

Improved Access to Foreign Markets Raises Plant-Level Productivity ... for Some Plants

Alla Lileeva*[‡]

York University

Daniel Trefler*[§]

University of Toronto and CIFAR

This version, July 4, 2007

ABSTRACT: We weigh into the debate about whether rising productivity is ever a consequence rather than a cause of exporting. Exporting and investing to raise productivity are complimentary activities. For lower-productivity firms, incurring the fixed costs of such investments is justifiable only if accompanied by the larger sales volumes that come with exporting. Lower foreign tariffs will induce these firms to simultaneously export and invest in productivity. In contrast, lower foreign tariffs will induce higher-productivity firms to export without investing, as in Melitz (2003). We model this econometrically using a heterogeneous response model. Unique 'plant-specific' tariff cuts serve as our instrument for the decision of Canadian plants to start exporting to the United States. We find that those lower-productivity Canadian plants that were induced by the tariff cuts to start exporting (*a*) increased their labour productivity, (*b*) engaged in more product innovation, and (*c*) had high adoption rates of advanced manufacturing technologies. These new exporters also increased their domestic (Canadian) market share at the expense of non-exporters, which suggests that the labour productivity gains reflect underlying gains in TFP. In contrast, we find no effects for higher-productivity plants, just as predicted by our complementarity theory.

Key words: international trade, exporting, productivity, innovation

JEL classification: F1

*We are grateful to Daron Acemoglu, Andy Bernard, Christian Broda, Xavier Gabaix, Bob Gibson, Elhanan Helpman, Sam Kortum, Kala Krishna, Marc Melitz, Diego Puga, John Sutton and Jim Tybout as well as seminar participants at Chicago, CIFAR, Harvard, MIT Sloan, NBER, Princeton, Statistics Canada and Toronto. We are particularly indebted to John Baldwin for his rich insights into the data. Trefler gratefully acknowledges funding from the Social Sciences and Humanities Research Council of Canada (SSHRC) and the tremendous support of the Canadian Institute for Advanced Research (CIFAR). Lileeva gratefully acknowledges funding from Statistics Canada.

[‡]Lileeva: Department of Economics, York University. Address: 1144 Vari Hall, 4700 Keele Street, Toronto, Ontario, M3J 1P3, Canada. (e-mail: lileeva@dept.econ.yorku.ca; website: <http://dept.econ.yorku.ca/~lileeva/>). Also affiliated with Statistics Canada.

[§]Rotman School of Management and Department of Economics, University of Toronto, 105 Saint George Street, Toronto, Ontario M5S 3E6, Canada (e-mail: dtrefler@rotman.utoronto.ca; website: <http://www.rotman.utoronto.ca/~dtrefler/>). Also affiliated with the NBER.

Does exporting raise productivity? The seminal contributions to the topic by Clerides, Lach and Tybout (1998) and Bernard and Jensen (1999) provide a clear 'no' to this question. The same conclusion appears in many subsequent contributions, including Bernard and Wagner (1997), Delgado, Fariñas and Ruano (2002) and Bernard and Jensen (2004). However, a large number of researchers have found varying degrees of support for a positive effect of exporting on productivity e.g., Aw, Chung and Roberts (2000), Baldwin and Gu (2003), Van Biesebroeck (2004), Lileeva (2004), Hallward-Driemeier, Iarossi and Sokoloff (2005), Fernandes and Isgut (2006), Park, Yang, Shi and Jiang (2006), Aw, Roberts and Winston (2007) and De Loecker (forthcoming). Any study examining whether starting to export raises productivity must confront two issues. First, there is an ironclad consensus in the literature that starting to export is endogenous: more productive plants choose to export. Second, if starting to export raises productivity, what are the mechanisms? It is hard to believe that there could be large productivity gains unless firms actively engaged in costly productivity-enhancing investments such as the adoption of advanced manufacturing technologies, the use of just-in-time production techniques, and product restructuring (the elimination of less successful products and the improvement of more successful products). When the decisions to start exporting and to invest in raising productivity are *both* endogenous, a set of econometric and policy issues emerge that have been ignored to date. These issues are similar to those raised in other contexts by Imbens and Angrist (1994), Card (2001) and Heckman and Vytlacil (2005).

The reason for linking the decisions to export and invest is that they are complimentary activities. When a firm does not export, the productivity gains from investing raise profits only on domestic sales. On the other hand, when a firm exports, the productivity gains from investing raise profits on both domestic and foreign sales. Thus, exporting raises the returns to investing in productivity. This complementarity appears in Ekholm and Midelfart (2005), Yeaple (2005), Bustos (2005), Atkeson and Burstein (2006), Costantini and Melitz (2007) and Ederington and McCalman (forthcoming) who provide conditions under which a reduction in the foreign tariff induces firms to simultaneously export and invest.

All of these papers allow for heterogeneity in initial productivity as in Melitz (2003). Productivity gains are then uniquely determined by the firm's initial productivity. (Atkeson and Burstein 2006 and Costantini and Melitz 2007 are exceptions.) In practice, however, we observe substantial heterogeneity in productivity gains even after conditioning on initial productivity, exporter status, and the decision to invest. Stories abound of firms that fail to implement new technologies as

successfully as their competitors — one need only think of Ford versus Toyota — and these stories are confirmed by econometric studies e.g., Aw et al. (2007, table 6). Once one allows for two sources of heterogeneity, in initial productivity and in productivity growth from investing, things quickly become complicated. In particular, if the productivity benefits of starting to export vary across firms then many of the parameters of interest for policy are not identified. For example, the impact of exporting on productivity (the average treatment effect) and the impact of exporting on productivity for those who export (the effect of treatment on the treated) are not identified. Imbens and Angrist (1994) show that if there is a valid instrument for exporting then what is identified is the average productivity gains from exporting *for firms that are induced to export because of the instrument*. In this study we will be able to identify the average productivity gains for Canadian firms that were induced to export to the United States because of the U.S. tariff cuts mandated by the Canada-U.S. Free Trade Agreement.

One way of thinking about the role of instruments is in terms of the very different conclusions drawn by Bernard and Jensen (1999) for the United States and De Loecker (forthcoming) for Slovenia. Slovenian firms likely started exporting because of improved access to the European Union and, as a pre-requisite to joining European Union supply chains, Slovenian firms likely invested heavily in reducing product defect rates and lowering costs. The implicit instrument — entry into the European Union — picks off new exporters that were investing. In contrast, most U.S. plants find themselves in a domestic market that is large enough to justify investing even without access to foreign markets. As Bernard and Jensen showed, plants in their U.S. sample likely started exporting because improved productivity from previous investing pushed them past the Melitz (2003) cut-off. These new exporters thus did not experience additional productivity gains from starting to export. The implicit instrument — past productivity growth — picks off new exporters that started investing before exporting. More generally, different firms have different degrees of complementarity between exporting and investing and thus have heterogeneous post-exporting investment strategies and productivity responses. Since different instruments yield different predictions about who exports, different instruments yield different results about the relationship between exporting, investing and productivity.

This observation about instruments is central to work by Imbens and Angrist (1994), Card (2001) and Heckman and Vytlačil (2005). It means that the policy conclusions drawn about the benefits of starting to export will depend on the choice of instrument. *It is thus surprising that not a single*

existing study has used a policy variable as an instrument for starting to export. We will use the U.S. tariff cuts mandated by the Canada-U.S. Free Trade Agreement as an instrument for the decision of Canadian plants to start exporting. The tariff cut is plant-specific. That is, we link the tariff-cut data to a plant's commodity data in order to compute the average tariff cut experienced by the plant.

We are not, of course, the first to use instruments. Clerides et al. (1998) and subsequent papers that adopted their pioneering methodology brilliantly use the dynamic structure of panel data to generate instruments. Clerides et al. instrument for exporting using exchange rates, plant age, plant business type and lagged values of capital. The literature spawned by the pioneering methodology of Bernard and Jensen (1999) rarely uses any instruments. Two exceptions are Van Biesebroeck (2004) and Park et al. (2006), but their instruments are also not policy variables.¹

This paper is related to a number of empirical studies that connect starting to export with investing. Hallward-Driemeier et al. (2005) and Alvarez and Lopez (2005) provide some weak evidence that their estimates of positive impacts of exporting on productivity are mediated by higher levels of investment and worker training. Bustos (2005) finds a correlation between changes in technology spending and starting to export for Argentinean plants. Aw et al. (2007) find that firms which both export and do R&D have higher subsequent productivity growth. They also find a positive cross-equation correlation between their export and R&D probits, which suggests that exporting and investing are simultaneous decisions. Baldwin, Beckstead and Caves (2002) and Baldwin and Gu (2004) find that exporters are relatively more specialized in the products they produce, invest more in R&D and training, and adopt more advanced manufacturing technologies. Feinberg and Keane (2006) and Keane and Feinberg (forthcoming) find that the 1983-96 increase in trade between U.S. multinationals and their Canadian affiliates was driven by technology adoption (the adoption of just-in-time techniques). Thus, there is evidence that various types of investment act as mediators between starting to export and rising productivity.

Our paper is also related to plant-level studies of the impact of the Canada-U.S. Free Trade Agreement. See Trefler (2004), Lileeva (2004), Baldwin and Gu (2003, 2006), Baldwin, Caves and Gu (2005), Baggs (2005) and Baggs and Brander (2006). Of particular interest, Baldwin and Gu (2006) find that exporters produce fewer products and have larger production runs. This may

¹Our work is best thought of as falling into the Bernard and Jensen branch of the literature rather than the Clerides et al. branch. We do not have the annual data on exporting needed to implement the Clerides et al. approach.

explain the positive productivity effects of exporting found by Baldwin and Gu (2003). See Gaston and Trefler (1997) and Head and Ries (1999, 2001) for industry-level analyses.

Finally, the literature spawned by Clerides et al. (1998) and Bernard and Jensen (1999) asks a bigger and more difficult question than the one posed here. It asks about the effect of exporting on productivity. We ask about the effect of exporting on productivity for those Canadian plants that were induced to export because of U.S. tariff cuts. Our results suggest that the larger question is in fact very difficult to answer. This is discussed in the conclusions where we argue that our results in no way contradict those of Clerides et al. and Bernard and Jensen.

The paper is structured as follows. In section 2 we describe our sample of 5,247 Canadian plants that were not exporters before the Free Trade Agreement (FTA) was implemented. Almost half of these began exporting after the FTA was implemented. In sections 3 and 5 we outline our empirical strategy for identifying and estimating the heterogeneous productivity effects of improved access to the U.S. market. The strategy is based on the Marginal Treatment Effect of Heckman and Vytlacil (1999, 2005). In sections 4 and 6 we provide estimates of the impact of the FTA tariff cuts on labour productivity growth and show that there is indeed heterogeneity. Lower productivity plants experienced large gains while higher productivity plants experienced no gains. In section 7 we show that the group of plants that experienced the large labour productivity gains were also the group of plants that invested most heavily in innovation and technology adoption. In section 8 we discuss the most significant of several weaknesses of our empirical work, namely, that we measure productivity by value added per worker rather than TFP. To partially address this, we show that the same plants that were induced to start exporting, to raise their labour productivity and to invest in new products and technologies were also the same plants that grabbed substantial *domestic* market share away from non-exporters. This suggests that these new exporters did indeed increase their TFP. Finally, a model of the complementarity between starting to export and investing in productivity is needed in order to interpret the form of heterogeneity that we estimate. This appears in the next section.

1. A Model of Selection into Investing and Exporting

Consider a model with two countries, home (Canada) and foreign (United States). Foreign values are denoted with an asterisk. Consumers have CES preferences and the market structure is monopolistic competition. A home firm producing variety i faces home demand $q(i) = p(i)^{-\sigma} A$ and

foreign demand $q^*(i) = p^*(i)^{-\sigma} A^*$ where $\sigma > 1$ is the elasticity of substitution between varieties, A is a measure of domestic market size, A^* is a measure of foreign market size, $p(i)$ is the price charged at home, and $p^*(i)$ is the price (inclusive of tariff) charged abroad. Let $\tau(i) - 1$ be the *ad valorem* tariff the firm faces when selling into the foreign market. Turning to costs, a standardized bundle of inputs costs c and produces $\varphi'_0(i)$ units of output. $\varphi'_0(i)$ measures productivity. However, it is easier to work with a transformation of productivity, namely, $\varphi_0 \equiv (\sigma - 1)^{\sigma-1} \sigma^{-\sigma} (\varphi'_0)^{\sigma-1}$. We are only interested in the firm's static optimization problem. We therefore treat the equilibrium outcomes A , A^* and $c = 1$ as exogenous parameters. In what follows we drop all i indices.

Consider the standard Melitz (2003) problem as described in Helpman (2006). For a fixed cost F^E the firm can export. Let $E = 1$ if the firm exports and $E = 0$ otherwise. Then the firm's maximum profits as a function of its exporting decision are

$$\pi_0(E) = \varphi_0 [A + E\tau^{-\sigma} A^*] - EF^E \quad (1)$$

for $E = 0,1$. See Helpman (2006, equations 1-2). It follows that the firm exports when φ_0 exceeds the Melitz cut-off $F^E / (\tau^{-\sigma} A^*)$.

In addition to an exporting decision, we assume that for a fixed cost F^I the firm can raise its productivity from φ_0 to φ_1 .² The firm's maximum profits when investing in productivity are

$$\pi_1(E) = \varphi_1 [A + E\tau^{-\sigma} A^*] - EF^E - F^I. \quad (2)$$


The firm's problem is most succinctly characterized by considering the difference between profits for (i) exporting and investing and (ii) neither exporting nor investing. From equations (1)-(2), this difference is

$$\pi_1(1) - \pi_0(0) = [\varphi_0 \tau^{-\sigma} A^* - F^E] + [(\varphi_1 - \varphi_0)A - F^I] + [(\varphi_1 - \varphi_0)\tau^{-\sigma} A^*]. \quad (3)$$

The first term in brackets equals the increase in profits from exporting without investing in productivity. The second term in brackets equals the increase in profits from investing in productivity without exporting. The third term captures the complementarity between investing and exporting – it is the increase in variable profits that results from both exporting and investing as opposed to doing just one or the other. It is necessarily positive because productivity gains raise profits on

²It makes no difference to our conclusions if there are only marginal costs of investing or both marginal and fixed costs of investing.

