

Economies of Scale and the Size of Exporters*

Roc Armenter

Federal Reserve Bank of Philadelphia

Miklós Koren

Central European University, IEHAS and CEPR

February 11, 2010

Abstract

Exporters are few—less than one-fifth among U.S. manufacturing firms—and are larger than non-exporting firms—about 4-5 times more total sales per firm. These facts are often cited as support for models with economies of scale and firm heterogeneity as in Melitz (2003). We find that the basic Melitz model cannot simultaneously match the size and share of exporters given the observed distribution of total sales. Instead exporters are expected to be between 90 and 100 times larger than non-exporters. It is easy to reconcile the model with the data. However, a lot of variation independent of firm size is needed to do so. This suggests that economies of scale play only a minor role in determining a firm's export status. We show that the augmented model also has markedly different implications in the event of a trade liberalization. Most of the adjustment is through the intensive margin and productivity gains due to reallocation are halved.

*We thank George Alessandria, Holger Breinlich, Ariel Burstein, Thomas Chaney, Alejandro Cuñat, Marc Melitz, and Kim Ruhl as well as seminar participants at the Central European University, the NY Fed Quantitative Macroeconomic workshop, IEHAS, the Philadelphia Fed, the Bank of Spain, the University of Texas at Austin, the University of Essex, and UCLA for comments. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

Exporters represent less than one-fifth of all U.S. manufacturing firms and are larger than non-exporters—total sales per firm are about 4-5 times higher for exporters.¹ These stylized facts are often cited as support for models with economies of scale and firm heterogeneity. In Melitz (2003), only larger, more productive firms generate enough net revenues abroad to cover the fixed costs associated with exporting. A key implication is that, in the event of a trade liberalization, resources will be reallocated toward the more productive firms, raising the average productivity in the industry.

In this paper we explore the *quantitative* implications of the Melitz model for the frequency and characteristics of exporters. We focus on the size of exporters since economies of scale are at the core of the theory in Melitz (2003). According to the basic version of the Melitz model, we should observe strict sorting by size; that is, the smallest exporter should be larger than the largest non-exporter. The set of exporters is thus easily characterized with a cut-off rule in terms of total sales. We use this equilibrium restriction to derive the quantitative predictions of the theory without having to specify the full model.

We find that the basic Melitz model cannot simultaneously match the share and size of exporters in the data given the large skewness observed in the distribution of total sales.² Given that exporters are roughly one-fifth of all firms, strict sorting suggests that exporters should be concentrated in the top quintile of the firm size distribution. We can thus obtain the model prediction for the size of exporters by comparing the top quintile firms with the rest. We find that the firms below the top quintile are quite small (an average of \$740,000 total sales), while the firms in the top quintile are much larger (\$70 million on average). Hence, the basic Melitz model greatly overpredicts the exporter size premium: exporters should have between 90 and 100 times more total sales than non-exporters.

Of course, we did not expect the strict sorting of exporters to hold *exactly* in the data. Melitz (2003) certainly does not intend to preclude the importance of other idiosyncratic factors, unrelated to size, in the firm’s decision to export. For example, the costs of exporting—variable or fixed—may vary from firm to firm.

We proceed to reconcile the model with the data by introducing further firm-level heterogeneity. At first pass we take a latent variable approach so there is no need to specify which are the additional sources of variation in the export decision—only that these factors are independent of firm size. We find that we need a lot of independent variation in order to bridge the large gap between the model and the data. This suggests that economies of scale, as captured by fixed costs, are not the main determinant of a firm’s export status.

Importantly, we find that the large amount of independent variation needed “waters down” the core mechanism in the Melitz model. To illustrate this we set up a very simple

¹Both facts are documented in Bernard, Jensen, Redding and Schott (2007) for the universe of U.S. manufacturing firms operating in 2002.

²We use the distribution of total sales for manufacturing firms (NAICS code 31-33) as given by the 2002 Statistics of U.S. Businesses of the Census.

partial equilibrium model based on Melitz (2003). We then explore two versions of the model. In the first version productivity is the only source of variation across firms. The model displays the strict sorting property and, as we discussed above, overstates the size of exporters. In the second version we assume firms face heterogeneous fixed costs. We then calibrate the distribution of fixed costs to match the share and size of exporters.³ From the previous discussion, we know we need a huge amount of dispersion in the fixed costs. We then compare the models in the event of a trade liberalization, with trade costs falling by half.

