energy-use into CO_2 emissions, the units of energy consumption in Prowess must match the units of energy used in Table A.1 (i.e., the denominator of the CO_2 emissions factor). Units of energy in the EPA data are either scf, short ton, or gallon, but firms in Prowess report units in a much wider range of measurements. In the firm-level data, firms report in any of 55 different units, for a total of 412 source-unit combinations. We standardize units when possible to match to the EPA data, but some source-unit combinations cannot be converted to a usable figure. For example, one firm reports cubic meters of biomass consumed. While it seems reasonable to assign "biomass" to "Agricultural Byproducts" in Table A.1, the unit in Prowess is denominated in volume, while the unit from the EPA is in mass. Without assuming a density of the "biomass", there is no way to convert the energy source quantity into CO_2. We drop all such cases, which amounts to a little under 1% of the data. After standardizing units, we multiply the consumption of physical units of energy source in Prowess by the CO_2 intensity in Table A.1. We then sum across energy sources in a year to compute CO_2 pollution for each firm-year.

One refinement we make is to leverage information about own-generation of electricity contained in Prowess itself. In the product-specific data, sometimes the firm reports "Electricity" as an output. In this case, we know the firm-specific CO_2 emission intensity of electricity production from above. We merge in this firm-specific CO_2 emission intensity...

Finally, we clean the data for outliers. Upon inspection of excessively large emissions or emission intensity values, it appears in many cases as if decimals have been transposed or units mis-reported. We adopt the standard approach of dropping the top and bottom 1% of values for emission intensity for most of the analysis. Additionally, to address the problem of egregious measurement errors, we drop firms that exhibit excessive variation in emission intensity over the period. If a firm's total emission intensity in value increases by more than a factor of 10 between two years, than we drop the firm from the sample.

Table A.1: CO_2 emission factors

Energy Source	Kg CO_2 per Unit of Energy Source	Unit of Energy Source	Kg CO_2 per MMBTU of Energy Source
Acetylene	.1053	scf	71.61
Agricultural Byproducts	974.9	short ton	118.17
Anthracite	2597.82	short ton	103.54
Biogas (Captured Methane)	.0438	scf	52.07
Coke	2530.59	short ton	102.04
Coke Oven Gas	.0281	scf	46.85
Distillate Fuel Oil No. 1	10.18	gallon	73.25
Distillate Fuel Oil No. 2	10.21	gallon	73.96
Electricity			278
Fuel Gas	.0819	scf	59
Kerosene	10.15	gallon	75.2
Kraft Black Liquor	1131.11	short ton	94.42
LPG	5.79	gallon	62.98
Lignite	1369.28	short ton	96.36
Lubricants	10.69	gallon	74.27
Motor Gasoline	8.78	gallon	70.22
Naptha (<401 deg F)	8.5	gallon	68.02
Natural Gas (US average)	.0545	scf	53.02
Petroleum Coke (Liquid)	14.64	gallon	102.41
Petroleum Coke (Solid)	3072.3	short ton	102.41
Propane (Liquid)	5.59	gallon	61.46
Residual Fuel Oil No. 6	11.27	gallon	75.1
Solid Byproducts	2725.32	short ton	105.51
Wastewater Treatment Biogas			52.07
Waxes	9.57	gallon	72.6
Wood and Wood Residuals	1442.64	short ton	93.8

Notes: The first column lists the energy source as named by the EPA. Prowess does not use exactly the same naming convention, so we mapped by hand these energy types to the energy types listed in Prowess. The second column reports kg CO_2 associated with a given unit of energy type in column 1, where the unit is reported in column 3. For most energy types, we use the CO_2 intensity listed in column 2. However, for some observations, we were unable to standardize units across the two datasets. In some cases, we were able to use an alternative CO_2 intensity reported per mmBTU. We list this alternative CO_2 intensity in column 4.

Table A.2: High Income Countries

OECD	Other
Australia	Andorra
Austria	Antigua and barbuda
Belgium-Luxembourg	Aruba
Canada	Bahamas
Denmark	Bahrain
Finland	Bermuda
France	Brunei darussalam
Germany	Cayman islands
Greece	Cyprus
Iceland	French polynesia
Ireland	Greenland
Italy	Guam
Japan	Hong kong
Korea, republic of	Israel
Netherlands	Kuwait
New zealand	Macau
Norway	Malta
Portugal	Netherlands antilles
Spain	New caledonia
Sweden	Qatar
Switzerland, Liechtenstein	Saudi arabia
USA, Puerto Rico and US Virgin Islands	Singapore
United kingdom	Slovenia
	United arab emirates

Notes: This table reports the countries identified as "High Income" destinations. The designations come from the World Bank, based on 2006 GNI per capita. The first columns reports high-income OECD countries. The second column reports "other" countries defined by the World Bank as "High Income."

Analysis of Emissions data

Aggregate Emissions in Prowess

The distribution of CO_2 emissions produced by energy source, derived from the firm-level energy consumption data, are reported in Figure A.3. We have aggregated all 140 energy sources into 5 broad groupings. We calculate that total CO_2 emissions in manufacturing increased from 119 MT to 482 MT between 1990 and 2010 using the firm-level dataset. By comparison, over the same period, total CO_2 from India have increased from 690 MT to 2,009 MT, so the firms in Prowess account for about 1/4 total CO_2 emissions from India.[1] In terms of distribution, we calculate that in 2010 coal accounts for 45% of total CO_2 emissions in the firm-level energy reports, gas, diesel, and electricity each account for between 15-20%, and biofuel only 2%.

Comparison to WIOD

We also compare emission totals and intensities by industry to the recently constructed World Input Output Database (WIOD) database, which reports emission and output by industry for most large countries. Comparing the overall figure to the total CO_2 emissions estimates from the same 12 industry groupings in the WIOD, we find that Prowess firms account for 80% of manufacturing-based emissions in WIOD.[2] Furthermore, in Figure A.4 we compare emission intensities by industry in both the firm-level database and the product-specific database to the WIOD database. The x-axis records the average emission intensity in value from the WIOD database, while the y-axis reports the median emission intensity of production for either the firm-level dataset (blue diamonds) or the product specific dataset (red dots). Most industries lie very close to the 45-degree line, indicating a high correlation between the three reports. The two outliers are "Paper, Pulp, and Wood products", for which we compute a much higher emission intensity in Prowess than in WIOD, and "Minerals," for which we compute a much lower emission intensity in Prowess. These comparisons gives us reassurance that the data cleaning procedure generates plausible values.

Comparison between Firm-level and Product-level

While the aggregate quantities and intensities in the firm-level and product-specific datasets match fairly well to other known datasets, a natural concern with respect to the product-specific dataset is how the firms assess the energy intensity of individual product. Though the firms are required by law to report these data, how could they actually compute them? If production happens at the same location on the same machines, it seems unlikely that a firm would be able to meaningfully distinguish between the energy intensity used to make one product versus another. However, if production of different goods happens sequentially, or is segregated between machines, or plant, or time of year, then it seems much more likely that a firm could compute different energy intensities per product. For example, in the case presented in Appendix Figure A.2, we see that Ballarpur Industries Limited in the fiscal year 2000-2001 reports separate amounts of electricity, coal, furnace oil, and other/internal