Figure 1. The Optimal Choices of Exporting and Investing



all units sold, including foreign sales, and hence raise the profits from exporting. This complementarity can also be thought of as a familiar market size effect that appears in many different models.

The firm's optimal choices are illustrated in figure 1 where initial productivity φ_0 is plotted against the productivity gains from investing $\varphi_1 - \varphi_0$. When productivity gains are small the firm never invests and we are in a Melitz world: the firm exports if and only if initial productivity is above the Melitz threshold. The Melitz threshold is the vertical line in figure 1. Given that the firm is exporting, it will invest if and only if the productivity gains are above some threshold.³ This threshold is the horizontal line in figure 1. For the empirical work to follow, the interesting region is where the first two terms in equation (3) are negative so that the firm will not export without investing and will not invest without exporting. In this region the complementarity between exporting and investing may nevertheless make it worthwhile for the firm to export and invest. To pin this down more precisely, suppose that in this region the firm must choose either (i) to export and invest or (ii) to do neither. The firm is indifferent between these two choices when $\pi_1(1) = \pi_0(0)$ or, from equation (3), when

$$\varphi_1 - \varphi_0 = -\varphi_0 \frac{\tau^{-\sigma} A^*}{A + \tau^{-\sigma} A^*} + \frac{F^I + F^E}{A + \tau^{-\sigma} A^*}. \quad (4)$$

Above this line the firm prefers to export and invest. Below it, the firm prefers to do neither.

³ $\pi_1(1) > \pi_0(1)$ iff $(\varphi_1 - \varphi_0)(A + \tau^{-\sigma} A^*) - F^I > 0$. Re-stated, $\pi_1(1) > \pi_0(1)$ iff $\varphi_1 - \varphi_0$ is above the threshold $F^I / (A + \tau^{-\sigma} A^*)$.

In figure 1 the firm never invests without exporting. Because investing without exporting is of no interest for the empirical work on exporting that follows, we have assumed implicitly that F^I is so large that the firm never invests without exporting. The appendix provides the analysis for the case where F^I is small. This leads to only a minor modification of the figure 1 analysis.⁴

We turn next to the effects of an improvement in access to the foreign market because of a fall in the foreign tariff τ . As shown in figure 2, the downward-sloping equation (4) rotates clockwise around its fixed vertical intercept. Thus, some firms that previously neither exported nor invested now find themselves choosing to both export and invest. For this group, the causal effect on productivity of improved market access is given by equation (4). The fall in τ also causes a leftward shift of the Melitz cut-off. See figure 2. Thus, some firms that previously neither exported nor invested now find themselves exporting without investing. For these firms improved market access has no causal effect on productivity. We will sometimes refer to the firms in the shaded regions as ‘switchers.’

The primary result of this section is summarized in figure 3. It shows that improved access to foreign markets raises productivity for some plants and not for others i.e., productivity responses are heterogeneous. This has important implications for empirical work. No researcher has ever adequately reported how productivity responses vary with initial productivity.⁵

A much less important result of this section is that the complementarity between exporting and investing leads to the particular form of heterogeneity shown in figure 3. *A priori* there is no reason to think that this will be a dominant effect in a richer model that allows for other factors and other sources of heterogeneity. For now we simply note that the form of heterogeneity displayed in figure 3 is what we find empirically. We also note that the empirical analysis to come imposes *none* of the theoretical structure developed in this section.⁶

Relationship to the Literature

⁴Specifically, define $\varphi_0 = F^E / (\tau^{-\sigma} A^*) - F^I / A$. For $\varphi_0 \geq \varphi_0$ the analysis presented in the main text is complete. In particular, the choice of investing without exporting is always dominated. For $\varphi_0 < \varphi_0$ there is a region of the parameter space for which investing without exporting is preferred. This region disappears when F^I is high enough that $\varphi_0 < 0$.

⁵Delgado et al. (2002) comes closest in estimating pre- and post-entry distribution functions of productivity growth separately for small and large firms. Hallward-Driemeier et al. (2005) and Park et al. (2006) imaginatively estimate effects for young and old firms.

⁶In figure 3, to the immediate left of the Melitz cut-off we have drawn a zero productivity response. In fact, the response is an average of the zero responses of those who start exporting without investing and the positive responses of those who invest when they start exporting. For ease of exposition and because the latter group is likely smaller, we have drawn the productivity effects as zero in this region.

Figure 2. Switching Behaviour Induced by Improved Foreign Market Access

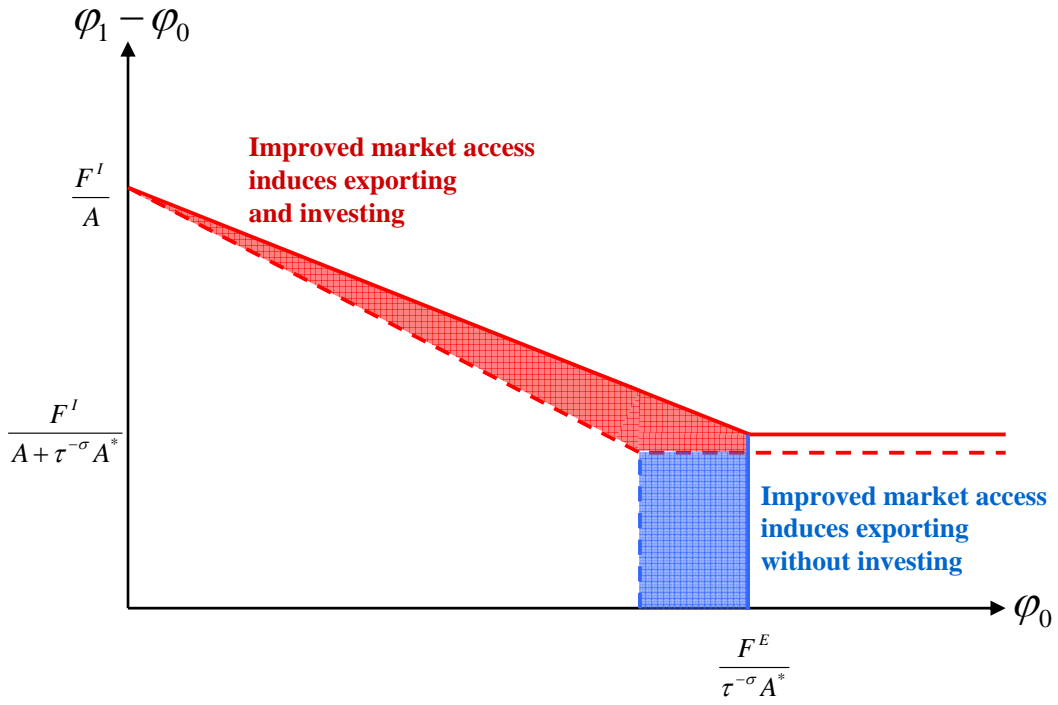



Figure 3. The Causal Effect on Productivity of Improved Foreign Market Access



Our paper is related to a growing literature on exporting and investing in productivity. See Ekholm and Midelfart (2005), Yeaple (2005), Bustos (2005) and Ederington and McCalman (forthcoming). These authors all assume that there is heterogeneity in initial productivity φ_0 . However, none of them allows for heterogeneity in productivity gains $\varphi_1 - \varphi_0$. We require heterogeneity in both φ_0 and $\varphi_1 - \varphi_0$. Implicitly, we have been assuming that there are firms scattered over all the regions in figures 1-3. If there are no firms in a given region then that region is irrelevant and the data will tell us this. In contrast, previous papers assume that there are firms in only a small part of figures 1-3, specifically, along some line $\varphi_1 - \varphi_0 = a + b\varphi_0$.⁷ The result obtained in these papers depends on where the line is assumed to lie in figures 1-3.⁸ The assumption of heterogeneity in φ_0 but only limited heterogeneity in $\varphi_1 - \varphi_0$ is helpful in allowing these authors to address a set of questions not considered in our paper. However, for our purposes, ruling out heterogeneity in $\varphi_1 - \varphi_0$ leads to a prediction that is inconsistent with the data. As we will see, for all but the smallest and largest φ_0 we observe exporters, non-exporters and switchers. This observation is inconsistent with the predictions of these earlier models, but is fully consistent with our model featuring two-dimensional heterogeneity.

2. The Canada-U.S. Free Trade Agreement and the Data

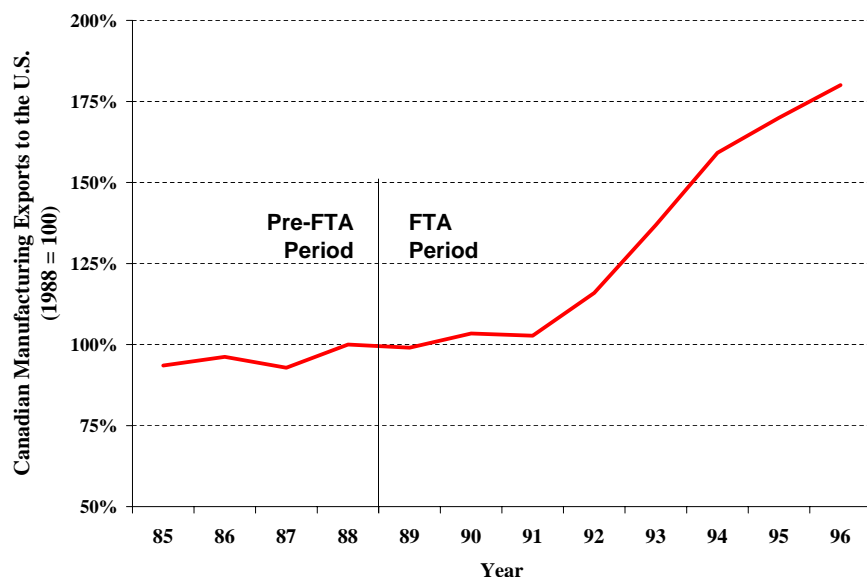
A. A Brief History of the FTA

We are interested in the effects of improved market access on firms' decisions to export and invest. We use the U.S. tariff reductions mandated under the Canada-U.S. Free Trade Agreement (FTA) to examine these effects. Negotiations for the FTA began in September, 1985. There was considerable uncertainty about whether there would be an Agreement until after the November 1988 general election brought the Conservatives back for a second term. The Agreement went into effect on January 1, 1989. By 1996, the last year for which we have plant-level data, each HS10 tariff was down to less than one-fifth of its 1988 level and by 1998 all tariffs were eliminated. See Brander (1991) and Thompson (1993) for details.

⁷For example, in Yeaple (2005) and Bustos (2005) when a firm invests it raises its productivity to $\varphi_1 = (b + 1)\varphi_0$ where $b > 0$. Thus, they only consider firms that are on the line $\varphi_1 - \varphi_0 = b\varphi_0$. In Ederington and McCalman (forthcoming), when a firm invests it raises its productivity to $\varphi_1 = a$ where $a > \varphi_0$. Thus, they only consider firms on the line $\varphi_1 - \varphi_0 = a - \varphi_0$.

⁸For concreteness, suppose that this line is upward-sloping and crosses the downward-sloping line in figure 3 at some point φ_0^* . Since all firms lie on this upward-sloping line, it follows that all firms with $\varphi_0 < \varphi_0^*$ are below our figure 3 line and hence neither export nor invest. All firms with $\varphi_0 > \varphi_0^*$ are above our line and hence export and invest. The only firms that are induced to export as a result of a marginal foreign tariff cut are firms with $\varphi_0 = \varphi_0^*$.

Figure 4. Canadian Manufacturing Exports to the United States



We have plant-level export status for 1984 and 1996. This means that we cannot examine the annual dynamics that are the focus of the literature spawned by the seminal papers of Roberts and Tybout (1997) and Clerides et al. (1998). This also means that we do not know the plant's first export date, information that is central in Bernard and Jensen (1999). Fortunately, there is abundant evidence that most exporting began only after implementation of the FTA and this will be enough information for our purposes. Figure 4 plots real Canadian manufacturing exports to the United States. Data are from Trefler (2004). These exports changed little in the 1985-88 period. They also changed little during the severe 1989-91 recession, the worst recession in Canadian manufacturing since the 1930s. However, exports climbed spectacularly after 1991, almost doubling in just five years. Romalis (forthcoming) shows a similar time profile for exports of goods that were subject to the largest tariff cuts. Feinberg and Keane (2005, 2006) use Bureau of Economic Analysis (BEA) data on shipments to the United States by U.S.-owned Canadian affiliates and find that there was no increase until 1988 and that most of the increase was from 1992 onwards. Thus, for plants that began exporting between 1984 and 1996, most likely started after implementation of the FTA and indeed, after 1991.

A second problem is that the 1984 export status data do not indicate the destination of exports. However, 83% of Canadian manufacturing exports in 1984 went to the United States and this number rose after implementation of the FTA. Thus, the vast majority of new entry into export

markets during the FTA period likely involved entry into the U.S. market.

In 1984 there were 5,417 plants that (a) did not export in 1984 and (b) survived until 1996. Of these, only 170 have missing data, leaving us with 5,247 plants. Criterion (b) is needed in order for us to observe productivity growth until 1996. In appendix B we include plants that did not survive until 1996 and find larger productivity estimates from improved market access. We define an exporter as a plant that exports *any* amount. Appendix G shows that our results are not sensitive to exporter definitions involving higher export thresholds.⁹

In addition to the 1984 and 1996 annual surveys of manufacturing, the 1979 and 1990 surveys were the only other surveys to ask the export question. The 1979 and 1990 data are used in appendix C. A very small number of the 1984 nonexporters exported in 1979. Our results are virtually unchanged when these 1979 exporters are excluded. Likewise, our results are virtually unchanged when we exclude plants that started exporting by 1990 but stopped exporting by 1996. Appendix C also includes additional analysis of export stoppers.