We find that, in the aggregate, both models are indistinguishable from a representative-firm model. Free entry amplifies the response of exports to a trade liberalization through the love-of-variety effect. However, firm-level heterogeneity virtually cancels all the amplification introduced by free entry.⁴

Behind the similarities in the aggregates there are large differences in the margins of adjustment. Critically, the augmented model has only a minor role for the extensive margin. First, there is much less entry than in the standard Melitz model in response to a fall in trade costs. We find that the growth rate of the number of exporters is more than 60% in the strict sorting version, but only 15% when all of the latent heterogeneity is accounted for. Second, new and existing exporters have relatively similar size in the latent heterogeneity model. This stands in marked contrast with the standard Melitz model in which new exporters are 10 to 12 times smaller than the average exporter prior to trade liberalization. The differences across models are more striking given that both models have very similar implications for the aggregate trade variables.

We seek to quantify further the role of the extensive margin in both models. To this end we decompose total export growth in an intensive and extensive margin. In the standard version the extensive margin accounts for more than 60% percent of the trade growth. Once we match the size of exporters, the extensive margin accounts for less than 20%. The muted response of export participation is quite surprising given the Melitz model is by now considered the workhorse model for the fast-growing literature on the extensive margin.

We also look at the productivity gains due to the reallocation effect across both models. We find that the productivity gains due to trade liberalization are halved in the augmented model. This is perhaps not surprising given the previous results. However, it must be noted that we miss on the productivity gains due to exit in the domestic market, since ours is partial equilibrium model and the wage rate—or more generally, unit costs—is taken as given.

³The distribution of fixed costs is the only difference between the two models. We use a common value for all remaining parameters, including those governing the distribution of productivity.

⁴In a recent paper, Atkeson and Burstein (2008) find that the firms' responses to a trade liberalization do not quantitatively impact the implications for aggregate productivity and welfare in a general equilibrium model. We must emphasize that ours is a partial equilibrium model, so the offsetting effects are not due to factor prices or aggregate demand adjusting in general equilibrium. See also Arkolakis et al. (2008).

Finally, we explore alternative specifications for the additional heterogeneity. First, we discuss the presence of small exporters in the data, that is, exporters with few export sales. This observation has spurred some recent research, most notably Arkolakis (2008). We argue that, as long as the strict sorting is preserved, the size of exporters will be overstated in the model. Moreover, we show that introducing latent heterogeneity gives rise to a significant fraction of small exporters even with a constant export intensity. We also ask whether the latent heterogeneity stems from differences across broadly defined industries. If that were the case, we would find that the size premium of exporters within each industry would be closer to the data. Instead, we find that for each 3-digit manufacturing industry, the implied size premium remains implausibly large.

We then move to the possibility that the fixed costs of exporting are sunk. In this case there is exporter hysteresis and the particular history of each firm can be the source of heterogeneity at play. This form of heterogeneity would be compatible with large adjustments along the extensive margin. We show that sunk costs can easily match the share and size of exporters. However, it is now the size of the *new* exporters that presents a puzzle: new exporters would be significantly larger than incumbent exporters.

Our paper relates to a growing body of work that studies the extensive margin and its quantitative implications in open macroeconomic models. In an early contribution, Ruhl (2003) shows how the extensive margin leads to significantly different elasticities in the short and long run. Ghironi and Melitz (2005, 2007) explore how the dynamics of an international real business cycle model are shaped by the extensive margin. Alessandria and Choi (2007a, 2008) have instead pursued the implications of sunk costs for trade dynamics and estimates of trade costs. Atkeson and Burstein (2008) explore how entry in export markets shapes the innovation decisions of the firms.

We are also not the first to introduce additional heterogeneity in the Melitz model. Das, Roberts, and Tybout (2007) estimate a partial equilibrium model with entry and exit for Colombian data. The estimation assumes a very rich error structure and thus many possible sources of heterogeneity, including stochastic fixed costs. Das, Roberts, and Tybout (2007) report that fixed costs are, on average, quite large. Our focus is instead on the dispersion needed to match the data rather than the point estimates. Eaton, Kortum, and Kramarz (2008) also estimate a model of firm heterogeneity and export participation with several sources of variation.