[1] World Development Indicators Table 3.9, http://wdi.worldbank.org/table/3.9

[2] WIOD reports 586 MT of CO_2 in 2009, compared to 467 MT in Prowess

Figure A.3: Total CO_2 Emissions by Energy Source

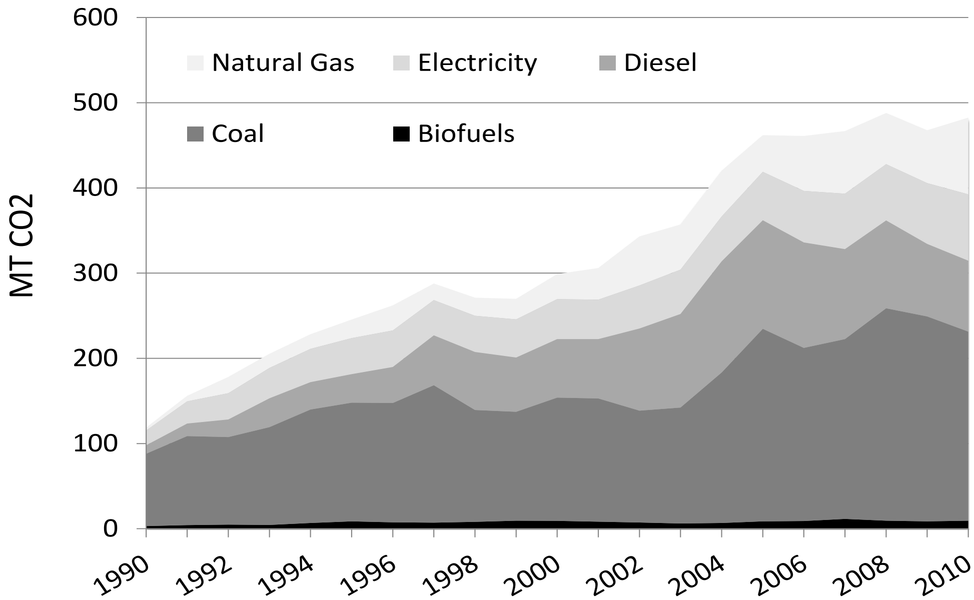

Notes: Figure reports total CO_2 emissions by year broken down by energy source as computed from the firm-level dataset. Energy sources have been aggregated into five major groupings.

generation each to make 1 tonne of paper, caustic soda, and rayon grade pulp. If paper is manufactured at a different time of day than rayon grade pulp, and the firm used machinery 40 percent longer to make a tonne of paper compared to rayon grade pulp, then it does not seem implausible that the firm could deduce that it requires roughly 40 percent more electricity to produce a tonne of paper compared to rayon grade pulp, as Ballarpur in fact reports. Still, one might be concerned about the signal to noise ratio in these data. To address this concern, we present several diagnostic tests that further bolster our confidence in the data.

First, we cross-check the implied emissions profile from the product-specific energy with the firm-level energy reports. The product-specific energy data requires some calculation on the part of the firm to determine energy intensity, while the firm-level energy report is a mere inventory of fuels consumed. If the two reports yield similar emissions profiles, we take it as evidence that there is some signal in the data. Aggregating the emissions in the product-specific dataset to the firm-year level and merging to the firm-level dataset, we have 7,777 firm-year pairs of emissions values to compare. We take the log ratio of the two values, order from lowest to highest, and plot in Figure A.5. Data points along the 0-line indicate agreement between the two datasets. Points off the 0-line indicate divergence between the two datasets. We find that 72% of the firm-year pairs have a ratio between 0.5 and 2, which we find reassuring, especially since some of the disagreement in the tails is due to differential cleaning procedures (see Appendix A.2).

Figure A.4: Emissions Intensity Across Datasets

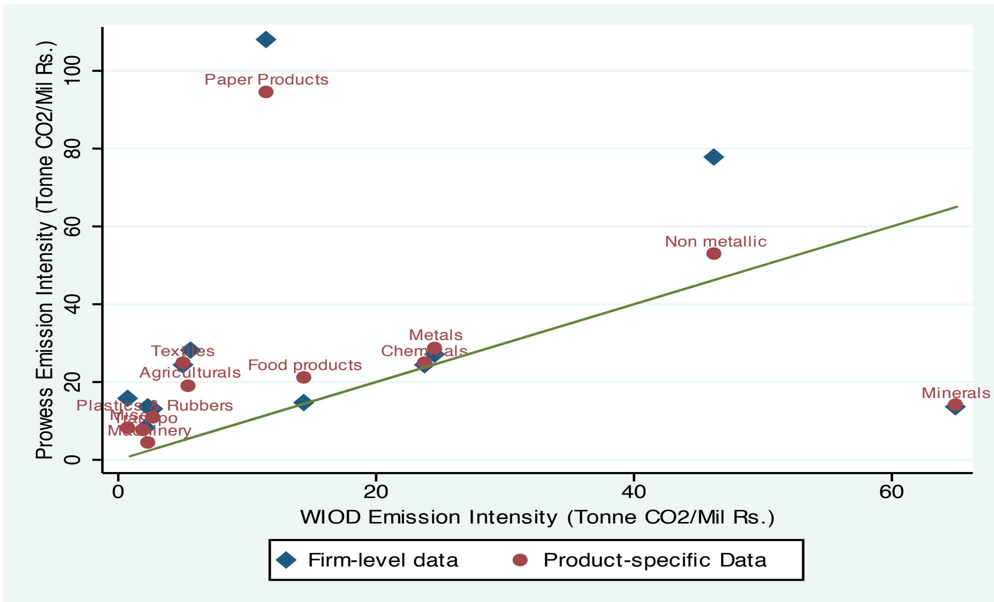

Notes: Emission intensity in tonnes of CO_2 per million Rs. from Prowess are reported on the y-axis, and from the WIOD database on the x-axis. Emission intensities from Prowess are the median value for firms in the given industry in the firm-level database (diamonds) and the median firm-product value in the product-specific database (circles). All values include observations from the entire period 1995-2009.

Testing Alternative Hypothesis

Finally, we examine alternative explanations for how firms compute product-specific energy shares. Even if the firm-level emissions reports match fairly well with the product-specific energy reports, it could be that firms still do not actually know the product-specific energy intensity, but instead just divide total energy-use along some convenient heuristic. Under this hypothesis, the aggregate of product-specific energy use would match the firm-level energy use, but the product-specific reports would still not reflect true emission intensity. In Figure A.6, we test the null hypothesis that energy shares in the product-specific data are based purely on the revenue share of different products within the firm. If it were the case, then the energy share should correlate perfectly with the revenue share.

Let Z_{eij} indicate the quantity of energy (in physical units - e.g., tons of coal, KwH of electricity, etc.) firm i uses from energy source e (e.g., coal, electricity, etc.) to manufacture Q_{ij} units of product j. The revenue earned from product j is $R_{ij} = Q_{ij} * P_{ij}$. Denote the revenue share and energy-type share associated with each product j:

$$r_{ij} = \frac{R_{ij}}{\sum_{j \in \Delta_i} R_{ij}} \quad , \quad z_{eij} = \frac{Z_{eij}}{\sum_{j \in \Delta_i} Z_{eij}} \tag{A.28}$$

for each energy source e, and where Δ_i is as before the set of products manufactured by a

Figure A.5: Comparing Emissions Profiles From Different Energy Reports

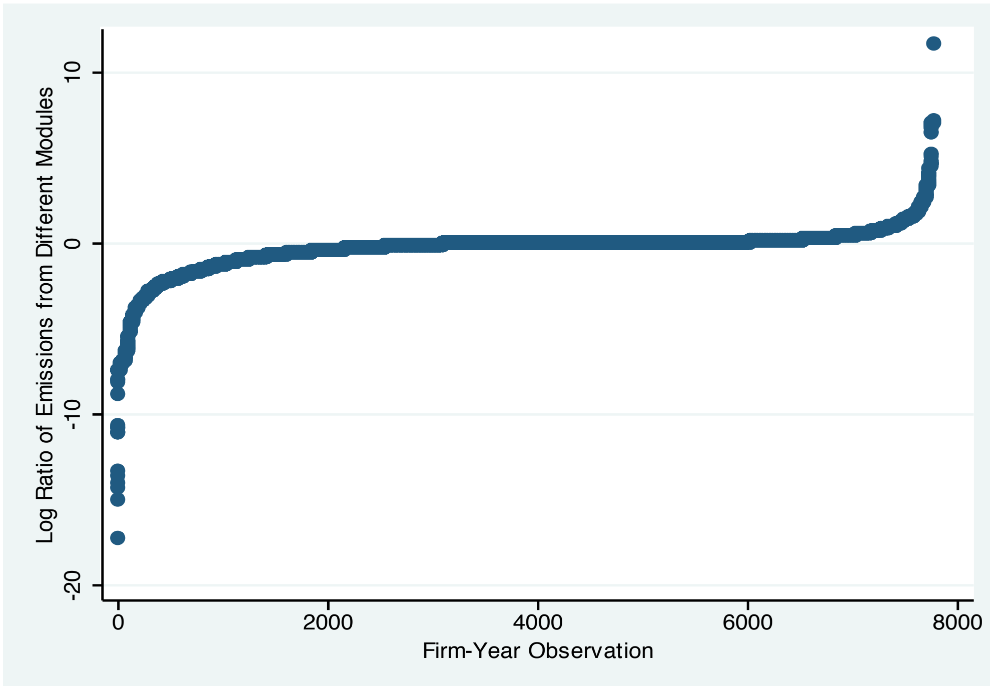

Notes: Each observation is a firm-year for which we have an emissions value from both the product-specific dataset and the firm-level dataset. Firm-year emission from the product-specific dataset are aggregated over products produced within the firm-year. Observations are ordered by the ratio of the two reports product-specific/firm-level. Y-axis reports the log ratio.