B. Description of the plant-specific tariff variable

We construct the FTA-mandated change in the plant-specific, commodity-weighted average U.S. tariff faced by Canadian plants. Let τ_{jt} be the U.S. tariff against Canadian imports of HS6 commodity j in year t . τ_{jt} is aggregated up from the underlying HS10 data using import weights. Let q_{ijt} be plant i 's sales of commodity j in year t . The FTA-mandated average tariff cut is $\sum_j (\tau_{j,1988} - \tau_{j,1996}) \omega_{ij}$ where $\omega_{ij} \equiv q_{ij,1996} / \sum_j q_{ij,1996}$. We use 1996 output weights in order to avoid the usual downward bias caused by contemporaneous weights i.e., the higher was the tariff in 1988 the lower were Canadian sales $q_{ij,1988}$. Since the first full year that the U.S. reported trade and tariffs in the HS classification was 1989, we use $\tau_{j,1989}$ in place of $\tau_{j,1988}$. Also, because tariffs were very close to 0 by 1996, we simply set $\tau_{j,1996} = 0$. This has the added advantage that it makes our measure of tariff changes invariant to compositional changes in the HS10 import weights used to aggregate up to HS6. Note that there will be additional issues with our tariff change measure that will be explained when we come to the empirical results. Also, the reader may have concerns about whether the tariff cuts are exogenous to the firm. We will provide evidence on this below.¹⁰

⁹The 1984 survey was administered to plants that accounted for a remarkable 91% of total manufacturing output.

¹⁰Tariffs are defined as duties divided by imports. Statutory rates would have been better, but we do not have access to an electronic file of statutory rates. Further, Trefler (2004) shows that results based on statutory rates are similar to those based on duties divided by imports.

Table 1. Average Plant Characteristics (Deviations from Industry Means)

	New Exporters $E = 1$	Non- Exporters $E = 0$	Difference $\mu_1 - \mu_0$ t	
N	2,133	3,114		
Log employment, 1984	-.36	-.82	.462	16.94
Log labour productivity, 1984	-.05	-.11	.064	4.15
Annual labour productivity growth, 1988-1996	.01	-.01	.019	7.86
Annual labour productivity growth, 1984-1988	.01	-.01	.026	6.11

Notes: This table reports the means for plants that did not export in 1984 and survived to 1996. Data are expressed as deviations from the industry mean where the industry consists of all plants in the same SIC4 industry that existed in both 1984 and 1996. It thus also includes continuous exporters.

C. Sample Moments

Table 1 reports some basic sample statistics. Of our 5,247 plants that did not export in 1984, 2,133 reported positive exports in 1996 ($E = 1$) and 3,114 reported zero exports in 1996 ($E = 0$). Table 1 provides additional sample statistics that are expressed as differences from the 4-digit SIC industry mean. The industry is defined as all plants that survived from 1984 to 1996. Subtracting off industry means partially controls for industry characteristics.

Table 1 makes it clear that new exporters and non-exporters were very different even before the FTA. Already in 1988, new exporters were only 0.36 log points smaller than the industry mean whereas nonexporters were 0.82 log points smaller. The difference of 0.462 is statistically significant ($t = 16.94$). As discussed in the introduction, we do not have capital stock or investment data and so use labour productivity (value added per worker). New exporters were 0.05 log points less productive than the industry mean, but 0.064 log points more productive than nonexporters ($t = 4.15$). Baldwin and Gu (2003) found that new Canadian exporters grew faster than nonexporters. In our sample, the growth differential was 0.019 log points per year in the 1988-96 period ($t = 7.86$). However, as one might expect from Bernard and Jensen (1999), the growth differential was similar in the 1984-88 period (0.026 log points, $t = 6.11$).¹¹

¹¹We started thinking about heterogeneous labour productivity responses because the double difference of labour productivity growth (1988-96 growth less 1984-88 growth for new exporters less nonexporters) is large for smaller, less productive plants and zero for larger, more productive plants. We will show this more formally below.

3. An Econometric Model

We next turn to specifying the simplest possible econometric model that allows us to identify and estimate heterogenous productivity responses to increased market access. Essentially, we regress 1988-96 productivity growth on a starting-to-export dummy and use the tariff as an instrument. We then use a much smaller sample of 521 plants matched to an innovation and technology survey to confirm that the group of plants for which productivity responses were largest is also the group of plants for which innovation and technology-adoption responses were largest.

Rather than working in productivity levels φ_0 and φ_1 we work in productivity changes so as to net out contaminating level fixed effects. Let $\Delta\varphi_1$ be the plant's productivity growth had it started exporting ($E = 1$). Let $\Delta\varphi_0$ be the plant's productivity growth had it never exported ($E = 0$). Our unit of analysis is the plant, but we suppress plant subscripts. As is standard, we decompose productivity growth into observable and unobservable components:

$$\Delta\varphi_0 = \beta_0(X) + U_0 \quad \text{for } E = 0 \quad (5)$$

$$\Delta\varphi_1 = \beta_0(X) + \beta_1(X) + U_0 + U_1 \quad \text{for } E = 1 \quad (6)$$

where $\beta_0(X)$ and $\beta_1(X)$ are components that vary with observables X . X will include initial productivity φ_0 . U_0 and U_1 are residual components. We assume that they are mean zero. We are interested in the causal effect of improved market access, via exporting, on productivity growth.

This causal effect is

$$\Delta\varphi_1 - \Delta\varphi_0 = \beta_1(X) + U_1 \quad (7)$$

The core econometric problem is unobserved heterogeneity: U_1 varies across plants. However, even observed heterogeneity $\beta_1(X)$ has not been adequately addressed in the literature.

For any one plant, if $E = 1$ we observe $\Delta\varphi_1$ and if $E = 0$ we observe $\Delta\varphi_0$. That is, observed productivity growth $\Delta\varphi$ is given by

$$\Delta\varphi = \Delta\varphi_1 E + \Delta\varphi_0(1 - E) = \beta_0(X) + \beta_1(X)E + (U_0 + U_1 E) \quad (8)$$

where we have used equations (5)-(6). This is our first estimating equation.

We model the exporting decision as a probit:

$$E = \begin{cases} 1 & P(X, \Delta\tau) \geq U_E \\ 0 & P(X, \Delta\tau) < U_E \end{cases} \quad \text{Export Probit.} \quad (9)$$

Table 2. Probit of the Probability of Starting to Export Between 1984 and 1996

Independent Variables	Coefficient	χ^2	p -value	Marginal Effect
$\Delta\tau$	0.89	265	0.000	.12
Log labour productivity, 1984	0.26	33	0.000	.05
Log employment, 1984	0.36	233	0.000	.12
Annual log labour prod. growth, 1984-88	0.89	30	0.000	.04
SIC4 fixed effects (208 industries)	Yes			

Notes: This table reports estimates of equation (9) for the 5,247 plants that did not export in 1984. $E = 1$ for the 2,133 plants that exported in 1996. $E = 0$ for the remaining 3,114 plants that did not export in 1996.

$P(X, \Delta\tau) \equiv \text{Prob}\{E = 1 | X, \Delta\tau\}$ is the probability of becoming an exporter given plant-level characteristics X and the U.S. tariff cut $\Delta\tau$.

Equations (8) and (9) form our econometric model. $\Delta\tau$ is the excluded variable in equation (8) and will serve as our instrument. Notice that none of our earlier theory is being imposed on the econometric model.

A. *Estimates of the Probit Model*

In 1984 there were 5,247 Canadian plants that were non-exporters. For each of these plants, let $E = 1$ if the plant was exporting in 1996 and let $E = 0$ if the plant was not exporting in 1996. Table 2 provides the results of a probit on E . In addition to the plant-specific tariff, the regressors include the log of labour productivity and employment in 1984, annual log labour productivity growth in 1984-88 and fixed effects for each of the 208 industries in Canada's 4-digit Standard Industrial Classification (SIC). Higher productivity plants, larger plants, and plants with rapidly growing productivity were all more likely to become new exporters. These are not surprising results. The new result is about $\Delta\tau$. As judged by the χ^2 -test statistic and the marginal effect, the U.S. tariff cut is at least as important as previously considered variables.

We must now provide more information about our tariff variable $\Delta\tau$. When estimating the probit using the tariff cut variable described in the data section above, we find that it is not statistically significant and has a very small marginal effect. This puzzled us at first until we realized that the distribution of the tariff cuts is very skewed – a few plants received very large tariff cuts that were often in excess of 50%. Such tariffs were likely well above the level needed to choke off imports, especially since the largest tariff cuts were in 'low-end' manufacturing where profit margins are often less than 10%. Thus, for tariffs that exceed a prohibitory threshold, tariff

variation is meaningless. As a result, when we use our continuous tariff-cut variable in the probit we find that it is not statistically significant. It turns out that the problem is easily solved by redefining $\Delta\tau$ to be a binary variable: $\Delta\tau = 1$ if the plant's tariff cut exceeds some threshold $\overline{\Delta\tau}$ and $\Delta\tau = 0$ otherwise. In table 2, the threshold $\overline{\Delta\tau}$ is the average tariff cut across all 5,247 plants. However, when we present our core results below, we will consider six other very different alternative definitions of the threshold $\overline{\Delta\tau}$ and show that our results are insensitive to the choice of threshold.¹²

Let $\widehat{P}(X, \Delta\tau)$ be the table 2 estimate of $P(X, \Delta\tau)$. It is common to estimate the effects of treatment either by matching treated and untreated units based on the 'propensity score' \widehat{P} or by using \widehat{P} as an instrument. This will be part of what we do. We therefore describe \widehat{P} in much more detail than is the reporting norm. This is done in table 3 where we group plants according to their predicted probabilities of entry \widehat{P} . The table reports sample statistics for each group. For example, the third row deals with those plants whose predicted probability of entry lies between 0.15 and 0.25. There are five points highlighted by the table.

First, the estimated probit does not appear to suffer any systematic mis-prediction. For example, in the third row, 148 plants began exporting (column 3), 668 did not start exporting (column 4) and the rate of entry was $0.18 = 148 / (148 + 668)$ (column 2). Since the actual entry rate of 0.18 lies in the interval of predicted entry rates (0.15, 0.25) we conclude that the probit did a good job of prediction. Comparing columns 1 and 2, the actual entry rate is in the predicted interval for 10 of 11 groups. Based on this we conclude that the probit does not have any systematic biases in predicting entry.

Second, some of the groups are very thin on either new exporters or non-exporters. Above the upper horizontal line in table 3 there are very few new exporters and below the lower horizontal line there are very few non-exporters. This makes it difficult to reliably estimate differences between new exporters and non-exporters. Results for these low and high values of \widehat{P} should therefore be treated with caution.

Third, \widehat{P} is highly correlated with initial labour productivity and employment size. Columns 5 and 6 show initial labour productivity and employment. The baseline is plants in the first row. For

¹²There is something uncomfortable about converting a continuous variable into a binary one. However, it is important here to keep the aim in focus. We are not interested in accurately estimating the elasticity of export participation with respect to tariff cuts. Rather, we are interested in an instrument that has decent explanatory power in the first stage i.e., in the probit.

Table 3. Features of the Probit $P(X, \Delta\tau)$

	Predicted Prob. of Entry $P(X, \Delta\tau)$	Entry Rate	No. of Exporters ($E = 1$)	No. of Nonexporters ($E = 0$)	Average Log Productivity	Average Log Employment	Average $\Delta\tau$	Marginal Effect of $\Delta\tau$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	.00-.05	.06	15	243	0.00	0.00	0.34	0.05
2	.05-.15	.08	51	596	0.07	0.44	0.49	0.15
3	.15-.25	.18	148	668	0.13	0.50	0.45	0.25
4	.25-.35	.29	206	506	0.16	0.71	0.48	0.31
5	.35-.45	.43	274	369	0.18	1.05	0.52	0.34
6	.45-.55	.50	293	289	0.20	1.23	0.49	0.36
7	.55-.65	.62	335	204	0.20	1.41	0.44	0.34
8	.65-.75	.72	339	131	0.22	1.60	0.39	0.31
9	.75-.85	.77	289	88	0.24	2.09	0.23	0.25
10	.85-.95	.87	130	19	0.45	2.35	0.19	0.16
11	.95-1.00	.98	53	1	0.35	2.41	0.62	0.02

Notes: This table reports characteristics of plants by $\hat{P}(X, \Delta\tau)$. Column 1 indicates the plant type. For example, the first row lists statistics for all plants with $\hat{P}(X, \Delta\tau) \in (0.00, 0.05)$. Column 2 is the proportion of plants that started exporting between 1984 and 1996. It equals column 3 divided by the sum of columns 3 and 4. In column 5, log productivity in 1984 is the difference between log productivity in the indicated row minus log productivity in row 1. Similarly for column 6 log employment in 1984. Column 7 is the average of $\Delta\tau$. Column 8 is the average marginal effect from the probit.

example, plants with $\hat{P} \in (0.15, 0.25)$ had 1984 labour productivity that was 0.13 log points higher than plants with $\hat{P} \in (0.00, 0.05)$ and 1984 employment that was 0.50 log points higher. Looking down the rows we see that initial labour productivity and employment size rise sharply with \hat{P} . Indeed, these two variables are the primary drivers of between-row variation in \hat{P} . Thus, \hat{P} is acting as an index of labour productivity and employment size. Ricardo's logic implies that more productive plants will be larger and hence employs more workers. We therefore sometimes treat \hat{P} as an aggregator of two correlates of underlying productivity — labour productivity and size — where the weights are chosen to best predict the impact of productivity on exporting. That is, we will sometimes interpret \hat{P} as a proxy for our theory's φ_0 .

Fourth, while the between-row variation in \hat{P} is driven by initial labour productivity and size, the within-row variation is driven by $\Delta\tau$. This can be seen from column 8 which shows the estimated marginal effects of $\Delta\tau$. These effects are very large. For example, in row 3 the marginal

effect more than doubles the probability of entry from 0.18 to $0.43 = 0.18 + 0.25$.