Some researchers have also extended the Melitz model seeking to explain some particular facts but without pursuing the quantitative implications of the theory. Johnson (2007) and Baldwin and Harrigan (2007), among others, consider firm-level variation in output quality in order to explain the correlation between unit values and distance. Kugler and Verhoogen (2008) further expand the model to explain the correlation of output and input prices with plant size observed in Colombia. In a similar setting, Hallak and Sivadasan (2008) talk of differences in “caliber” among firms and note that such differences are necessary to break the monotone relationship between size and export status.

The paper is organized as follows. Section 2 derives the size of exporters in the basic Melitz model. Section 3 introduces a latent variable to reconcile the model with the data. In Section 4 we present a simple structural framework that is flexible enough to encompass strict sorting and additional sources of heterogeneity. The implications of the latent heterogeneity are then explored in Section 5. We discuss some alternative specifications for the heterogeneity in Section 6. We also include a brief discussion on the empirical literature on the extensive margin. Section 8 concludes.

2 On the Size of Exporters

In the Melitz model, firm i will export only if its foreign sales, net of the associated variable costs, would cover the fixed costs associated with exporting,

$$r_i^* - c_i^* \geq f. \tag{1}$$

Firms differ in their labor productivity or in the quality of goods produced. Either way, more efficient firms can generate more net income abroad and are thus more likely to be exporters. Hence, the model predicts that exporters are more productive and sell higher quality goods than non-exporters. Incidentally, more efficient firms also sell more in the domestic market, so exporters are unambiguously larger than non-exporters in terms of total sales. These qualitative predictions are borne in the data and often cited as support for the Melitz model.

We seek to explore the *quantitative* predictions of the model on the exporters' characteristics. The first step is to rewrite the entry condition (1) in terms of total sales. Data on total firm revenue are easily accessible and we can observe a firm's total sales independently of whether the firm exports. In the basic Melitz model, there is a tight relationship between a firm's total sales and its underlying efficiency parameter: a firm with higher productivity will always have more total sales in equilibrium. Since net income abroad is also strictly increasing in productivity, we have an increasing, monotone relationship between total sales and net income abroad—even if the latter is a counterfactual because the firm does not export.

We can thus summarize the model's predictions for the set of exporters with a simple threshold rule in place of (1): firm i will export only if its total sales r_i are above some level t ,

$$r_i \geq t. \tag{2}$$

It is immediate that exporters and non-exporters are strictly sorted by size; that is, the smallest exporting firm has more total sales than the largest non-exporting firm. The value of the threshold level t is determined in equilibrium and is bound to depend on the model's parameters.

We can, though, easily derive the model's predictions using the share of exporters and the distribution of total sales observed in the data. The threshold condition (2) implies that

the share of exporters is equal to the fraction of firms with total sales equal or larger than t . We have thus that

$$s_x = 1 - \Psi_r(t),$$

where Ψ_r is the empirical c.d.f. of total sales, and s_x is the fraction of firms with positive export sales. In other words, we can solve for the $(1 - s_x)$ th percentile in the distribution of total sales and obtain the threshold t consistent with the observed share of exporters. We can then easily compute the truncated mean,

$$E\{r_i | r_i \geq t\} = \int_t^\infty r_i \frac{d\Psi_r(r_i)}{s_x},$$

so we can compare the model's implications for the average total sales for exporters and non-exporters.