firm in a given year. If the firm assigns energy shares z_{eij} purely based on the revenue share r_{ij}, then the two should be perfectly correlated. To test this hypothesis, we estimate a linear regression model

$$r_{ijt} = \sum_{e \in \Gamma_{it}} \phi_e z_{eijt} + \gamma' W_{eijt} + \epsilon_{ijt}, \qquad \text{(A.29)}$$

where t denotes years and Γ_{it} represents the set of energy sources used by the firm in year t, and W_{eijt} represents fixed effect controls for year, energy source, and product category j. We merge the product-specific dataset to the output data according to the algorithm described in Appendix A.2, and estimate (A.29) for multi-product observations only. Under the null hypothesis, $\phi_e = 1$ for all e.

We present results graphically in Figure A.6. We estimate the model separately for products using 1, 2, or 3 different types of energy. That is, a firm uses only one energy source to produce a 1-source product, two sources of energy to produce a 2-source product, etc. The point estimate for the raw correlations are depicted with round dots, while squares indicate that the regression controls for the full suite of fixed effects. In each case, standard errors have been clustered at the firm-level and 95% confidence interval are depicted graphically by

Figure A.6: Correlation Between Revenue Share and Energy Shares in Product-Specific Data

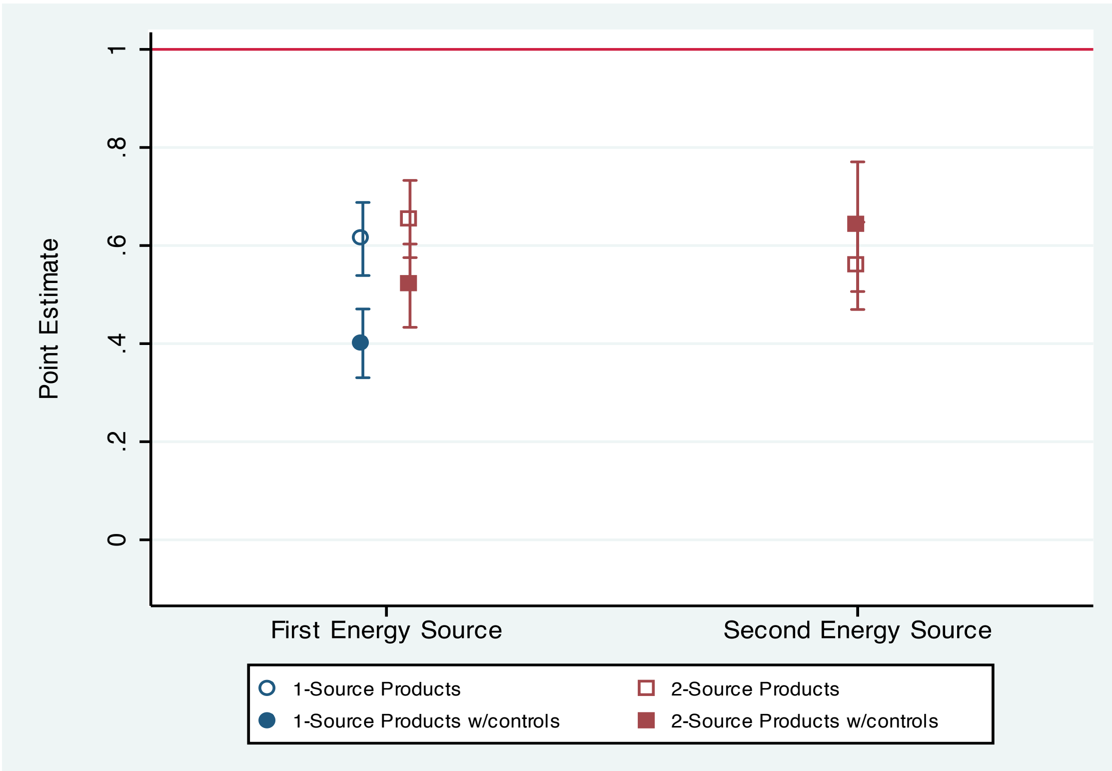

Notes: Point estimates from linear regressions are depicted by circles (for the regressions without controls) and squares (for regressions with controls). 95% confidence intervals are plotted by whiskers. Regressions are estimated separately depending on the number of energy sources utilized by the firm in the given year to produce the good.

the whiskers. In the case a firm uses more than one energy type to manufacture a product, the "First Energy Source," corresponds to the energy-type that accounts for the largest share of millions of British Thermal Units (mmBTU) in the production of the good, while "Second Energy Source" corresponds to the next largest, etc. Point estimates resulting from the same regression are connected by solid lines (for the raw correlations) and dashed lines (for the model with controls).

Figure A.6 shows that, in every case, we can reject a unit elasticity at the 1% level. Additionally, for the case of 2-source products, we can reject the equality of coefficients. The estimates for 3-source products becomes quite noisy as the sample size drops to less than 2000 data points, so we can no longer reject a null of equality, but we can still reject the null of unity. The estimates are in each case greater than 0, as one would expect that output share is increasing in input share, but it does not appear as though product-specific energy intensities merely reflect the revenue share of the product, which one might conclude from failing to reject $\phi_e = 1$. We consider Figure A.6 a reassuring check that the product-specific data is based on actual input usage.

Merging Product-specific data to Output data

In Figure A.4 in the main text, we compare emission intensity in value in the product-specific energy reports to the industry average emission intensity in value reported in the WIOD. In order to compute emission intensity in value, we must merge the product-specific energy data to the output data. As described in the main text, the merge is problematic because often times, neither product names nor Prowess product ID's are consistent between the input data and the output data. Lacking a common unique identifier, we design an alternative strategy for merging the two datasets.

To illustrate the problem, consider an example from the data. In 1994 ABT Industries Ltd reports "Fruit Juices" as an output in the output module (and no other output products), but reports energy intensity for manufacturing "Fruit & vegetable juices, concentrates" in the product-specific energy module (and no other output products). CMIE codes the former as Prowess ID 0511050100000000, while it codes the latter as Prowess ID 0511050000000000. When we merge on Prowess ID, we fail to merge these data, since they have different identifying codes, though clearly they describe the same product. If cases such as this were rare in the data, we could proceed by merging on the Prowess ID, but we find that near merges such as these represent a large share of the data.

Our strategy for dropping as little information as possible while merging the two datasets is to leverage the tree-structure of the CMIE Prowess ID. In the product classification, similar products share the same beginning digits. For example, in the fruit juice case, both products are given prids that begin with "051105." Thus, from information contained within the Prowess ID, we can assess the similarity of products. If two products do not merge for a firm-yr initially, we try successively to merge on higher levels of aggregation as indicted in the Prowess ID. For example, in the Fruit Juice case we would ultimately merge on the 6-digit identifier 051105. This procedure increases the sample size by about 30%.

Merging Trade Data to Prowess

In this section, we discuss how we merge the trade instruments to Prowess. Merging trade data to Prowess is problematic because CMIE classifies products in Prowess according to its own 16-digit codes, which do not map directly to any other classification system at a disaggregated level. We generate our own mapping that connects the Prowess ID code directly to HS trade classifications. This mapping allows for a tighter link between (HS) product-specific shocks and production activity in Prowess firms.