Fifth, in the Rosenbaum and Rubin (1983) propensity score matching view of \hat{P} , \hat{P} is a sufficient statistic for all information about plant characteristics that is relevant to the exporting decision i.e., the distribution of X conditional on \hat{P} is independent of E . This is the testable ‘balancing’ hypothesis. To examine it, for each of our 11 \hat{P} rows we tested whether the mean of 1984 log labour productivity is the same for $E = 0$ and $E = 1$ plants. There was no difference in any of the 11 rows. We repeated this for 1984 log employment and 1984-88 log labour productivity growth and again always rejected the hypothesis of a difference in X between exporters and non-exporters. See appendix table 7. This means that once we control for \hat{P} there are no differences between new exporters and nonexporters.¹³

4. Preliminary Evidence on the Heterogeneity of Productivity Responses

Since heterogeneous response models with unobserved heterogeneity can be complex econometrically, we start with the simpler case of observed heterogeneity only. With no unobserved heterogeneity, $U_1 = 0$ and equation (8) becomes $\Delta\varphi = \beta_0(X) + \beta_1(X)E + U_0$. E is obviously endogenous and we instrument it with $\Delta\tau$. Recall that $\Delta\varphi$ is annual log labour productivity growth averaged over the 1988-96 period.

Table 4 reports the IV estimates of

$$\Delta\varphi = \beta_{0q}X + \beta_{1q}E + U_0 \quad q = 1, \dots, 5 \quad (10)$$

separately for each of the five quintiles of $\hat{P}(X, \Delta\tau = 0)$. q indexes quintiles. \hat{P} is the probit of table 2 and X collects the regressors in that probit (including 4-digit SIC fixed effects). The first column of table 4 shows the range of \hat{P} in each quintile. The second column shows the number of plants. The third column presents the IV estimates of the β_{1q} when $\Delta\tau$ is the instrument. We do not want to take these results seriously just yet. Instead, we wish to use them to make three points.

First, we do not interact E with X . Obviously we do not want to interact E with the fixed effects, but it is of interest that the interactions of E with employment, productivity and productivity growth are all statistically insignificant. This shows that once we condition on \hat{P} no useful

¹³Table 3 provides some informal and preliminary evidence on the exogeneity of the tariff cuts. Looking down column 7, $\Delta\tau$ is not particularly correlated with \hat{P} . (Within rows, $\Delta\tau$ and \hat{P} are of course highly correlated.) That is, $\Delta\tau$ is uncorrelated with observable plant characteristics that are correlated with exporting. This is very informal evidence on the exogeneity of $\Delta\tau$. We will provide more formal evidence on the exogeneity of $\Delta\tau$ below.

Table 4. Preliminary IV Estimation of Labour Productivity Growth

$P(X, \Delta\tau)$	N	IV: $\Delta\tau$		IV: $P(X, \Delta\tau)$		OLS		Std. Dev. of $\Delta\varphi$
		β_{1q}	t	β_{1q}	t	β_{1q}	t	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
0.00-0.22	1,049	0.25	6.19	0.22	6.51	0.041	5.34	0.10
0.22-0.34	1,050	0.20	6.18	0.20	6.21	0.030	5.51	0.09
0.34-0.45	1,049	0.13	6.41	0.13	6.40	0.039	7.53	0.09
0.45-0.61	1,050	0.05	2.45	0.05	2.44	0.030	5.54	0.09
0.61-1.00	1,049	0.08	2.00	0.07	1.95	0.022	4.07	0.08

Notes: This table reports estimates of β_{1q} in equation (10). The dependent variable is average annual log labour productivity growth, 1988-96. The equation is estimated separately for each quintile of the distribution of $\hat{P}(X, \Delta\tau = 0)$. Column 1 gives the range of the quintile. Column 2 gives the number of observations. Column 3 gives the IV results using $\Delta\tau$ as the instrument. Column 5 gives the IV results using $\hat{P}(X, \Delta\tau)$ as the instrument. Column 7 gives the OLS results. Column 9 gives the standard deviation of the dependent variable.

correlation between E and X remains. This is a restatement of our balancing tests in appendix D. It provides additional support for our Rosenbaum and Rubin (1983) propensity score structure and thus for our interpretation of \hat{P} as a proxy for initial productivity φ_0 .

Second, the hypothesis of homogeneity across quintiles is easily rejected: there is response heterogeneity that is correlated with observables i.e., with \hat{P} . Interestingly, the estimated β_{1q} are decreasing over the first four quintiles. Interpreting \hat{P} as initial productivity, this pattern is consistent with our figure 3 prediction. The last quintile is not consistent with our prediction.

Third, the estimates are huge. The first quintile coefficient of 0.25 means that exporting raises labour productivity by 25% a year in each of 8 years. This is much larger than the 0.10 within-quintile standard deviation of annual productivity growth shown in column 9. More importantly, it is also much larger than the OLS estimate shown in column 7. The two other papers in this line of research that use IV approaches also find IV estimates that are much larger than OLS estimates. See Van Biesebroeck (2004) and Park et al. (2006).¹⁴

An IV estimate that is larger than its OLS counterpart is usually taken as evidence of unob-

¹⁴We are referring only to papers in the Bernard and Jensen branch of the literature. See footnote 1 above.

served response heterogeneity e.g., Card (2001). To understand why, assume for simplicity that equation (10) has unobserved heterogeneity and no covariates: $\Delta\varphi = \beta_{0q} + \beta_{1q}E + U_0 + U_1E$. Let $\Delta P \equiv P(X,1) - P(X,0)$ be the amount by which the tariff cut increases the probability of starting to export. Then $\text{plim}\beta_{1q}^{IV} = \beta_{1q} + \text{plim}\mathbf{E}[U_1 \cdot \Delta P] / \mathbf{E}[\Delta P]$ i.e., the IV estimator is the weighted average of the productivity effects where the weights are ΔP .¹⁵ The complementarity between exporting and investing in productivity would suggest that U_1 should be positively correlated with ΔP : a firm that expects a larger productivity effect (U_1 large) gains a lot from exporting and therefore would be sensitive to a tariff cut (ΔP large). If this complementarity effect is very strong then $\text{plim}\beta_{1q}^{IV}$ will be large, just as in table 4. To deal with this problem we will move beyond quintiles to finer gradations of P in the next section.

We make one last point that will help for what comes in the next section. We have used $\Delta\tau$ as the instrument. Econometrically, any monotonic function of $\Delta\tau$ can be used as an instrument and $\widehat{P}(X,\Delta\tau)$ is a typical choice e.g., Imbens and Angrist (1994). If \widehat{P} were linear in its arguments then $\Delta\tau$ and \widehat{P} would yield identical IV estimates. However, the nonlinearity means that $\Delta\tau$ and \widehat{P} are not perfectly collinear. Hence, \widehat{P} is a second valid instrument. This use of the functional form of the probit to create a second instrument is not a desirable property of \widehat{P} as an instrument. Fortunately, table 4 shows that it is also not an important property empirically. The IV estimates of β_{1q} using \widehat{P} as an instrument (column 5) are almost identical to the IV estimates of β_{1q} using $\Delta\tau$ as an instrument (column 3). Thus, we can use either $\Delta\tau$ or \widehat{P} as an instrument.

5. The Marginal Treatment Effect

Consider a plant whose manager is unusually good in two senses. First, she jumps on opportunities in the U.S. market as they present themselves. In terms of our model, recall that a plant exports when $P > U_E$. Thus our good manager has a small U_E . Second, she squeezes large productivity gains out of a given investment i.e., she has a large U_1 . Together, these imply that we should expect a negative correlation between U_E and U_1 and hence a negative correlation between U_E and the productivity effect of improved market access $\Delta\varphi_1 - \Delta\varphi_0 = \beta_1 + U_1$. More formally, we should expect $\mathbf{E}[\beta_1 + U_1 | U_E]$ to be decreasing in U_E .

¹⁵See Card (2001, page 1142). To hint at why this is the case we adopt the following notation. For any variable y let \bar{y} be the mean of y , let y^1 be the mean of y for those plants with $\Delta\tau = 1$ and let y^0 be the mean of y for those plants with $\Delta\tau = 0$. Then the IV estimator is $\Sigma[(\Delta\varphi - \overline{\Delta\varphi}) \cdot \Delta\tau] / \Sigma[(E - \bar{E}) \cdot \Delta\tau] = [\Delta\varphi^1 - \Delta\varphi^0] / [E^1 - E^0] \approx [\Delta\varphi^1 - \Delta\varphi^0] / [\widehat{P}(1) - \widehat{P}(0)]$. The last expression is the *LATE* estimator. ΔP is the population counterpart to the denominator. The discussion following equation (11) below links the population moment $\mathbf{E}[U_1\Delta P]$ to the numerator.

Of course U_E is not observed so it is useless to condition on it. However, for a plant manager that is just indifferent between exporting and not exporting it must be that $U_E = P(X, \Delta\tau)$. See equation (9). Restated, we know U_E for plants that are induced to export because of the U.S. tariff cuts. Thus, we are interested in $\mathbf{E}[\beta_1 + U_1|P(X, \Delta\tau)]$. This expectation is called the Marginal Treatment Effect (*MTE*) and is the effect of exporting on productivity for those plants that were induced to export because of the U.S. tariff cuts. The Marginal Treatment Effect is due to Heckman and Vytlacil (2005). See also Heckman and Vytlacil (1999) and Carneiro, Heckman and Vytlacil (2003).

In terms of figure 1, $P(X, \Delta\tau) = U_E$ corresponds to the lines which demarcate no exporting from exporting, including the downward-sloping ‘complementarity’ line. The fact that the Marginal Treatment Effect can be linked directly to our theory is the primary reason that we burdened the reader with any theory at all.

To see that the Marginal Treatment Effect is identified, consider the equation (8) expression $\Delta\varphi = \beta_0 + \beta_1 E + (U_0 + U_1 E)$. Since P is an instrument, $\mathbf{E}[U_0|P] = 0$. In addition, the average value of E given P is just the probability of E : $\mathbf{E}[E|P] = P$.¹⁶ Hence,

$$\mathbf{E}[\Delta\varphi|P] = \beta_0 + \beta_1 P + \mathbf{E}[U_1 E|P]. \quad (11)$$

Consider a plant that has never exported and which lies on our line $P = U_E$. For such a plant, the tariff cut has two effects. First, it raises P which raises $\mathbf{E}[\Delta\varphi|P]$ by β_1 . Second, it induces the plant to switch from $E = 0$ to $E = 1$. This raises $\mathbf{E}[\Delta\varphi|P]$ from $\mathbf{E}[U_1 \cdot 0|P] = 0$ to $\mathbf{E}[U_1 \cdot 1|P] = \mathbf{E}[U_1|P]$. Combining these two effects, an increase in P raises $\mathbf{E}[\Delta\varphi|P]$ by the Marginal Treatment Effect $\mathbf{E}[\beta_1 + U_1|P]$. Restated, the Marginal Treatment Effect is identified by the derivative of $\mathbf{E}[\Delta\varphi|P]$:

$$\mathbf{E}[\beta_1 + U_1|P = U_E] = \left. \frac{\partial \mathbf{E}[\Delta\varphi|P]}{\partial P} \right|_{P=U_E}. \quad (12)$$

Since we observe $\Delta\varphi$ and can estimate P , we can estimate $\mathbf{E}[\Delta\varphi|P]$ and then differentiate it to obtain an estimate of $\mathbf{E}[\beta_1 + U_1|P]$. This establishes identification of the Marginal Treatment Effect. See Heckman and Vytlacil (1999) for a formal proof.

Turning to estimation, we follow the three-step procedure of Carneiro et al. (2003). First, we estimate equation (11) using a fully non-parametric IV procedure with \hat{P} as the instrument. This is the non-parametric equivalent of what appears in column 5 of table 4. As in that table we continue

¹⁶ $\mathbf{E}[E|P] = \mathbf{E}[1|P] \cdot P + \mathbf{E}[0|P] \cdot (1 - P) = 1 \cdot P + 0 \cdot (1 - P) = P$.

to assume $\beta_0(X) = \beta_0 X$ and $\beta_1(X)P = \beta_1 P$.¹⁷ The procedure returns estimates of β_0 , β_1 and the residuals i.e., of $\hat{\beta}_0$, $\hat{\beta}_1$ and $\hat{\varepsilon}$, respectively. Second, we non-parametrically regress $\hat{\varepsilon}$ on \hat{P} . The resulting function $\hat{\varepsilon}_{U_1E}(\hat{P})$ is an estimate of $\mathbf{E}[U_1E|\hat{P}(X,\Delta\tau)]$. Together these two steps provide an estimate of $\mathbf{E}[\Delta\varphi|\hat{P}]$ i.e., of $\hat{\beta}_0 X + \hat{\beta}_1 \hat{P} + \hat{\varepsilon}_{U_1E}(\hat{P})$. Third, we numerically differentiate this estimate to obtain an estimate of the Marginal Treatment Effect. See appendix E for details.

6. Estimates of the Marginal Treatment Effect

Figure 5 reports our baseline estimate of the Marginal Treatment Effect i.e., of the impact of starting to export on productivity for those plants that were induced to export because of the U.S. tariff cuts. The horizontal axis is $\hat{P}(X,\Delta\tau)$.¹⁸ The left-hand vertical axis is the average annual log point change in labour productivity and is comparable to the parameter estimates of β_{1q} in table 4. The right-hand axis transforms this into the percentage change in labour productivity over the eight years 1988-96.¹⁹ For example, plants with $\hat{P} = 0.35$ that started exporting because of improved market access are estimated to have had labour productivity growth that was 40% higher than nonexporters by 1996. On the other hand, plants with $\hat{P} \geq 0.6$ are estimated to have had virtually 0 labour productivity growth. The dashed lines are 95% confidence intervals. These are calculated by bootstrapping the entire estimation procedure, including the probit stage, using 1,000 draws. The only statistically significant labour productivity effects are for $\hat{P} \leq 0.6$. Overall, the labour productivity impacts for plants with $\hat{P} \leq 0.6$ are both economically large and statistically significant.²⁰

Interpreting \hat{P} as a measure of initial productivity φ_0 , we find it of considerable interest that the profile in figure 5 is so similar to what is predicted by the theory in figure 3.

These results are very different from the OLS and IV results in table 4. If figure 5 had included the OLS results from table 4, they would appear as a curve starting at 0.041 and ending at 0.022. OLS thus underestimates the low- \hat{P} gains and over-estimates the high- \hat{P} gains. If figure 5 had


¹⁷This restriction that $\beta_1(X)$ be independent of X is tested and accepted (not rejected) in appendix E.