Bernard, Jensen, Redding and Schott (2007) report that only 18% of the U.S. manufacturing firms had positive sales abroad in 2002.⁵ Exporters, thus, are relatively rare. For the distribution of total sales we look at the 2002 Statistics of U.S. Businesses of the Census. Table 1 summarizes the data for manufacturing firms. As is well known, there is an enormous amount of skewness in the firm-size distribution. The average firm sells \$13.2 million, and yet 45% of the firms sell less than \$500,000. In short, there are many, many small firms and a few very, very large ones.

| Size bin | Frequency | Cumulative Frequency | Average sales |
|-----------------------|-----------|----------------------|-----------------------|
| 0–\$100,000 | 0.145 | 0.145 | \$55,600 |
| \$100,000–\$500,000 | 0.305 | 0.450 | \$257,000 |
| \$500,000–\$1 million | 0.144 | 0.594 | \$718,000 |
| \$1–5 million | 0.257 | 0.851 | \$2.26 million |
| \$5–10 million | 0.060 | 0.911 | \$6.84 million |
| \$10–50 million | 0.063 | 0.974 | \$19.3 million |
| \$50–100 million | 0.010 | 0.984 | \$56.4 million |
| over \$100 million | 0.015 | 1.000 | \$670 million |
| <i>All Firms</i> | | | <i>\$13.2 million</i> |

Table 1: **The distribution of firm sales in manufacturing** – *Census*

From Table 1 we see that the 82nd percentile falls somewhere between \$1 and \$5 million sales, definitively closer to the latter. This already suggests that non-exporters are expected to be quite small under strict sorting: firms below \$1 million sales average less than \$320,000 in sales and would constitute over 70% of all non-exporters. At the same time, exporters

⁵Exporters are similarly scarce if we look at plants or establishments rather than firms—see Bernard, Eaton, Jensen, and Kortum (2003), for example. The scarcity of exporters has also been confirmed in a variety of countries.

will be pretty large. Firms above \$1 million average \$82 million in total sales and would represent over 80% of all exporters. We want to be a little bit more precise than this, though. For the firms within the range of \$1 – \$5 million we do not know the exact distribution: we assume these firms follow a two-side truncated Pareto distribution, parameterized to match the average total sales in the range (\$2.26 million).⁶

We find that strict sorting by size implies that exporters should sell, on average, between 90 and 100 times more than non-exporters. The exporter size premium remains very large for whatever distribution one assumes for the firms in the \$1 – \$5 million range. As a check we computed a lower bound on the exporter’s size by taking all firms with sales between \$1 million and \$5 million to be identical, with total sales equal to the bin’s average (\$2.26 million). This would be the smallest exporter size premium compatible with strict sorting and the data. Even in this case we find that exporters are predicted to be more than 80 times larger than non-exporters.

How does the implied exporter size premium compare with the data? Bernard et al. (2007) report that exporters are 4 to 5 times larger than non-exporters.⁷ Strict sorting thus greatly overpredicts the size differences between exporters and non-exporters. Table 2 compares the size of exporters and non-exporters in the model with the data for U.S. manufacturing firms in 2002. Under the hypothesis of strict sorting, exporters should have, on average, \$70.1 million in total sales—double what we actually observe. In the model non-exporters are expected to be very small, just \$740,000 in total sales. In the data, though, there are clearly some non-exporters that are large enough to bring their average total sales above \$8 million.

| | Data | Strict Sorting |
|-------------------------------|----------------|----------------|
| Average Sales - Non-exporters | \$8.1 million | \$740,000 |
| Average Sales - Exporters | \$36.4 million | \$70.1 million |
| Exporter Size Premium | 4.5 | 95 |

Table 2: **Exporters and non-exporters in the model and the data**

The model cannot match the size of exporters and non-exporters even if we allow some leeway in the fraction of exporters. Table 3 reproduces the previous exercise for different assumptions on the fraction of exporters, s_x . As we increase the share of exporters, the model’s prediction for the size premium declines but does not get anywhere close to the actual number. For example, say we take exporters to be twice as frequent as they actually

⁶There is no need to assume anything about the distribution of firms outside this size bin: they are all squarely in the exporter or non-exporter category.