There are 3,340 distinct 16-digit Prowess ID codes in the Prowess dataset which we aim to map to 5,108 HS6 revision 1996 codes in the trade data. We assign correspondences between the two by hand, exploiting the fact that both Prowess and the HS system hew fairly closely to the ISIC classification. Since both classification systems spring from a common source, names and orderings are fairly similar between the two.

To begin, we map every one of the 16-digit Prowess ID that we could to HS trade classifications at either the 2-digit, 4-digit or 6-digit level. The mapping utilizes both the product name in the two datasets as well as the numerical ordering to generate correspondences. For example, consider the six products recorded in Prowess under the heading "Silk and silk textiles":

Prowess ID	Product Name
601010000000000	Silk worm cocoons
601020000000000	Raw silk
601030000000000	Silk waste
601040000000000	Silk yarns
601050000000000	Woven fabrics of silk
601060000000000	Silk fabrics, processed

Now consider the seven HS4 products reported under the heading "Silk" in the HS trade data:

HS4	product name
5001	Silk-worm cocoons suitable for reeling
5002	Raw silk (not thrown)
5003	Silk waste (including cocoons unsuitable for reeling, yarn waste and garneted stock)
5004	Silk yarn (other than yarn spun from silk waste) not put up for retail sale
5005	Yarn spun from silk waste, not put up for retail sale
5006	Silk yarn and yarn spun from silk waste, put up for retail sale; silk-worm gut
5007	Woven fabrics of silk or of silk waste

Examining the names of the products in the two classifications, it seems obvious that there is a correspondence and that order is well preserved: in both systems, "Silk worm cocoons" comes first, directly followed by "Raw silk", directly followed by "Silk waste", etc. The mapping is not simply a matter of harmonizing names, however. We can see, for example, that the HS system distinguished between three kinds of silk yarn (HS4 5004, 5005, 5006), whereas CMIE lumps them all together into a single category (Prowess ID 601040000000000). And disaggregating the opposite way, CMIE discriminates between "Woven fabrics of silk" (601050000000000) and "Silk fabrics, processed" (601060000000000), whereas HS aggregates to the single HS4 "Woven fabrics of silk or of silk waste" (5007). Mapping requires some judgment, but it seems fairly obvious that the correspondence should be:

Prowess ID	Product Name	HS4
601010000000000	Silk worm cocoons	5001
601020000000000	Raw silk (not thrown)	5002
601030000000000	Silk waste	5003
601040000000000	Silk yarn	5004, 5005, 5006
601050000000000	Woven fabrics of silk	5007
601060000000000	Silk fabrics, processed	5007

In the case of silk, products map relatively well on HS4. However, in other cases, the finer HS6 classification is needed. For example, CMIE distinguishes between "Floor coverings of coir" (608010000000000) and "Floor coverings of wool" (608020000000000). The logical HS4 category would be "Carpets & other textile floor coverings, woven, of pile construction, not made up" (5702). However, moving to the finer HS6 classification, we find that "Floor coverings of coconut fibres (coir)" are coded as HS 570220, whereas "Carpets...of wool/fine animal hair" are coded HS 570231. At this level of disaggregation, we can properly map:

Prowess ID	Product Name	HS6
608010000000000	Floor coverings of coir	570220
608020000000000	Floor coverings of wool	570231

The procedure works as follows: First, we map each Prowess ID (if possible) to the appropriate HS2, HS4, or HS6 code. Note that there could be more than one HS code that maps to a given Prowess ID, as in the case of "silk yarn" above. We have enough detail to match 1,497 Prowess IDs to either one or possibly multiple HS codes. When the matched HS code is at the 2-digit or 4-digit level, we take the simple average of $\widetilde{D}_{jt}$ over the HS2 or HS4. When there are multiple HS codes that merge to the same Prowess ID, we take the simple average over all HS codes to compute a unique demand shock for each Prowess ID-year. When the demand shocks have been computed at the Prowess ID level, we will denote them $\widetilde{D}_{pt}$, as opposed to $\widetilde{D}_{jt}$ when they are at the HS6-level.

Appendix B

Appendix to Chapter 2

B.1 Data Appendix

Trade Data

Quota restrictions under the MFA for textile and apparel imports into the US were managed by the Commerce Department's Office of Textile and Apparel (OTEXA). Quota limits were product-country-year specific and tended to persist through time, so if a quota was binding when it was first established, it tended to bind until liberalization. OTEXA kept annual "progress reports" on each quota going back to 1984, listing the exporting country and year along with the quota limit and fill rate, i.e. how much of the quota was used. All progress reports were obtained in text file by Peter Schott, digitized, and published on his website along with the code book (Brambilla, Khandelwal, and Schott, 2010).

OTEXA assigned each textile and apparel import product, classified according to HS10, to one of 167 3-digit quota category groups (g). The raw data from Brambilla, Khandelwal, and Schott (2010) report fill rates for 16,416 quota group-country-year observations, across 64 countries and 21 years, where "quota group" could signify a single 3-digit OTEXA category, or a partial or merged category. In most cases, these 3-digit categories constitute the unit of quota administration, with a single quantity restriction applied to each category-country pair in each year, though in some cases quotas were assigned to partial or merged categories. In some cases, quotas were assigned to subsets or aggregates of the 167 3-digit categories. For example, in 2004, the US restricted imports from India in the partial category "341-Y", where the parent category "341" indicates "Women's cotton non-knitted blouses" and the "-Y" indicates that the quota only applies to "blouses of warp/fill material." In another example, in 2004, imports from India in the categories "Men and Boy's cotton trousers"(347) and "Women and Girl's cotton trousers"(348) were regulated together with a joint quota applied to merged category "347/348."

Computing $India_{gt}$ is a simple matter of assigning any quota-group-year observation from India with a fill rate greater than 90% to take the value 1, and 0 otherwise. However, computing ROW_{gt} requires merging the quota data with trade flows, which are classified at the HS10 level. Additionally, both variables must then be merged to production data. A problem arises here in that while 3-digit quota categories map to HS10 products (via a mapping published on the OTEXA website), partial categories do not. For example, in the case above, there is no way to know precisely which HS10 products in the 3-digit category 341 map to the partial category "341-Y." There is no straightforward way to map this quota group to trade data or production data. Thus, for the purposes of calculating ROW_{gt} and for merging to production data, we must translate raw fill rates at the quota-group level (e.g. 341-Y) into fill rates at the 3-digit category level (e.g. 341). Protection indicators can then be merged to trade and production data.

We proceed in two steps, treating first partial then merged quota-groups. In the event that a partial group is subject to a quota, there are two cases. In one case, both the partial group and the full "parent" quota group are both subject to quotas for a given country-year. For example, in 2001, Taiwan was constrained in both the partial category "mmf not-knit shirts, mb, yn-dyed" (640-Y), and the parent category "m&b not-knit mmf shirts" (640). In this case, we drop the "child" category 640-Y and assign the fill rate from the parent 640 to all HS10s in the 3-digit category. Our reasoning here is that regardless of the fill rate for the partial category, all HS10s in 640 (even those from 640-Y) are subject to the quota from the higher level of aggregation. Thus, we know at minimum, if the parent quota is binding, then the HS10s covered by the partial-category quota are also binding. In a second case, the child group is regulated, but the parent group is not. For example, in 2002 in Bangladesh, 369-S is subject to quota, while the parent category 369 is not. In this case, we know at minimum that the HS10s comprised by the partial 369-S are subject to the quota from 369-S. However, not knowing precisely which HS10s within 369 those are, we assign the fill rate from the partial 369-S to all HS10s in the parent 369.

Next, when OTEXA merges two categories and assigns a single quota level to the merged group, again, two cases might occur. In one case, OTEXA assigns a quota to both the merged quota group and one or both of the individual component groups. For example, imports from South Korea in 2003 were constrained in category 347 individually, and the merged group "347/348." In this case, we take the maximum of the fill rate between these two quotas. The logic here is that products in 347 can only be subject to one level of restraint. If the quota 347 is binding, even if the quota 347/348 is not binding, products in 347 are still constrained (through the individual quota). Whereas if the quota 347/348 is binding and 347 is not, products in 347 are still subject to binding constraints (through 347/348). In the final case, in which products are regulated under a merged quota and there is no corresponding quota for the individual categories in the country-year, we simply assign the fill rate from the merged category to each component category.