¹⁸To understand the labels on the horizontal axis, note that after estimating equation (11) we grouped plants into 11 groups based on their $\hat{P}(X,\Delta\tau)$: (.00,.05), (.05,.15), (.15,.25), ..., (.95,1.00). We then differentiated by differencing neighbouring groups. Thus, after differentiating we have 10 cells and the horizontal axis of figure 5 labels the boundaries between these cells i.e., .05, .15, .25 etc.

¹⁹If $\Delta\varphi \equiv (\ln \varphi_{1996} - \ln \varphi_{1988})/8$ is the value on the left-hand axis then $100 \cdot [e^{8 \cdot \Delta\varphi} - 1] = 100 \cdot (\varphi_{1996} - \varphi_{1988})/\varphi_{1988}$ is the value on the right-hand axis.

²⁰We have been reporting the effect of exporting on productivity. If one is interested in the effect of the tariff cut on productivity then one must multiply the Marginal Treatment Effect by the induced probability of exporting because of the tariff cut. That is, one must multiply figure 5 by the marginal effect of $\Delta\tau$ reported in column 8 of table 3.

Figure 5. The Marginal Treatment Effect: Baseline Specification



Notes: This figure provides estimates of the Marginal Treatment Effect i.e., of the productivity gains over the 1988-96 period for those plants that were induced to start exporting as a result of the U.S. tariff cuts. The dashed lines are 95% confidence intervals.

included the table 4 IV results, the shape of the profile would be similar to the Marginal Treatment Effect profile except much higher and, for $\hat{P} < 0.6$, much steeper.

Heterogeneity means that the Marginal Treatment Effect varies across plants i.e., it means that the figure 5 profile has a non-zero slope. The confidence intervals in figure 5 suggest that the slope is significantly different from 0. Appendix F provides a parametric slope test which confirms this insight ($p = 0.0000$).

Appendix F also provides a weak parametric test of overidentification which exploits the fact that $\hat{P}(X, \Delta\tau)$ is not linear in $\Delta\tau$. The test shows that $\Delta\tau$ is statistically insignificant when included

directly into the productivity growth equation. Its coefficient has a t -statistic of 0.41. Thus, the tariff cuts do not directly affect productivity growth. This is useful evidence supporting the exogeneity of the U.S. tariff cuts.

A. Sensitivity to Specification of Tariff Cuts $\Delta\tau$ and Exporter Status


Recall that a plant has $\Delta\tau = 1$ if its plant-specific tariff cut exceeds a threshold $\overline{\Delta\tau}$ which is the average tariff cut across all 5,247 plants. Our results are not sensitive to the choice of threshold. Figure 11 in appendix G plots the Marginal Treatment Effect for three alternative thresholds: (1) 50% below the average across all plants, (2) 50% above the average across all plants, and (3) $\overline{\Delta\tau}$ equal to the median tariff cut across all plants. There are no statistically significant differences between these and the baseline results of figure 5.

A threshold that is the same for all industries will tend to have more variation across industries than within industries: high-tariff industries will have many plants above the threshold and low-tariff industries will have many plants below the threshold. An alternative that brings out the within-industry variation in tariffs is a threshold $\overline{\Delta\tau}$ that is the average tariff cut for plants *within* the industry. The left panel of figure 6 displays the results for this case along with the baseline results from figure 5. The two results are similar, though the within results are somewhat smaller. A problem with this ‘within’ approach is that all industries will have plants with $\Delta\tau = 1$, even industries that had very low tariffs before 1988. To eliminate this problem we re-estimated the model using only those industries that had deep tariff cuts. The middle panel of figure 6 plots the Marginal Treatment Effect for the one third of industries that experienced the deepest tariff cuts (70 industries, $70 \approx 209/3$). The right panel plots it for the one half of industries that experienced the deepest tariff cuts (105 industries, $105 \approx 209/3$). These results are very similar to the baseline results. This suggests that it is within-industry variation in high-tariff industries that is driving the results.

7. Investing in Productivity

We have now accomplished the first of two major goals of this paper: we have shown that there were indeed labour productivity gains for low- and medium-productivity plants ($\hat{P} < 0.6$) that were induced to export as a result of improved access to U.S. markets. Our second goal is to link these labour productivity gains to active investments in productivity. We will show in this section

Figure 6. Marginal Treatment Effect: Within-Industry Tariff Threshold



Notes: This figure provides estimates of the Marginal Treatment Effect using a within-industry definition of the U.S. tariff cut. A plant has $\Delta\tau = 1$ if its tariff cut was deeper than the average tariff cut in the industry. The left panel reports the Marginal Treatment Effect when estimated using all industries. The middle (right) panel reports the Marginal Treatment Effect when estimated using the one third (one half) of industries that experienced the deepest U.S. tariff cuts. $P(X, \Delta\tau)$ is the estimated probability of exporting.

that the same plants that benefited from being induced to export — plants with $\hat{P} < 0.6$ — were also the plants that invested in product innovation and the adoption of advanced manufacturing technologies. This is a long paper so it is perhaps useful at this point to flag the importance of this section.

Data are from the 1993 Survey of Innovation and Advanced Technologies (SIAT). The surveyed plants include 521 plants that are in our group of 5,247 plants. The two-part survey deals with (a) innovation and (b) the adoption of advanced manufacturing technologies. See Baldwin and Hanel (2003) for a description of the survey. We start with the technology-adoption questions. The survey asks plants about their current use of various types of technologies and year of initial adoption. The most important of these is manufacturing information systems (MIS) which deals with computer-based production management and scheduling systems for orders, inventory and finished goods. MIS also deals with computer-based management of machine loading, production scheduling, inventory control and material handling. These systems are necessary for a variety of productivity-enhancing production techniques such as just-in-time inventory and lean manufacturing. Investments in MIS are thus a central component of any productivity-enhancing change in production techniques.

The first set of results in table 5 reports data on MIS adoption rates over the 1989-93 period. We start at 1989 because the Agreement came into effect on January 1, 1989. With only 521 plants we cannot use the data-intensive non-parametric approaches used above. We start simply with summaries of the raw adoption rates. We stratified the sample into two groups of plants, those with $\hat{P} < 0.6$ and those with $\hat{P} > 0.6$. \hat{P} is from the table 2 probit. Within each of these two groups, table 5 compares the adoption rates of new exporters and non-exporters. Among low- \hat{P} plants, 23% of new exporters adopted MIS between 1989 and 1993 whereas only 7% of non-exporters had done so. Thus, new exporters were 215% ($= (23 - 7)/7$) more likely than non-exporters to have adopted at least one advanced manufacturing technology by 1993. Among high- \hat{P} plants, 22% of new exporters had adopted at least one technology by 1993 and an almost identical 20% of non-exporters had done so. Thus, among high- \hat{P} plants new exporters were only 11% more likely than non-exporters to have adopted MIS. This means that among the group of plants where we found higher productivity gains for new exporters than non-exporters (plants with $\hat{P} < 0.6$) new exporters were adopting advanced technologies more frequently than non-exporters. In contrast, among the group of plants where we did *not* find differential productivity gains (plants with $\hat{P} >$

Table 5. Post-Agreement Technology Adoption and Product Innovation

$P(X, \Delta\tau)$	Without Plant Controls			With Plant Controls				N
	New	Non-	%	Predicted %		Double Difference		
	Exporters	Exporters	Difference	δ	p -value	δ_p	p -value	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1. Manufacturing Information Systems, Adopted in 1989-93								
$P < 0.6$	0.23	0.07	215%	0.88	0.00	1.06	0.01	309
$P > 0.6$	0.22	0.20	11%	-0.17	0.63			135
2. Inspection and Communications, Adopted in 1989-93								
$P < 0.6$	0.26	0.15	77%	0.46	0.02	1.00	0.02	291
$P > 0.6$	0.19	0.24	-22%	-0.56	0.15			119
3. Computer Aided Design and Engineering, Adopted in 1988-93								
$P < 0.6$	0.25	0.18	36%	0.18	0.03	0.61	0.07	322
$P > 0.6$	0.33	0.42	-22%	-0.45	0.11			148
4. Product Innovation without Process Innovation, Activities in 1989-91								
$P < 0.6$	0.32	0.12	160%	0.65	0.00	0.89	0.01	350
$P > 0.6$	0.20	0.18	12%	-0.29	0.36			171

Notes: Each plant was placed in either the $\hat{P}(X, \Delta\tau) < 0.6$ group or the $\hat{P}(X, \Delta\tau) > 0.6$ group where \hat{P} is the table 2 probit. Within each group new exporters are compared to non-exporters. Columns 2 and 3 report raw adoption rates. Column 4 is $100 \cdot [(\text{column 2}) / (\text{column 1}) - 1]$. Column 5 reports the coefficient on E in a probit of technology adoption on E (exporter status) and X (employment and productivity in 1984, productivity growth in 1984-88, and 2-digit SIC fixed effects). The technology adoption probit was estimated separately for the two \hat{P} groups. Letting $D_P = 1$ if $\hat{P} < 0.6$ and $D_P = 0$ otherwise, column 7 reports the coefficient δ_p on $E \cdot D_P$ in a probit of technology adoption on $E, X, E \cdot D_P$ and $X \cdot D_P$. This probit pools across the two \hat{P} groups. A low p -value indicates statistical significance.

0.6), new exporters were adopting advanced technologies about as frequently as non-exporters. Column 5 of the table provides the p -value for a test that new exporters and non-exporters adopted at the same rates. The adoption rates differed significantly only for low- \hat{P} plants ($p = 0.00$), exactly as predicted. We will explain how this p -value was estimated shortly.²¹

Looking at the other technologies, a similar pattern emerges. Inspection and communications was, together with MIS, the big innovation being adopted in our period.²² The second set of results in table 5 shows that for low- \hat{P} plants, new exporters were 77% more likely to have adopted inspection and communications technologies during 1989-93 ($p = 0.02$). For high- \hat{P} plants, new exporters were 22% less likely to adopt, a statistically insignificant difference.

The third set of results in table 5 deals with computer aided design and engineering. While design and engineering differences appear in the 1989-93 period, a considerable amount of adoption started in 1988. We therefore report results for 1988-93. The results are as expected, though somewhat weaker than for MIS. Consistent with our theory, this likely reflects the fact that these technologies are relatively inexpensive and hence are affordable even to nonexporters.²³

Turning from processes to product innovation, the fourth set of results in table 5 is from the 1989-91 innovation component of the SIAT survey. The survey asks plants whether they were active in product innovation during the 1989-91 period and if this innovation occurred without any corresponding process innovation.²⁴ For low- \hat{P} plants, new exporters were 160% more likely than non-exporters to have engaged in such activities ($p = 0.00$). Again, there is no statistically significant difference for high- \hat{P} plants ($p = 0.36$).

We next turn to explaining how the statistical tests of the reported differences were estimated. Let T be a plant-level binary indicator of MIS adoption during 1989-93. Let X be as in the probit

²¹One incorrect explanation of these results is that most high- \hat{P} plants had already adopted MIS and that what we are picking up is technology stragglers. In fact, adoption of this technology (and of the others to be discussed) was below 20% in 1988 for all four types of plants. On a separate note, this 20% is lower than one might guess from the non-survey based evidence reported by Feinberg and Keane (2006) and Keane and Feinberg (forthcoming), but this likely reflects their focus on U.S.-owned multinationals. These are particularly large and advanced *continuous* exporters such as General Motors and are therefore not in our sample.

²²Inspection and communications includes (a) automated sensor-based equipment used for inspection/or testing of incoming materials, in-process materials and final products (e.g., tests of failure rates); (b) local area networks for technical data and factory use; inter-company computer networks linking the plant to subcontractors, suppliers and/or customers; (c) programmable controllers; and (d) computers used for control on the factory floor.

²³The survey also asks about automated material handling, integration and control software, and fabrication and assembly. However, adoption rates for these technologies over the 1989-93 period were too infrequent (less than 10%) to be used for inference.

²⁴The question has a vague feeling to it, but this is the nature of questionnaires about innovation. The precise question is as follows: "Please indicate the categories of your innovation activity for the period 1989-1991: Product innovations without change in manufacturing technology."

Table 6. Investing in Productivity: Sensitivity Analysis

$P(X, \Delta\tau)$	Without Plant Controls			With Plant Controls				N
	New	Non-	Difference	Difference		Double Difference		
	Exporters	Exporters		δ	t -stat	δ_p	t -stat	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1. Log Productivity Growth, 1988-96: Technology Subsample								
$P < 0.6$	0.025	-0.004	0.029	0.036	3.70	0.044	2.73	350
$P > 0.6$	-0.004	0.011	-0.015	-0.008	-0.76			171
2. Log Productivity Growth, 1988-96: Full Sample								
$P < 0.6$	0.018	-0.011	0.029	0.037	12.69	0.036	5.97	3,114
$P > 0.6$	-0.003	-0.008	0.005	0.001	0.25			2,133
3. Change in Output per Commodity, 1988-96								
$P < 0.6$	0.059	-0.011	0.070	0.060	7.71	0.048	3.85	1,738
$P > 0.6$	0.040	0.028	0.012	0.012	1.45			1,084

Notes: See the text and notes to table 5 for a full explanation. Column 5 reports OLS estimates of the coefficient on E in a regression of the dependent variable on X and E . Column 7 reports the coefficient on $E \cdot D_p$ in a regression of the dependent variable on X , E , D_p , $X \cdot D_p$ and $E \cdot D_p$. 2-digit SIC fixed effects are used in the first set of results and 4-digit SIC fixed effects are used in the second and third sets of results.

of table 2, but with 2-digit SIC fixed effects. With so few plants, we cannot use the 208 4-digit SIC fixed effects that we have used elsewhere. We estimated a probit of T on X and E separately for plants with $\hat{P} < 0.6$ and $\hat{P} > 0.6$. Let δ be the coefficient on E . It measures the average adoption rate difference between new exporters and non-exporters after controlling for plant characteristics X . Columns 5-6 of table 5 report estimates of δ and their p -values. In all cases, the estimated differences are statistically significant for low- \hat{P} plants and statistically insignificant for high- \hat{P} plants.