Production Data

In this section, we describe how we map between firm-product output data and firm-product energy input usage. In this step, the main challenge is that product names are often inconsistent across the two modules (output and energy). Treating the inconsistency as reporting error (e.g., different people within the firm entered data for the differnt modules and did not apodt a common naming convention), we seek to reconstruct the true production information from the data available.

A candidate procedure would be to merge on the Prowess classification ID ("prid"). However, while prid helps in some cases, there can be multiple products within a firm-year assigned to the same prid, so merging on prid does not provide a unique mapping.

Additionally, assignment of product names to prids appears to be fairly inconsistent across the two modules as well, even for similarly named products.

Lacking a pre-existing variable on which to merge the datasets, we investigate each firm in the dataset and match the energy data to the output data by hand. We utilize both the reported names and sometimes the prid for guidance. An important point to note is that the dimension of the output data and the energy data is not always the same. Based on the naming and coding conventions, we infer in some cases that a single entry in the energy data refers to multiple products in the output data and vice versa. We illustrate the procedure with an example.

Consider the products reported in the output data by the firm Arihant Industries Ltd:

prid	product name (output data)
601060000000000	Processing Of Art Silk Fabrics
602060000000000	Worsted Yarn
603070101010000	Grey Cloth
605010203000000	Polyester Filament Yarn
605010203000000	Texturised Yarn
605010204040000	Acrylic Yarn

compared to the products reported in the energy data

prid	product name (energy data)
603050000000000	Texturised Yarn
605010200000000	Polyester/Viscose & Acrylic Worsted Yarn
603080000000000	Processed Cloth
603070500000000	Synthetic Fabrics

Neither prids nor product names provide direct matches for any products beyond one – “Texturised Yarn.” Also, the dimensions of the datasets are not the same: we have 6 products in the output data, but only 4 products in the energy data. Clearly, some judgment will be required to merge these two reports.

We proceed as follows. We assume “Texturised Yarn” refers to the same product in the two datasets, eventhough CMIE coded the two entries with different prids. We suppose “Polyester/Viscose & Acrylic Worsted Yarn” in the energy data is an aggregate of “Worsted Yarn,” “Polyester Filament Yarn,” and “ Acrylic Yarn” in the output data and assign its energy intensity to all three of these products. With energy intensities assigned to 4 of the 6 output data products, we have left “Grey Cloth” and “Processing Of Art Silk Fabrics” in the output data and “Processed Cloth” and “Synthetic Fabrics” in the energy data. We match “Grey Cloth” to “Processed Cloth” and ‘Processing Of Art Silk Fabrics” to “Synthetic Fabrics,” noting that silk fabric is a synthetic. To each product in the output data, we assign a firm-product idenitfyier (“id”) and match the appropriate energy product name as follows:

id	Output product name	Energy product name
220	Acrylic Yarn	Polyester/Viscose & Acrylic Worsted Yarn
222	Grey Cloth	Processed Cloth
223	Polyester Filament Yarn	Polyester/Viscose & Acrylic Worsted Yarn
224	Processing Of Art Silk Fabrics	Synthetic Fabrics
226	Texturised Yarn	Texturised Yarn
228	Worsted Yarn	Polyester/Viscose & Acrylic Worsted Yarn

We construct analogous mappings for every firm in the dataset and merge energy intensities to the output data. Finally, we again need to standardize output units between the two datasets, dropping observations for which units cannot be converted (less than .01% of the data). The resulting dataset contains 12,071 firm-product-year observations, comprising 813 firms and 1,436 firm-products.

Mapping Prowess to MFA Categories

To estimate trade impacts, we must merge the quota constraints constructed in section 2.2 to the production data from section 3. The merging procedure is nontrivial, so we describe it here in detail. The challenge is that no mapping exists between Prowess product codes and the OTEXA 3-digit classification system. However, in constructing the Prowess classification system, CMIE hewed very close to the ISIC nomenclature, which itself maps to four or six digit HS codes ("HS4" or "HS6"). Our strategy is to map Prowess product codes ("prids") to HS4 or HS6 by hand, and then map to the OTEXA categories through the OTEXA-HS10 map listed on the OTEXA website. As one might imagine, the mapping will not be one-to-one. In many cases, multiple OTEXA categories will map to the same prid. This feature of the mapping introduces noise into the estimation, though we argue there is no reason to suspect the measurement error biases results in any direction.

To begin, we mapped every one of the 394 textile and apparel products in the Prowess prid classification that we could to HS trade classifications at either the 4-digit or 6-digit level. The mapping utilizes both the product name in the two datasets as well as the numerical ordering to generate correspondences.

For example, consider the six products recorded in prid under the heading "Silk and silk textiles":

prid	product name
601010000000000	Silk worm cocoons
601020000000000	Raw silk
601030000000000	Silk waste
601040000000000	Silk yarns
601050000000000	Woven fabrics of silk
601060000000000	Silk fabrics, processed

Now consider the seven HS4 products reported under the heading "Silk" in the HS trade data:

HS4	product name
5001	Silk-worm cocoons suitable for reeling
5002	Raw silk (not thrown)
5003	Silk waste (including cocoons unsuitable for reeling, yarn waste and garneted stock)
5004	Silk yarn (other than yarn spun from silk waste) not put up for retail sale
5005	Yarn spun from silk waste, not put up for retail sale
5006	Silk yarn and yarn spun from silk waste, put up for retail sale; silk-worm gut
5007	Woven fabrics of silk or of silk waste

Examining the names of the products in the two classifications, it seems obvious that there is a correspondence and that order is well preserved: in both systems, "Silk worm cocoons" comes first, directly followed by "Raw silk", directly followed by "Silk waste," etc. The mapping is not simply a matter of harmonizing names, however. We can see, for example, that the HS system distinguished between three kinds of silk yarn (HS4 5004, 5005, 5006), whereas prid lumps them all together into a single category (prid 601040000000000). And disaggregating the opposite way, pird discriminates between "Woven fabrics of silk" (601050000000000) and "Silk fabrics, processed" (601060000000000), whereas HS aggregates to the single HS4 "Woven fabrics of silk or of silk waste" (5007). Mapping requires some judgment, but it seems fairly obvious that the correspondence should be:

prid	prid name	HS4
601010000000000	Silk worm cocoons	5001
601020000000000	Raw silk (not thrown)	5002
601030000000000	Silk waste	5003
601040000000000	Silk yarn	5004, 5005, 5006
601050000000000	Woven fabrics of silk	5007
601060000000000	Silk fabrics, processed	5007

In the case of silk, products map relatively well on HS4. However, in other cases, the finer HS6 classification is needed. For example, Prowess distinguishes between "Floor coverings of coir" (608010000000000) and "Floor coverings of wool" (608020000000000). The logical HS4 category would be "Carpets & other textile floor coverings, woven, of pile construction, not made up" (5702), however, this level of aggregation does not respect the coir vs wool distinction that we have in the prid. However, moving to the finer HS6 classification, we find that "Floor coverings of coconut fibres (coir)" are coded as HS 570220, whereas "Carpets...of wool/fine animal hair" are coded HS 570231. At this level of disaggregation, we can properly map:

prid	prid name	HS6
608010000000000	Floor coverings of coir	570220
608020000000000	Floor coverings of wool	570231

The procedure works as follows: First, we map each prid (that we can) to the appropriate HS6 or HS4. We have enough detail to match 318 of the 394 prids to either an (or possibly multiple) HS4 or HS6 in the trade classification (we discuss the remaining 76 prids below). Next, we take the OTEXA mapping from 3-digit categories to HS10 and collapse to the HS4 and HS6 level. The first four or six digits of the HS10 code identifies the HS4 or HS6. While this truncation procedure maps each HS4 and HS6 in the trade data to the OTEXA categorization, it does not necessarily map to a *unique* category. Out of 812 HS6 codes, 397 map to a unique category. The median number of categories is 2 with a standard deviation of 3.2. At the HS4 level, 33 codes map to a unique category (out of 130 HS4). The median number of categories is 3 with a standard deviation of 8. We then map categories to prids through either the HS6 (when possible) or HS4 (when not) and take the simple average of $India_{gt}$ and ROW_{gt} over the matched categories.