We also examined whether the difference for low- \hat{P} plants was statistically larger than the difference for high- \hat{P} plants. Let D_p be an indicator variable for whether $\hat{P} < 0.6$ or $\hat{P} > 0.6$. We pooled all plants and estimated a probit of T on X , E , D_p , $X \cdot D_p$ and $E \cdot D_p$. Let δ_p be the coefficient on $E \cdot D_p$. A test of the difference of differences is the p -value on δ_p . See columns 7-8 of table 5. The difference in differences are statistically significant except for design and engineering.

Finally, we were concerned about the size and representativeness of the 521-plant sample from the SIAT survey. Table 6 addresses this concern. It has the same structure as table 5, but the

dependent variable is log labour productivity growth over 1988-96. Since this is a continuous variable, we use OLS rather than a probit and estimate t -statistics. The first group of results uses only the plants in the SIAT subsample. The second group of results uses our full set of 5,247 plants. As is apparent, the core results (columns 5 and 7) are almost identical for the two samples which suggests that sample selection is not a problem. Interestingly, the t -statistics in columns 6 and 8 are much larger for the larger sample. This suggests that the relatively low statistical significance in table 5 is attributable to the relatively small sample size.

8. Problems with Labour Productivity


We have shown that for plants that were induced by U.S. tariff cuts to export, those with a low \hat{P} experienced (1) high rates of investment in advanced technology adoption and product innovation and (2) high rates of labour productivity growth. It is possible that the labour productivity growth does not reflect any TFP growth, but instead reflects high rates of investment. This seems unlikely — there is abundant and growing evidence that it is precisely investments in MIS and information and communications technologies that drive TFP growth e.g., Stiroh (2002).

Unfortunately, we do not have the capital stock or investment data needed to back up this claim for our particular sample.²⁵ We do however have indirect ways. The first and most important way was suggested to us by Kala Krishna. New exporters obviously increased their sales relative to non-exporters because new exporters started selling into the U.S. market. However, if there were no difference in the TFP performance of new exporters relative to non-exporters, then we would not expect new exporters to increase their *domestic* (Canadian) sales relative to non-exporters. Yet this is exactly what happened.

We computed the Marginal Treatment Effect for the log change in domestic sales between 1984 and 1996. The methodology is identical to what we did for labour productivity in figure 5, with just one difference: the dependent variable in equation (8) is the average annual log change in domestic sales. Domestic sales are total sales less exports. Figure 7 shows the results. For $\hat{P} < 0.5$, plants that were induced to export because of improved market access experienced statistically significant increases in domestic sales. Further, these increases were large. For example, for plants

²⁵On purely theoretical grounds, with CES preferences and Cobb-Douglas production functions, value added per worker is independent of productivity φ . On purely empirical grounds, value added per worker is highly correlated with TFP. Not surprisingly, results using TFP typically carry over to labour productivity. See Amiti and Konings (forthcoming) for a recent example.

Figure 7. Marginal Treatment Effect: Changes in Domestic Sales



with $\hat{P} = 0.35$ the gains were 0.023 log points a year or 20% over 8 years. It is quite remarkable how similar figure 7 is to figure 5. Nothing in our non-parametric econometric model imposes this similarity.²⁶

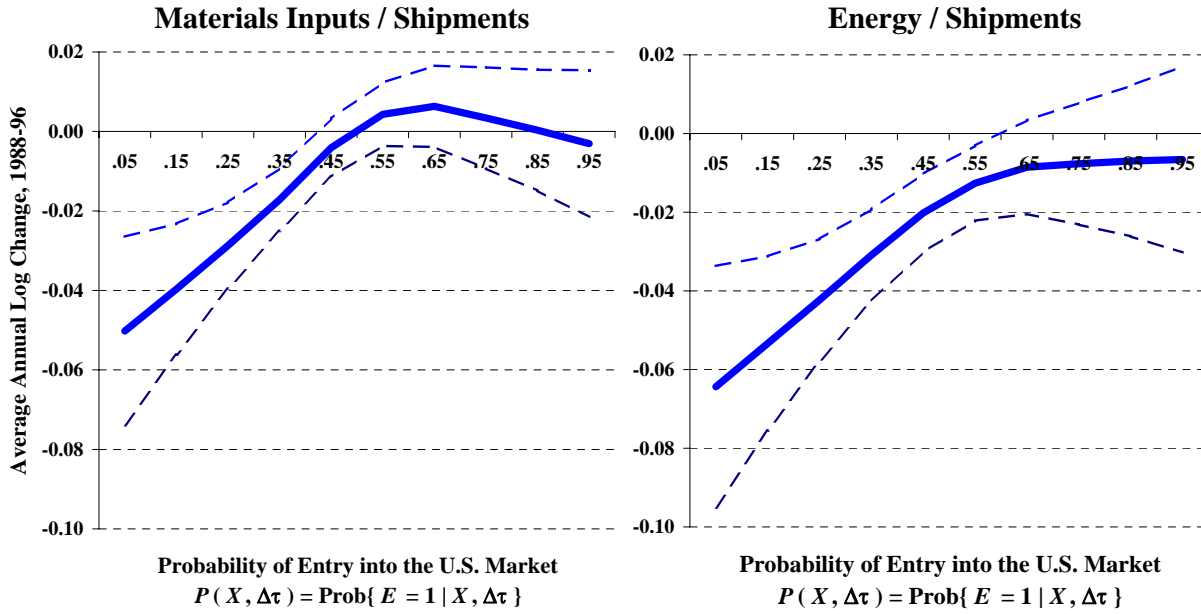
Figure 8 repeats the analysis using materials costs divided by shipments (middle panel) and energy costs divided by shipments (right panel) as the dependent variables. In the region of \hat{P} where we find labour productivity gains we also find reductions in input usage per unit of shipments. This is suggestive of TFP gains.

The last set of results in table 6 suggests another source of productivity gains. The same plants that experienced productivity gains also experienced economically and statistically significant gains in *output per commodity*.

One alternative explanation of our results is that they reflect systematic capacity utilization differences between nonexporters and new exporters. If a plant experiences excess capacity then its sales and variable input usage (labour, materials and energy) will be low per unit of capital. Thus,

²⁶It would have been preferable to work with the log change in domestic sales over 1988-96 rather than 1984-96. We do not have the export data for 1988 needed to compute 1988 domestic sales. However, if we assume that exports were 0 in 1988 and use this assumption to compute log changes in domestic sales 1988-96, we obtain very similar results to those reported in figure 7.

Figure 8. Marginal Treatment Effect: Material and Energy Inputs Per Unit of Output



its TFP and labour productivity will be low. Our results are therefore consistent with the following: for low- \hat{P} plants, nonexporters have excess capacity relative to new exporters while for high- \hat{P} plants, nonexporters have the same excess capacity as new exporters. We are dealing with eight-year changes and hence with long-term capacity under-utilization, which is a much longer time frame than is typical in discussions of procyclical productivity. This aside, Basu (1996, page 719) tackles procyclical productivity using the key observation that ‘material growth is a good measure of unobserved changes in capital and labor utilization.’ The capacity utilization explanation thus predicts that for low- \hat{P} plants, new exporters should have materials-to-sales growth that is the same as for nonexporters.²⁷ Yet as was shown in figure 8, this is not the case. There is thus not much support for an excess-capacity explanation of our results.²⁸

²⁷Or possibly higher due to diminishing returns to fixed capital.

²⁸Yet another explanation of our results is that low- \hat{P} new exporters are providing more goods purchased for resale. However, our results in figures 7 and 8 are virtually unchanged when we use what Statistics Canada refers to as ‘manufacturing activity’ i.e., good produced on the shop floor.

9. Conclusions

This paper presented three core empirical results.

1. Figure 5 showed that there were labour productivity gains from exporting and investing for Canadian manufacturing plants that were induced to export because of improved access to the U.S. market. Further, these labour productivity gains were heterogeneous: only those plants with low pre-Agreement productivity benefitted.
2. Table 5 showed that the labour productivity gainers also had high post-Agreement adoption rates of advanced manufacturing technologies and high post-Agreement levels of product innovation. That is, the new exporters who gained did so by investing in productivity.
3. Figure 7 showed that the pattern of productivity gains mirrored the pattern of domestic (Canadian) sales. The new exporters that experienced productivity gains increased their Canadian sales relative to non-exporters. This is exactly what one would expect if the labour productivity gains reflected underlying TFP gains.

We argued that these facts are consistent with a model featuring two-dimensional heterogeneity i.e., heterogeneity in initial productivity as in Clerides et al. (1998) and Melitz (2003) and heterogeneity in the productivity gains from investing. In particular, for high levels of initial productivity, firms sort into exporting without any implications for productivity growth, just as in the Melitz model. However, for low levels of initial productivity there is a fundamental complementarity between exporting and investing. Exporting makes it more profitable to improve productivity because it increases the output over which the productivity gains will be spread. Thus, there will be plants that find it profitable to export and invest even though it is not profitable only to export or only to invest.

An important feature of our work is the goal of estimating a policy-relevant response. To this end we identified a policy-relevant instrument for exporting, namely, plant-specific tariff cuts mandated under the terms of the Canada-U.S. Free Trade Agreement. This said, the fact that productivity responses were estimated to have a large unobserved component makes the implications for policy tricky. If there was no unobserved response heterogeneity ($U_1 = 0$) then the causal effect of interest would have been the IV estimate of the coefficient on E and we would have blithely claimed that this coefficient is the effect of exporting on productivity *for any plant that*

exports from any country. Obviously, this additional out-of-sample claim requires additional strong assumptions, but assumptions that economists are typically comfortable making. When there is unobserved heterogeneity, the assumptions are less comfortable. We know there were productivity gains for a particular group of plants — Canadian manufacturing plants that were induced to start exporting because of improved access to the U.S. market — but we are not claiming that these gains will accrue to any plant that starts exporting. One way of making this point is to return to our manager who is good both at exploiting export opportunities and at squeezing productivity gains out of new investments. For our results to apply out of sample we would have to claim that the out-of-sample distribution of managers (technically, the joint distribution of U_E and U_1) is the same as our in-sample distribution of managers. We simply have no evidence on this claim one way or the other. We are thus only answering a question about our sample.

This observation is important for thinking about how our results are related to those of Clerides et al. (1998) and Bernard and Jensen (1999). At first blush it would appear that we are contradicting their earlier findings. A more thoughtful interpretation definitely suggests otherwise. For one, the distribution of unobservables (U_E and U_1) in their samples may differ from ours. For another, the reason for exporting in Clerides et al. and Bernard and Jensen is certainly different from the reason for exporting in our Canadian context. Thus, even if the distribution of unobservables in the three samples were the same, the parts of the distribution that started to export are unlikely to be the same. Differences in results are thus to be expected. It would thus be fascinating to see whether our Canadian results carry over to other countries that actively pursued policies of opening up foreign markets, particularly in countries with small domestic markets and hence large complementarities between exporting and investing in productivity.

10. Appendix

A. Theory

Let I be a binary indicator of whether the firm invests ($I = 1$) or not ($I = 0$). Let $\pi_I(E)$ be profits as in equations (1)-(2). The firm chooses one of four alternatives, $(E, I) \in \{(0,0), (0,1), (1,0), (1,1)\}$. Each line in figure 9 corresponds to an indifference condition between two alternatives. For example, the comparison $\pi_1(1) = \pi_0(1)$ is the horizontal line to the right of the Melitz cut-off $F^E/\tau^{-\sigma}A^*$. The label is always above the line and indicates the region for which the inequality holds. For example, $\pi_1(1) > \pi_0(1)$ holds above the line and $\pi_1(1) < \pi_0(1)$ holds below the line. It is trivial to verify that the lines are correctly drawn.

Consider the region to the right of the Melitz cut-off. We know from equation (3) — see the first term and the discussion following the equation — that the firm always exports in this region. We therefore only have to consider alternatives $(E, I) = (1,1)$ and $(E, I) = (1,0)$ i.e., we only have to consider the horizontal line. Thus, the firm exports and invests above the horizontal line and exports without investing below the horizontal line. This completes the proof for the region to the right of the Melitz cut-off.


Now consider the region to the left of the Melitz cut-off, but to the right of $\frac{F^E}{\tau^{-\sigma}A^*} - \frac{F^I}{A}$. Since we are to the left of the Melitz cut-off, the firm will never export without investing i.e., we can ignore the choice $(E, I) = (1,0)$. Above the solid line we have $\pi_1(1) > \pi_0(0)$ and $\pi_1(1) > \pi_1(0)$ i.e., $(E, I) = (1,1)$ is preferred to $(0,0)$ and $(0,1)$. Hence, the firm exports and invests. Below the solid line we have $\pi_1(1) < \pi_0(0)$ and $\pi_1(0) < \pi_0(0)$. Hence the firm neither exports nor invests. This completes the proof of the theory in the main text, which assumed $\varphi_0 > \frac{F^E}{\tau^{-\sigma}A^*} - \frac{F^I}{A}$.

Finally, consider the region to the left of $\frac{F^E}{\tau^{-\sigma}A^*} - \frac{F^I}{A}$. As in the previous paragraph, we need not consider exporting without investing. Above the top solid line we have $\pi_1(1) > \pi_1(0)$ and $\pi_1(1) > \pi_0(0)$. Hence, the firm exports and invests. Below the bottom solid line we have $\pi_1(0) < \pi_0(0)$ and $\pi_1(1) < \pi_0(0)$. Hence the firm neither exports nor invests. Between the two solid lines we have $\pi_1(1) < \pi_1(0)$ and $\pi_1(0) > \pi_0(0)$. Hence the firm invests without exporting.

B. Survivor Bias

Our data consist of plants that were not exporting in 1984 and survived until 1996. We have thus dropped 1984 nonexporters that died before 1996. This creates the potential for survivor bias.