The non-uniqueness of HS-OTEXA mapping is the largest source of noise generated by the mapping. The multiplicity means that some prids have several potential protections rates that could characterize their trade barriers with the US. The uncertainty is not a true feature of the world, but rather a byproduct of mismatched data aggregation. However, as long as the noise enters randomly, the multiplicity is only a problem for statistical power, and does not bias the estimates. Furthermore, the multiplicity is again only a problem if the multiple categories in question have substantially different protection values. In many instances, this is not the case. For example, prid 603030100000000 "Cotton yarn" matches through HS4s 5205 and 5206 to OTEXA categories 300 and 301, "Carded cotton yarn" and "Combed cotton yarn." In 2002, India was subject to quota in neither category, so it does no violence to take the simple average of $India_{gt}$ in this case. The competition index scores 0.04 for Carded cotton yarn in this year and 0.09 in Combed cotton yarn. Assigning $ROW_{gt} = 0.065$ (ie the simple average) does not seem unreasonable. While this example was chosen to illustrate that averaging over multiple categories need not introduce noise into the instruments, obviously, in some cases it will. However, we mention it to impress the point that actual noise generated by averaging should be lower than might by inferred from the sheer count of non-unique matches.

Lastly, we must treat the remaining 76 prids that did not map to an HS4 or HS6. Most cases in which we cannot map an HS code to a prid occur because the prid classification is too broad. For example, 25 observations in the Prowess dataset are classified as simply "Silk & silk textiles" (601000000000000). We showed above that silk products map exclusively to one of six disaggregated prid codes, but sometimes, the firm does not give enough detail in it's description of the product for CMIE to distinguish among the finer disaggregated categories, though it can identify the product as "Silk." In such cases, we can either drop the observation entirely, or assign it an average protection value of all possible disaggregated product codes. For example, in the case CMIE codes a product 601000000000000, we assign the simple average of $India_{gt}$ and ROW_{gt} from the six silk products that share the same 3-digit heading as the observation in question (ie "601"). We chose to impute averages instead of dropping observations because while category-averages are not as informative as precisely matched values, they still carry some information; and there's no harm in imputing averages. At worst, it introduces more noise, but the increased sample size due to inclusion likely makes the estimate less noisy, rather than more noisy.

Appendix C

Appendix to Chapter 3

In this Appendix, we present the methodology for computing supply, price, land-saving, and GHG impacts of the GE technology. We also present the results in 4 Figures and 1 table that are discussed in the main text.

C.1 Supply Effect

We compute the supply effect of GE technology for the three principle GE crops as the percentage difference between observed 2010 production and two different counterfactual supplies corresponding to different assumptions about the extensive margin. Counterfactual supplies are computed country by country and then aggregated to a world figure.

We first compute the implied traditional variety yield $\widehat{y_{it0}}$ by solving

$$\begin{aligned} Q_{ct} &= y_{ct0}L_{ct0} + y_{ct1}L_{ct1} \\ &= y_{ct0}\left(L_{ct0} + \left(1+\hat{\beta}\right)L_{ct1}\right) \\ \Longrightarrow \quad & \widehat{y_{ct0}} = \frac{Q_{ct}}{L_{ct0} + \left(1+\hat{\beta}\right)L_{ct1}} \end{aligned} \tag{C.1}$$

where $\hat{\beta}$ represents the estimated yield effect of the GE technology for the given crop. In the estimated impacts that follow, we use both our own estimated yield impacts from section 3, and a range of other yield impacts from the literature. Assuming that production would have occurred on extensive margin lands even without the use of GE technology, then the counterfactual supply is given by

$$\widetilde{Q_{ct}} = \widehat{y_{ct0}}L_{ct} \tag{C.2}$$

We sum over country-specific counterfactual supplies to find the world total counterfactual supply $\widetilde{Q_t}$ and compute supply effect $\widetilde{s} = \frac{Q_t - \widetilde{Q_t}}{Q_t}$. If however, it is assumed that production on the extensive margin would not have occurred without the GE technology, i.e., that GE seeds cause the increase in hectarage, then the production on the extensive margin would have to be subtracted from $\widetilde{Q_{ct}}$ to yield counterfactual supply:

$$\widetilde{\widetilde{Q_{ct}}} = \widehat{y_{ct0}}\left[L_{ct} - L_{ct}^{ext_1}\right] \tag{C.3}$$

Figure C.1: Supply Effect of GE Corn

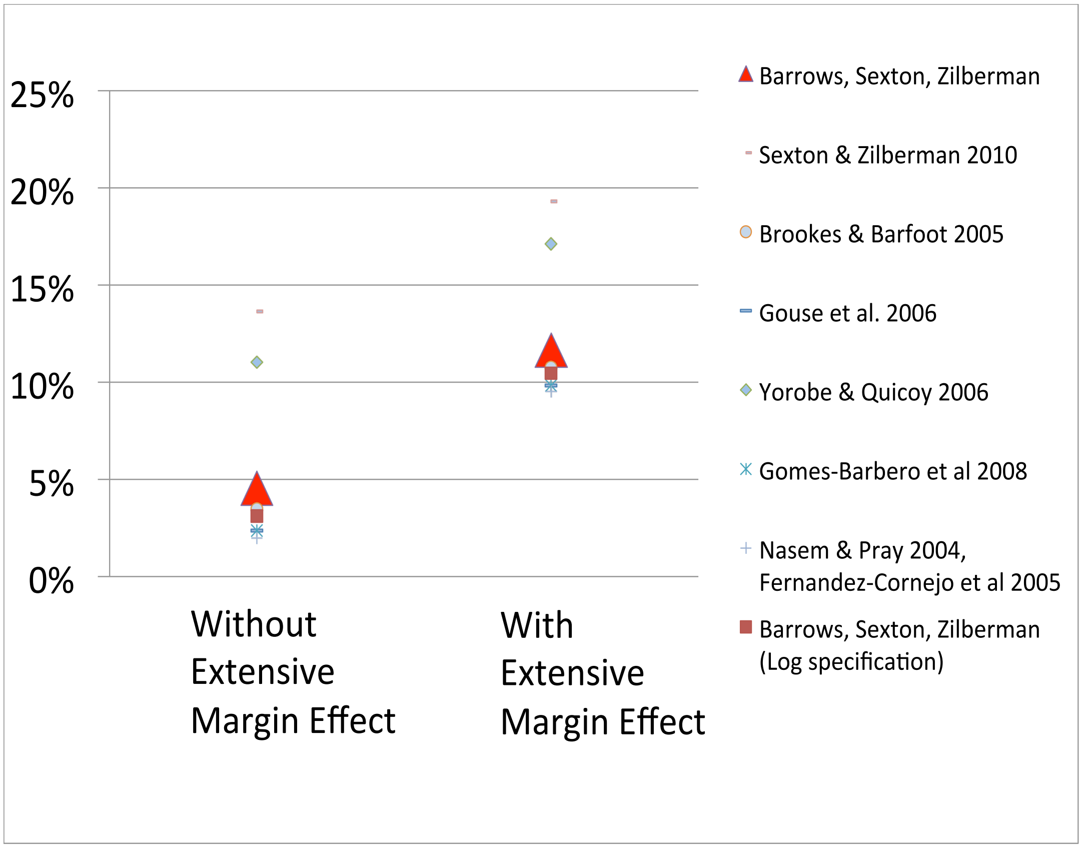

Notes: Supply effect of GE corn calculated as percentage difference between observed supply and counterfactual supply without GE technology. Estimate without extensive margin effect allow that production on the extensive margin would have occurred with the traditional technology as well. Estimates with extensive margin effect subtract all production on extensive margin in the counterfactual supply. Each point corresponds to estimates based on the yield effect from different studies in the literature. The "Barrows, Sexton, Zilberman" estimates are derived from our preferred yield estimates in Table 3.3 (column 4). "Barrows, Sexton, Zilberman (Log specification)" estimates are derived from the log specification in column 6 of Table 3.3.

where $L_{ct}^{ext_1}$ denotes the extensive margin computed in Section 4. The corresponding supply effect is defined analogously as above $\widetilde{\widetilde{s}} = \frac{Q_t - \widetilde{\widetilde{Q}}_t}{Q_t}$.