Figure 9. Proof of the Theory



However, any method for dealing with survivor bias will likely lead to larger estimates of the Marginal Treatment Effect. To understand why note that dying plants experience rapidly declining productivity. Griliches and Regev (1995) call this the ‘shadow of death’. If dying plants are also nonexporters — which is highly likely — then including dying plants in the analysis effectively lowers the productivity growth of nonexporters and hence raises the productivity growth of new exporters relative to nonexporters. To examine this more formally, we started by dividing the set of dying plants into two groups. The first group consists of plants that exited during 1989-91. This group exited during the severe 1989-91 recession and before exports started increasing in 1992. (Recall from figure 4 that exports were flat during 1989-91 and almost doubled during 1992-96.) These plants thus died before the effects of the FTA were felt and hence contain little information about these effects. Including them only spuriously biases up our estimates of the Marginal Treatment Effect. The second group exited in 1992-95.

To estimate a Marginal Treatment Effect with 1992-95 exiters we follow a common practice in the labour literature of imputing the dependent variable for exiters e.g., Baker and Benjamin (1997). In our context this means imputing a negative value to 1988-96 average annual log labour productivity growth. We report results for the case where this productivity growth is assumed to be -0.05 , which implies a modest productivity fall of 33% spread out over 8 years. The result appears in figure 10 as the curve labelled 'Include 1992-95 Exiters.' As expected, the effect is larger than our baseline result carried over from figure 5. See the curve labelled 'Baseline.' When we use a more realistic imputation (-0.173 , which implies a productivity fall of 75% over 8 years) the Marginal Treatment Effect doubles in size. In summary, accounting for survivor bias leads to larger estimates of the Marginal Treatment Effect.²⁹

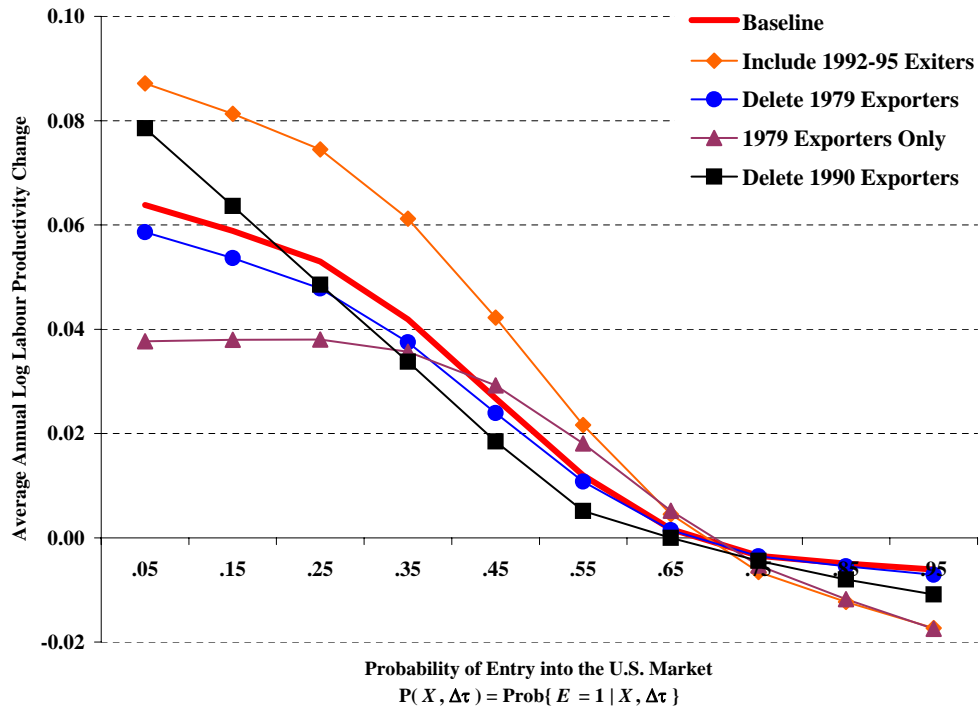
C. Stopping to Export

We next turn to plants that stopped exporting. In addition to the 1984 and 1996 annual surveys of manufacturing, the 1979 and 1990 surveys were the only other surveys that asked exporting questions. Of our 5,247 plants that were nonexporters in 1984, 615 exported in 1979. These plants thus stopped exporting between 1979 and 1984. To examine if this small group of plants has any impact on our conclusions we deleted them from our sample. The re-estimated Marginal Treatment Effect appears in figure 10 as the curve labelled 'Delete 1979 Exporters.' The re-estimated Marginal Treatment Effect is almost identical to our baseline specification.

Interestingly, when we estimate the Marginal Treatment Effect just for the 615 plants that exported in 1979 and stopped by 1984, the Marginal Treatment Effect is smaller for plants with $\hat{P} < 0.45$. This appears as the curve labelled '1979 Exporters Only.' It thus appears that less-productive plants which exported previous to 1984 obtained less of a productivity kick from the FTA inducement to export.

²⁹On a related note, there is the question of whether our sample is representative. The 1984 survey was administered to plants that accounted for a remarkable 91% of total manufacturing output. Thus, only 9% of 1984 output lies outside our scope. Of this 9%, 3% was produced by plants with annual sales of less than \$100,000, what Statistics Canada calls 'short-form' plants. The remaining 6% was produced by plants that primarily served as headquarters for multi-plant firms. Some of our 1984 plants became short-form plants by 1996 and we therefore have less information for them about exporting and commodity composition. These are plants that shrank since 1984, which is why they were demoted to short-form status. Deleting them somewhat biases downward our results because it removes the worst-performing nonexporters from the sample. Baldwin and Gu (2003, footnote 5) show that less than 1% of short-form plants export. We also know that short-form plants are virtually never induced to export because of tariff cuts. (We have considerable unreported evidence available on this point; however, one can already see it in table 3 which shows that the small plants in our sample — which are typically bigger than short-form plants — were unlikely to be induced to export.) We therefore follow Baldwin and Gu (2003) in assuming that 1984 nonexporters who became short-formers were not induced to export.

Figure 10. Marginal Treatment Effect: Deaths and Export ‘Stoppers’



We next turn to 1990 exporters. Using our sample of 1984 nonexporters, we excluded plants that exported in 1990 but stopped exporting by 1996. This appears as the figure 10 curve labelled ‘Delete 1990 Exporters.’ Once again, the results are very similar to our baseline specification.³⁰

D. Balancing

This section reports on the results discussed at the end of section 3. Each element of table 7 is a difference in means between new exporters and nonexporters. The column headings indicate the variable that is being differenced. The differences are computed separately for groups of plants grouped according to \hat{P} . For example, for plants with $\hat{P} \in (0.15, 0.25)$, new exporters have 1984 log labour productivity that is 0.04 log points less than nonexporters. The differences are never significant at the 1% level. Italics indicates significance at the 5% level.

³⁰Interestingly, we could not estimate a Marginal Treatment Effect separately for these deleted plants. In the starting-to-export probit (see table 2), $\Delta\tau$ is economically and statistically insignificant ($p = 0.24$). This is the *only* instance in this paper where $\Delta\tau$ is not an absolutely excellent predictor of starting to export. Apparently these plants entered for reasons that had little to do with the FTA and exited due to the severe 1989-91 recession.

Table 7. Tests of the Balancing Hypothesis

		Differences in Means Between $E = 1$ and $E = 0$		
	$P(X, \Delta\tau)$	Log Productivity	Log Employment	Change in Productivity (1984-88)
1	.00-.05	-0.55	<i>-0.52</i>	-0.04
2	.05-.15	0.32	<i>0.30</i>	0.02
3	.15-.25	-0.04	-0.05	0.01
4	.25-.35	-0.01	-0.02	0.00
5	.35-.45	0.10	0.08	0.02
6	.45-.55	0.11	0.14	<i>-0.03</i>
7	.55-.65	0.06	0.07	-0.01
8	.65-.75	-0.05	-0.06	0.00
9	.75-.85	-0.11	-0.10	-0.01
10	.85-.95	-0.61	<i>-0.64</i>	0.02
11	.95-1.00	-2.18	<i>-2.35</i>	0.17

Notes: This table reports differences in means between new exporters and nonexporters. The differenced variables are 1984 log labour productivity, 1984 log employment and 1984-88 log labour productivity growth. The differences are never significant at the 1% level. Italics indicates statistical significance at the 5% level.

E. Details of Nonparametric Estimation

This appendix describes details of the estimation of the Marginal Treatment Effect. We follow the multi-step procedure of Carneiro et al. (2003). The first group of steps estimate equation (11) non-parametrically and is essentially the non-parametric counterpart to the linear IV estimates of β_{1q} reported in column 5 of table 4. Let X' consist of log productivity in 1984, log employment in 1984 and log productivity growth in 1984-1988. X is X' plus the 4-digit industry fixed effects. We have assumed that $\beta_0(X) = \beta_0 X$ and $\beta_1(X) = \beta_1 + \beta'_1 X'$. Let $\widehat{P} = \widehat{P}(X, \Delta\tau)$ be our probit estimate from table 2. Plugging this information into equation (11) yields

$$\mathbf{E}[\Delta\varphi|\widehat{P}] = \beta_0 X + \beta_1 \widehat{P} + \beta'_1 X' \cdot \widehat{P} + \mathbf{E}[U_1 E|\widehat{P}].$$

We non-parametrically estimate this equation as follows. (1) Regress $\Delta\varphi$, X and $X' \cdot \widehat{P}$ on \widehat{P} using local linear regression. (2) Letting $\widehat{\varepsilon}_{\Delta\varphi}$, $\widehat{\varepsilon}_X$, and $\widehat{\varepsilon}_{X'P}$ be the respective residuals from these regressions, regress $\widehat{\varepsilon}_{\Delta\varphi}$ on $\widehat{\varepsilon}_X$ and $\widehat{\varepsilon}_{X'P}$ using OLS in order to estimate β_0 and β'_1 .³¹ Empirically we find $\beta'_1 = 0$.³² Accordingly, we repeated step (1) without $X' \cdot \widehat{P}$ and step (2) without $\widehat{\varepsilon}_{X'P}$. The resulting estimates of the elements of β_0 are -0.062 ($t = -25.04$) for 1984 log labour productivity, 0.004 ($t = 2.24$) for 1984 log employment and -0.329 ($t = -37.20$) for 1984-88 log labour productivity growth. We do not report the estimated industry fixed effects that are part of X . (3) Let $\widehat{\varepsilon}$ be the residual from the step (2) regression. It is an estimate of $\beta_1 \widehat{P} + \mathbf{E}[U_1 E|\widehat{P}]$.³³ Regressing $\widehat{\varepsilon}$ on \widehat{P} using local linear regression, we obtain a nonparametric estimate of $\widehat{\varepsilon}(\widehat{P})$. Putting steps (1)-(3) together provides an estimate of $\mathbf{E}[\Delta\varphi|\widehat{P}]$ which we denote by $\widehat{\mathbf{E}[\Delta\varphi|\widehat{P}]} \equiv \widehat{\beta}_0 X + \widehat{\varepsilon}(\widehat{P})$.

The second group of steps involves estimating the derivative in equation (12). We do this by numerically differentiating $\widehat{\mathbf{E}[\Delta\varphi|\widehat{P}]}$. Specifically, plants were divided into the 11 groups of table 3. Means of $\widehat{\mathbf{E}[\Delta\varphi|\widehat{P}]}$ were then calculated for each group and derivatives were calculated by finite differencing across neighbouring groups.³⁴

³¹If the reader finds this difficult to interpret, step 1 is related to the first stage of IV and step 2 is related to the second stage of IV.

³²The elements of β'_1 are individually insignificant ($t < 2.00$ in all three cases) and jointly insignificant ($F = 2.05$ which has a p -value of 0.031). Once again, this insignificance supports the balancing hypothesis.

³³In discussing β_1 we have been ignoring an identification issue. To see it simply, suppose that $U_1 = c_1 + c_2 \widehat{P}$ so that $\mathbf{E}[U_1 E|\widehat{P}] = c_1 \widehat{P} + c_2 (\widehat{P})^2$. Then $\beta_1 \widehat{P} + \mathbf{E}[U_1 E|\widehat{P}] = (\beta_1 + c_1) \widehat{P} + c_2 (\widehat{P})^2$. That is, only $\beta_1 + c_1$ is identified.

³⁴Local linear regressions were done using the SAS LOESS procedure. The procedure uses optimal smoothing based on Hurvich, Simonoff and Tsai (1998).

F. Two Parametric Specification Tests

The aim of estimating heterogeneous responses translates into estimating whether $\mathbf{E}[\Delta\varphi|\widehat{P}(X,\Delta\tau)]$ depends non-linearly on \widehat{P} . To see this in the simplest way possible, rather than estimating $\mathbf{E}[U_1E|\widehat{P}]$ non-parametrically, suppose we know that $\mathbf{E}[U_1E|\widehat{P}] = c_0 + c_1\widehat{P} + c_2(\widehat{P})^2 + c_3(\widehat{P})^3$ for some unknown coefficients $c_i, i = 0, \dots, 3$. Then from equation (11) with X suppressed,

$$\mathbf{E}[\Delta\varphi|\widehat{P}] = (\beta_0 + c_0) + (\beta_1 + c_1)\widehat{P} + c_2(\widehat{P})^2 + c_3(\widehat{P})^3. \quad (13)$$

From equation (12), the Marginal Treatment Effect is $(\beta_1 + c_1) + 2c_2\widehat{P} + 3c_3(\widehat{P})^2$. This means that there is heterogeneity only if c_2 and/or c_3 are not zero. In terms of figure 5, this means that the line is non-horizontal only if c_2 and/or c_3 are not zero.

This parametric example motivates the simple parametric test of heterogeneity suggested by Carneiro et al. (2003). Estimate equation (13) using OLS and test for $c_2 = c_3 = 0$. The t -statistics for c_2 and c_3 are -5.39 and 4.39 , respectively. The F -statistic for $c_2 = c_3 = 0$ is 18.31 ($p = 0.000$). Thus, we can reject homogeneity. Similar results obtain using a fourth-order polynomial in equation (13). In this case $F = 18.56$.