In Figures C.1 and C.2, we report world supply effect for GE corn and cotton for the year 2010 conditional on yield effects from several different studies. Supply effects based on our estimates from section 3 are denoted with large red triangles. Other markers correspond to the supply effects based on yield effects from Sexton and Zilberman (2011) along with all the studies reviewed in Qaim et al. (2009). Estimates are reported according to the extensive

Figure C.2: Supply Effect of GE Cotton

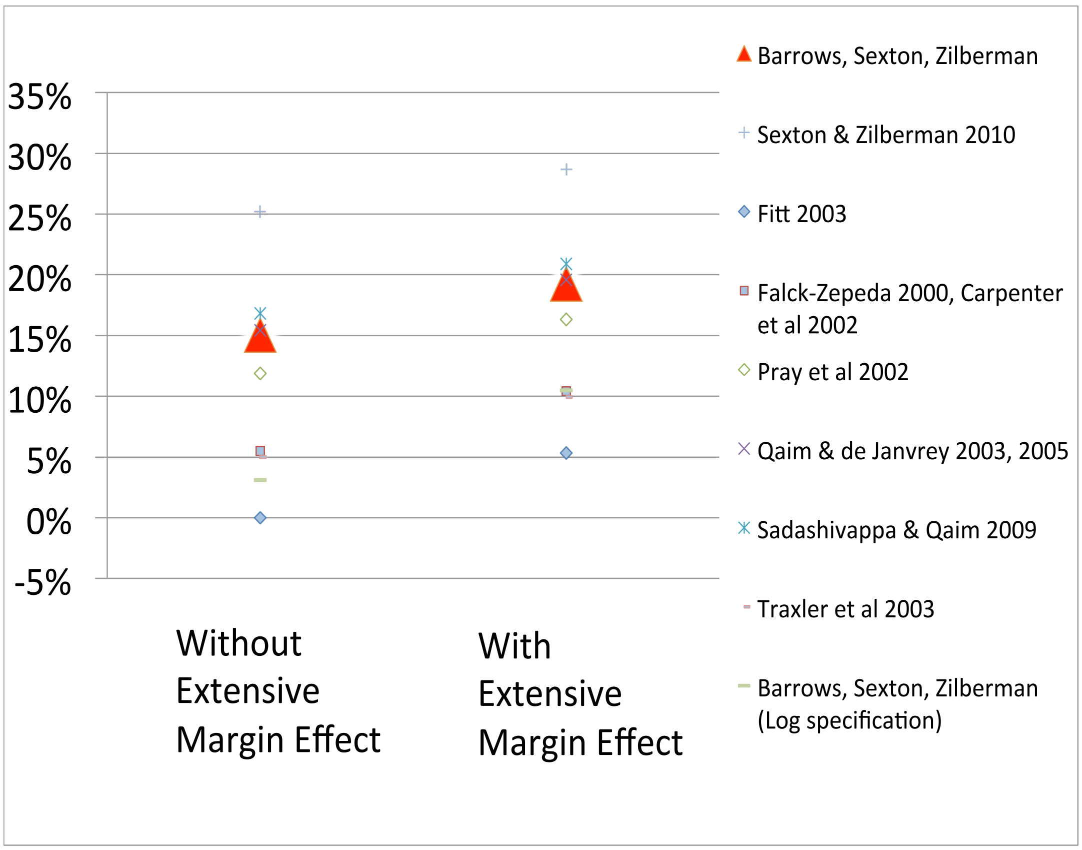

Notes: Supply effect of GE cotton calculated as percentage difference between observed supply and counterfactual supply without GE technology. Estimate without extensive margin effect allow that production on the extensive margin would have occurred with the traditional technology as well. Estimates with extensive margin effect subtract all production on extensive margin in the counterfactual supply. Each point corresponds to estimates based on the yield effect from different studies in the literature. The "Barrows, Sexton, Zilberman" estimates are derived from our preferred yield estimates in Table 3.3 (column 1). "Barrows, Sexton, Zilberman (Log specification)" estimates are derived from the log specification in column 3 of Table 3.3.

margin assumption. The left column, labeled "Without Extensive Margin Effect," reports the resulting supply effects when we assume that extensive margin lands could have been profitably farmed with traditional seeds. The right column, labeled "With Extensive Margin Effect," reports supply effects after subtracting all production on extensive margin lands. Results are discussed in the main text.

C.2 Price Effects

The supply effect from GE technology can be translated into price effects using a methodology from De Gorter and Zilberman (1990) and Alston, Norton, Pardey, et al. (1995). Suppose that without GE technology, the supply curve shifts in by a factor of η, where η corresponds to the supply effect from the previous section. In the new equilibrium:

$$(1-\eta)\,Q_s\,(p) = Q_d\,(p) \tag{C.4}$$

where $Q_s\,(p)$ and $Q_d\,(p)$ represent quantities supplied and demanded, respectively, as a function of output price p. Totally differentiating with respect to η and p, yields

$$\begin{aligned}(1-\eta)\frac{\partial Q_s\,(p)}{\partial p}dp - Q_s d\eta \quad &= \quad \frac{\partial Q_d\,(p)}{\partial p}dp \\ &\Longrightarrow \quad dp\left[(1-\eta)\frac{\partial Q_s\,(p)}{\partial p} - \frac{\partial Q_d\,(p)}{\partial p}\right] = Q_s d\eta \\ &\Longrightarrow \quad \frac{dp}{p} = \frac{\partial \eta}{\epsilon_s - \epsilon_d}\end{aligned} \tag{C.5}$$

where the last line follows from setting $\eta = 0$. Equation (C.5) states that the percentage change in equilibrium price (the price effect) is equal to the supply effect divided by the difference between price elasticity of supply and price elasticity of demand. Thus, estimating the price effect simply involves scaling the supply effect from the previous section by elasticities parameters readily obtained from the literature. In our estimates, $\epsilon_s = 0.3$, a low elasticity scenario is parameterized with $\epsilon_d = -0.3$, and a high elasticity scenario uses $\epsilon_d = -0.5$.[1] For each elasticity scenario, we also vary the assumption on the extensive margin as before. For each of these 4 scenarios {low elasticity, no extensive margin ; low elasticity with extensive margin; high elasticity, no extensive margin; high elasticity, with extensive margin} price effects are computed conditional on yield estimates and plotted in Figure C.3 for corn and Figure C.4 for cotton. We discuss results in the main text.

C.3 Land-Use Saving Effects

Lastly, we estimate land-use saving effects and the corresponding GHG emissions savings due to GE technology. We compute saved hectares as the difference between observed hectarage

[1] Roberts and Schlenker (2010) suggest that supply elasticities vary between 0.08 and 0.13 for supply of grain calories and demand elasticities vary between -0.05 and -0.08. Thus, the magnitude of the price effect should be greater than five times the magnitude of the supply effect, which are greater than the impacts estimated here.

Figure C.3: Price Effect of GE Corn

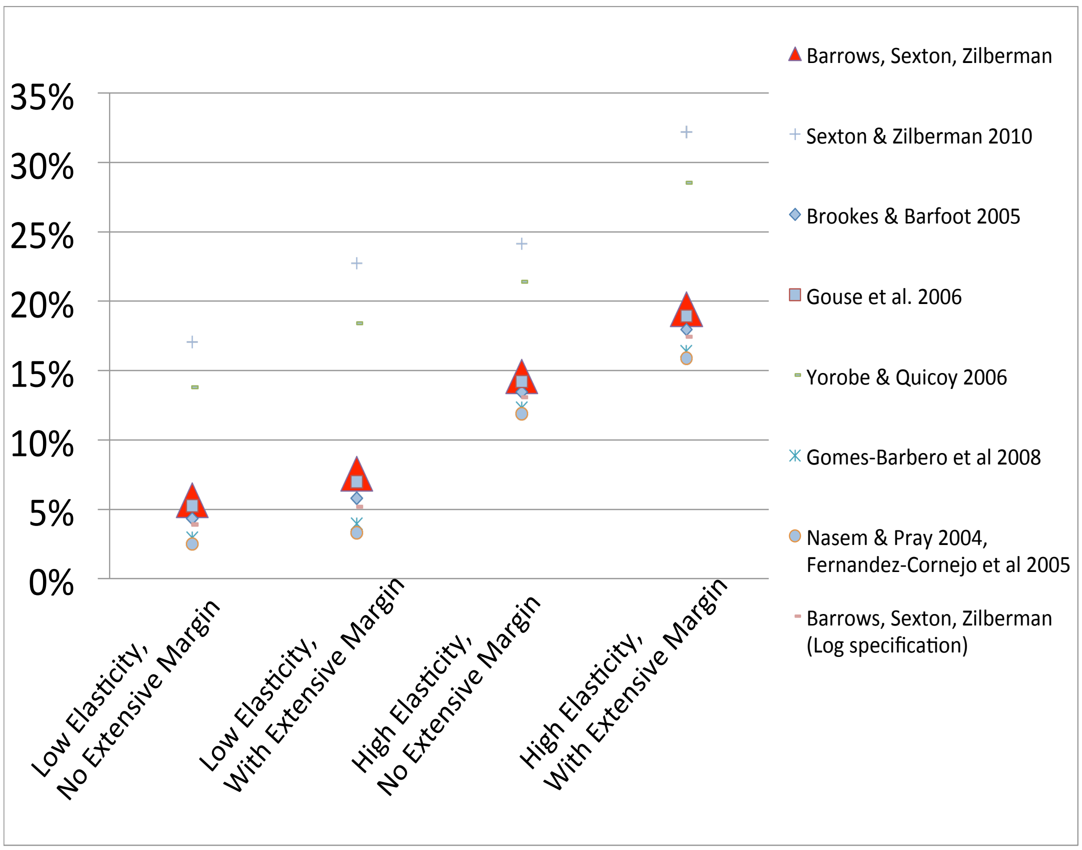

Notes: Price effect of GE corn calculated as percentage difference between observed price and counterfactual price without GE technology. Estimate without extensive margin effect allow that production on the extensive margin would have occurred with the traditional technology as well. Estimates with extensive margin effect subtract all production on extensive margin in the counterfactual supply. "Low elasticity" scenario sets elasticity of demand to -0.3, "high elasticity" scenario sets elasticity of demand to -0.5. Each point corresponds to estimates based on the yield effect from different studies in the literature. The "Barrows, Sexton, Zilberman" estimates are derived from our preferred yield estimates in Table 3.3 (column 4). "Barrows, Sexton, Zilberman (Log specification)" estimates are derived from the log specification in column 6 of Table 3.3.

Figure C.4: Price Effect of GE Cotton

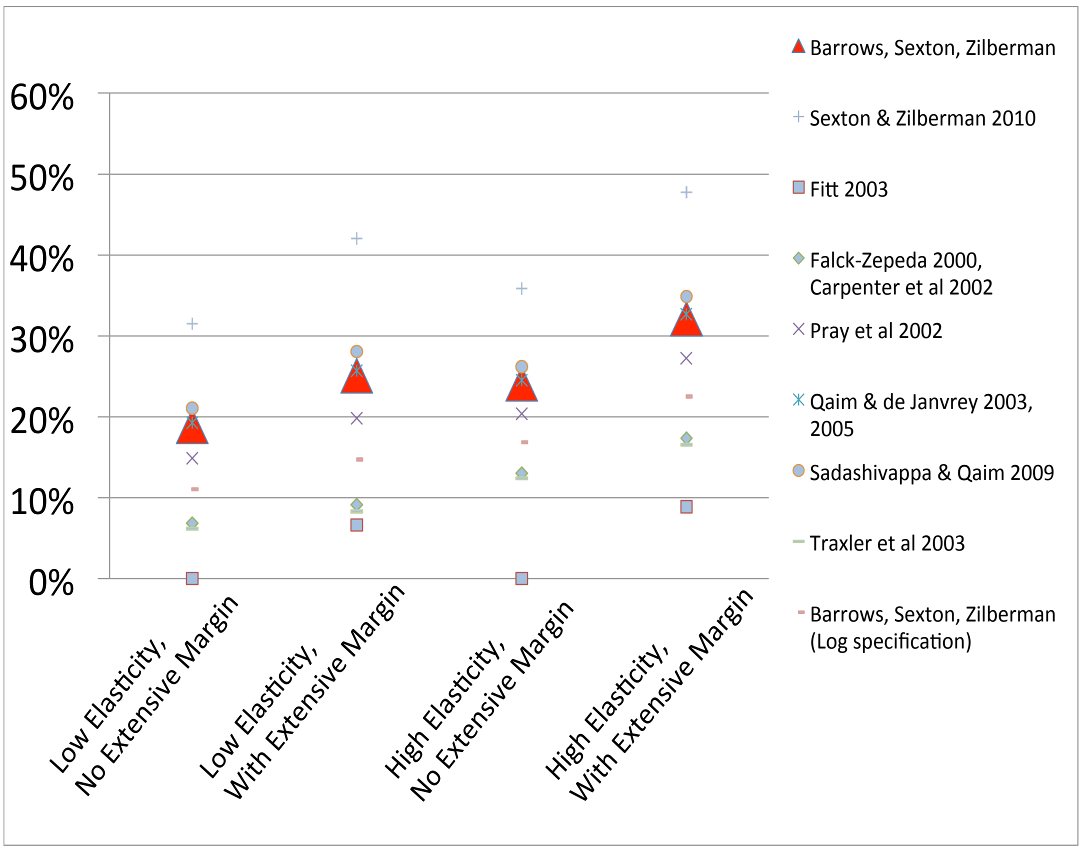

Notes: Price effect of GE cotton calculated as percentage difference between observed price and counterfactual price without GE technology. Estimate without extensive margin effect allow that production on the extensive margin would have occurred with the traditional technology as well. Estimates with extensive margin effect subtract all production on extensive margin in the counterfactual supply. "Low elasticity" scenario sets elasticity of demand to -0.3, "high elasticity" scenario sets elasticity of demand to -0.5. Each point corresponds to estimates based on the yield effect from different studies in the literature. The "Barrows, Sexton, Zilberman" estimates are derived from our preferred yield estimates in Table 3.3 (column 1). "Barrows, Sexton, Zilberman (Log specification)" estimates are derived from the log specification in column 3 of Table 3.3.

Table C.1: Land-Use Saving Effects in 2010

	(1) 2010 Harvested Area *(Millions of Ha)*	(2) Area Saved *(Millions of Ha)*	(3) GHG Saved *(Gt)*
Cotton	32	6	0.07
Corn	160	5	0.06
Soybeans	102	2	0.03
Total	294	13	0.15

Notes: 2010 Harvested Area are world aggregate from FAO Stat. "Area Saved" in column 2 represents the difference between observed area (column 1) and counterfactual area needed to meet observed 2010 demand without the intensive margin yield impact of GE. Column 3 multiplies "Area Saved" by a constant GHG/Ha/yr value of 11.7 metric tonnes, taken from the land-use literature (Searchinger et al., 2008).

in 2010 and counterfactual hectarage that would be needed to produce the same output without the GE supply effects.

Formally, counterfactual hectarage without considering the extensive margin effect is computed as

$$\widetilde{L_{ct}} = \frac{Q_{ct}}{\widehat{y_{ct0}}} \tag{C.6}$$

Country-specific hectarages are aggregated to the world level and observed 2010 hectarage is subtracted to compute world hectarage savings

$$\widetilde{L}_t = \sum_c \left(\widetilde{L_{ct}} - L_{ct}\right) \tag{C.7}$$

Estimates are reported by crop in the second column of Table C.1 and discussed in the main text. In the last column of Table C.1, we translate land-use savings into Gt of averted GHG emissions by multiplying the hectares saved by GHG emissions per hectare of land-use change per year.