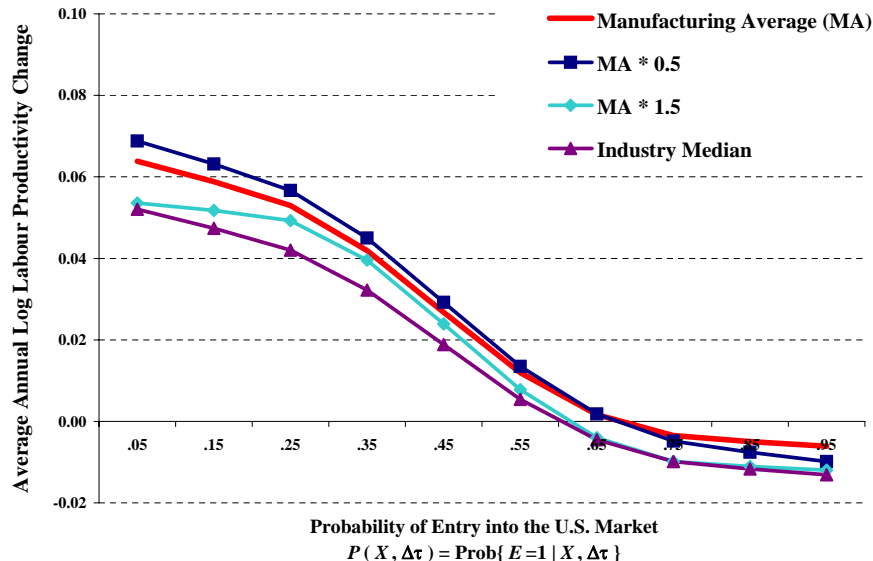
We can also use this parametric approach to construct an informal over-identification test. Although we have only one underlying instrument $\Delta\tau$, by transforming it *non-linearly* into $\widehat{P}(X, \Delta\tau)$ we can use the non-linear functional form to identify a second instrument \widehat{P} . See the discussion at the end of section 4. This means that the functional form provides us with over-identification. It thus allows us to include $\Delta\tau$ directly into the second stage i.e., into equation (13). When we do so the t -statistic on $\Delta\tau$ is 0.41 . Thus, this informal over-identification test leads us to reject the hypothesis that $\Delta\tau$ belongs in the second-stage productivity equation.

G. Sensitivity to the Tariff Threshold and Definition of New Exporters

Figure 11 shows that our results are not sensitive to the specification of $\Delta\tau$. See the figure notes and section 6 for a discussion.

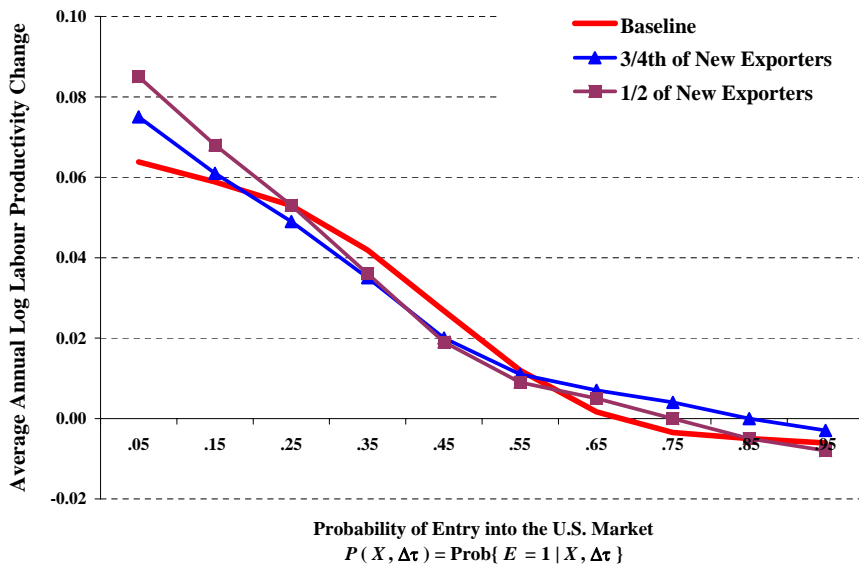
Our results are also not particularly sensitive to the definition of exporting. Let r be the ratio of exports to sales. We re-estimated the Marginal Treatment Effect for the one-half of all new exporters in an industry with the largest values of r . We also used three-quarters in place of one-half. The results appear in appendix figure 12 and are very similar to our figure 5 baseline results.

Figure 11. Marginal Treatment Effect: Sensitivity to Tariff Threshold



Notes: This figure provides estimates of the Marginal Treatment Effect using different definitions of the U.S. tariff cut. A plant has $\Delta\tau = 1$ if its tariff cut was deeper than some threshold $\bar{\Delta\tau}$. In our baseline specification of figure 5, here labelled 'Manufacturing Average (MA)', the threshold is the average tariff cut for all plants. In 'MA * 0.5' the threshold is 50% below the average tariff cut for all plants. In 'MA * 1.5' the threshold is 50% above the average tariff cut for all plants. In 'Industry Median' the threshold is the median tariff in the industry.

Figure 12. Marginal Treatment Effect: Sensitivity to Definition of New Exporters



Notes: See appendix G for a discussion.

References

- Alvarez, Roberto and Ricardo A. Lopez**, "Exporting and Performance: Evidence from Chilean Plants," *Canadian Journal of Economics*, November 2005, 38 (4), 1384–1400.
- Amiti, Mary and Jozef Konings**, "Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia," *American Economic Review*, forthcoming.
- Atkeson, Andrew and Ariel Burstein**, "Innovation, Firm Dynamics, and International Trade," November 2006. Mimeo, UCLA.
- Aw, Bee Yan, Mark J. Roberts, and Tor Winston**, "Export Market Participation, Investments in R&D and Worker Training, and the Evolution of Firm Productivity," *World Economy*, January 2007, 14 (1), 83–104.
- , **Sukkyun Chung, and Mark J. Roberts**, "Productivity and Turnover in the Export Market: Micro Evidence from Taiwan (China) and the Republic of Korea," *World Bank Economic Review*, January 2000, 14 (1), 65–90.
- Baggs, Jen**, "Firm Survival and Exit in Response to Trade Liberalization," *Canadian Journal of Economics*, November 2005, 38 (4), 1364–1383.
- **and James Brander**, "Trade Liberalization, Profitability, and Financial Leverage," *Journal of International Business Studies*, March 2006, 37 (2), 196–211.
- Baker, Michael and Dwayne Benjamin**, "The Role of the Family in Immigrants' Labor Market Activity: An Evaluation of Alternative Explanations," *American Economic Review*, September 1997, 87 (4), 705–707.
- Baldwin, John R. and Peter Hanel**, *Innovation and Knowledge Creation in an Open Economy*, Cambridge: Cambridge University Press, 2003.
- **and Wulong Gu**, "Participation in Export Markets and Productivity Performance in Canadian Manufacturing," *Canadian Journal of Economics*, 2003, 36 (3), 634–657.
- **and —**, "Trade Liberalization: Export-Market Participation, Productivity Growth and Innovation," *Oxford Review of Economic Policy*, Autumn 2004, 20 (3), 372–392.
- **and —**, "The Impact of Trade on Plant Scale, Production-Run Length and Diversification," Working Paper 11F0027MIE–038, Statistics Canada, May 2006.
- , **Desmond Beckstead, and Richard Caves**, "Changes in the Diversification of Canadian Manufacturing Firms (1973-1997): A Move to Specialization," Analytical Studies Branch Research Paper Series 11F0019MIE2002179, Statistics Canada, February 2002.
- Baldwin, John, Richard E. Caves, and Wulong Gu**, "Responses to Trade Liberalization: Changes in Product Diversification in Foreign and Domestic Controlled Plants," in Lorraine Eden and Wendy Dobson, eds., *Governance, Multinationals and Growth*, Cheltenham, UK: Edward Elgar Publishing, 2005.
- Basu, Susanto**, "Procyclical Productivity: Increasing Returns or Cyclical Utilization?," *Quarterly Journal of Economics*, August 1996, 111 (3), 719–751.
- Bernard, Andrew B. and J. Bradford Jensen**, "Exceptional Exporter Performance: Cause, Effect, or Both?," *Journal of International Economics*, February 1999, 47 (1), 1–25.

- **and** –, “Exporting and Productivity in the U.S.,” *Oxford Review of Economic Policy*, Autumn 2004, 20 (3), 343–357.
- **and Joachim Wagner**, “Exports and Success in German Manufacturing,” *Weltwirtschaftliches Archiv*, 1997, 133 (1), 134–157.
- Brander, James A.**, “Election Polls, Free Trade, and the Stock Market: Evidence from the 1988 Canadian General Election,” *Canadian Journal of Economics*, November 1991, 24 (4), 827–843.
- Bustos, Paula**, “Rising Wage Inequality in the Argentinean manufacturing Sector: The Impact of Trade and Foreign Investment on Technology and Skill Upgrading,” 2005. Mimeo, CREI, Universitat Pompeu Fabra.
- Card, David**, “Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems,” *Econometrica*, September 2001, 69 (5), 1127–1160.
- Carneiro, Pedro, James J. Heckman, and Edward Vytlačil**, “Understanding What Instrumental Variables Estimate: Estimating Marginal and Average Returns to Education,” 2003. Mimeo, Columbia University.
- Clerides, Sofronis, Saul Lach, and James R. Tybout**, “Is Learning by Exporting Important? Micro-dynamic Evidence from Colombia, Mexico, and Morocco,” *Quarterly Journal of Economics*, August 1998, 113 (3), 903–947.
- Costantini, James and Marc Melitz**, “The Dynamics of Firm Level Adjustment to Trade Liberalization,” 2007. Mimeo, CEPR.
- De Loecker, Jan**, “Do Exports Generate Higher Productivity? Evidence from Slovenia,” *Journal of International Economics*, forthcoming.
- Delgado, Miguel A., Jose C. Fariñas, and Sonia Ruano**, “Firm Productivity and Export Markets: A Non-parametric Approach,” *Journal of International Economics*, August 2002, 57 (2), 397–422.
- Ederington, Josh and Phillip McCalman**, “Endogenous Firm Heterogeneity and the Dynamics of Trade Liberalization,” *Journal of International Economics*, forthcoming.
- Ekholm, Karolina and Karen Helene Midelfart**, “Relative Wages and Trade-Induced Changes in Technology,” *European Economic Review*, August 2005, 49 (6), 1637–1663.
- Feinberg, Susan E. and Michael P. Keane**, “Tariff Effects on MNC Decisions to Engage in Intra-Firm and Arms-Length Trade,” 2005. Mimeo, Rutgers University.
- **and** –, “Accounting for the Growth of MNC-Based Trade Using a Structural Model of U.S. MNCs,” *American Economic Review*, December 2006, 96 (5), 1515–1558.
- Fernandes, Ana M. and Alberto E. Isgut**, “Learning-by-Exporting Effects: Are They for Real?,” 2006. Mimeo, The World Bank and Institute for Competitiveness & Prosperity.
- Gaston, Noel and Daniel Trefler**, “The Labour Market Consequences of the Canada-U.S. Free Trade Agreement,” *Canadian Journal of Economics*, February 1997, 30 (1), 18–41.
- Griliches, Zvi and Haim Regev**, “Firm Productivity in Israeli Industry: 1979-1988,” *Journal of Econometrics*, January 1995, 65 (1), 175–203.
- Hallward-Driemeier, Mary, Giuseppe Iarossi, and Kenneth L. Sokoloff**, “Exports and Manufacturing Productivity in East Asia: A Comparative Analysis with Firm-Level Data,” 2005. Mimeo, The World Bank and UCLA.

- Head, Keith and John Ries**, "Rationalization Effects of Tariff Reductions," *Journal of International Economics*, April 1999, 47 (2), 295–320.
- and —, "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade," *American Economic Review*, September 2001, 91 (4), 858–876.
- Heckman, James J. and Edward Vytlacil**, "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects," *Proceedings of the National Academy of Sciences*, April 1999, 96 (8), 4730–4734.
- and —, "Structural Equations, Treatment Effects and Econometric Policy Evaluation," *Econometrica*, April 2005, 73 (3), 669–738.
- Helpman, Elhanan**, "Trade, FDI and the Organization of Firms," *Journal of Economic Literature*, September 2006, XLIV (3), 589–630.
- Hurvich, Clifford M., Jeffrey S. Simonoff, and Chih-Ling Tsai**, "Smoothing Parameter Selection in Nonparametric Regression Using an Improved Akaike Information Criterion," *Journal of the Royal Statistical Society, Series B*, 1998, 60 (2), 271–293.
- Imbens, Guido W. and Joshua D. Angrist**, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, March 1994, 62 (2), 467–475.
- Keane, Michael P. and Susan E. Feinberg**, "Advances in Logistics and the Growth of Intra-Firm Trade: The Case of Canadian Affiliates of U.S. Multinationals, 1984-1995," *Journal of Industrial Economics*, forthcoming.
- Lileeva, Alla**, "Import Competition and Selection," Working Paper, York University, 2004.
- Melitz, Marc J.**, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, November 2003, 71 (6), 1695–1725.
- Park, Albert, Dean Yang, Xinzheng Shi, and Yuan Jiang**, "Exporting and Firm Performance: Chinese Exporters and the Asian Financial Crisis," 2006. Mimeo, University of Michigan.
- Roberts, Mark J. and James R. Tybout**, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, September 1997, 87 (4), 545–564.
- Romalis, John**, "NAFTA's and CUSFTA's Impact on International Trade," *Review of Economics and Statistics*, forthcoming.
- Rosenbaum, Paul R. and Donald B. Rubin**, "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, April 1983, 70 (1), 41–55.
- Stiroh, Kevin J.**, "Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?," *American Economic Review*, December 2002, 92 (5), 1559–1576.
- Thompson, Aileen J.**, "The Anticipated Sectoral Adjustment to the Canada - United States Free Trade Agreement: An Event Study Analysis," *Canadian Journal of Economics*, May 1993, 26 (2), 253–271.
- Trefler, Daniel**, "The Long and Short of the Canada-U.S. Free Trade Agreement," *American Economic Review*, September 2004, 94 (4), 870–895.
- Van Biesebroeck, Johannes**, "Exporting Raises Productivity in Sub-Saharan African Manufacturing Firms," *Journal of International Economics*, December 2004, 67 (2), 373–391.

Yeaple, Stephen Ross, "A Simple Model of Firm Heterogeneity, International Trade, and Wages,"
Journal of International Economics, January 2005, 65 (1), 1–20.