⁷Table 3 in Bernard et al. (2007) states that the difference in average log shipments between exporters and non-exporters is 1.48 for the same set of firms we used. The finding that exporters are larger than non-exporters has also been confirmed for plants and establishments, as well as for other countries. The size differences are all in the range between 4 and 6. In this section, we work with total sales rather than log total sales so we can use the information in Table 1 without any further assumptions.

are, $s_x = .35$. The average total sales of exporters is then very close to the data. Non-exporters, though, are then very small and the size premium is still 15 times larger than in the data.

| | Data | Strict Sorting | | |
|-------------------------------|----------------|----------------|----------------|----------------|
| Share of exporters s_x | .18 | .15 | .25 | .35 |
| Average Sales - Non-exporters | \$8.1 million | \$746,000 | \$637,000 | \$491,000 |
| Average Sales - Exporters | \$36.4 million | \$83.9 million | \$51.0 million | \$36.8 million |
| Exporter Size Premium | 4.5 | 114 | 89 | 75 |

Table 3: **Exporters and non-exporters given the share of exporters s_x**

The lack of fit due to strict sorting can also be understood in terms of the conditional probability of exporting given the firm size. In the basic Melitz model, the conditional probability is a step function of firm size, taking value 0 if r_i is below the cutoff t , and one otherwise. Larger firms are indeed more likely to export in the data, but only a little bit more. In other words, the conditional probability rises only slowly with firm size in the data, while it rises very sharply in the model. We find the differences in the conditional probability between the model and the data are somewhat harder to summarize and interpret in economic terms. In addition, any measure of fit would depend on the level of detail available in the data. Computing the exporter size premium in both model and data is easier, plus the differences are readily stated in economic terms.

We also experimented with a parameterized Pareto distribution for the total sales distribution. If Ψ_r is Pareto, we have that the cutoff t is given by $t = (s_x)^{-1/k}$ where k is the slope parameter of the Pareto distribution. The unconditional average firm size μ can be decomposed in the conditional averages

$$\mu = s_x \mu_x + (1 - s_x) \mu_{nx}$$

where μ_{nx} and μ_x are the average total sales for exporter and non-exporters respectively. For a Pareto distribution, μ_x/μ is equal to the cutoff t and we have that

$$\frac{\mu_{nx}}{\mu} = \frac{1 - (s_x)^{1-1/k}}{1 - s_x}.$$

Estimates of the slope parameter in Axtell (2001) and Luttmer (2007) are very close to one. It is thus clear that the term $(s_x)^{1-1/k}$ is very close to one and μ_{nx}/μ is bound to be very small. Axtell (2001) works with $k = 1.024$, which implies an exporter size premium of 110 under strict sorting. On the upper end of the estimates, a slope parameter $k = 1.065$ brings the exporter size premium down to 41. It must be emphasized that, while the Pareto distribution has been documented to fit the tail of firm-size distribution, it is not a good approximation of the size distribution of exporters in the data for the very reason that

strict sorting overstates the size premium: exporters are *not* concentrated in the tail of the firm-size distribution. This lack of fit with the empirical distribution explains why a Pareto distribution can imply a lower size premium than our non-parametric lower bound.

3 Introducing Latent Heterogeneity

Of course, we did not expect the strict sorting of exporters to hold *exactly* in the data. Melitz (2003) certainly does not intend to preclude the importance of other idiosyncratic factors, unrelated to size, in the firm’s decision to export. For example, the costs of exporting—variable or fixed—may vary from firm to firm. It is possible that a low-productivity firm finds exporting profitable simply because it is quite cheap to transport its product. Alternatively we could think that whatever makes a product successful in the domestic market does not always translate into more sales in the foreign market. Clearly, once we consider these additional factors, the strict sorting of exporters no longer holds.

3.1 Fitting the data

Here we proceed to reconcile the model with the data by introducing the necessary firm-level heterogeneity in the export decision. We rewrite the threshold condition (2) in terms of a latent variable, t_i . A firm i exports if its total sales satisfy

$$r_i \geq t_i. \tag{3}$$

The latent variable is identically and independently distributed across firms, with a c.d.f. Ψ_t over support \mathfrak{R}_+ .

With the latent variable condition (3) we can capture all of the underlying heterogeneity without having to specify the sources of variation. Indeed, the only structure imposed on the latent variable is that it is independent of firm size. As in the previous section, we can postpone laying out and solving the structural model while seeing through the model’s implications given the data.

It is now useful to fit the empirical distribution of total sales with a parametric distribution. We use a lognormal distribution with mean $\mu_r = 6.3$ and standard deviation $\sigma_r = 2.6$ —so we reproduce the average total sales (in thousands of dollars) as well as the approximated location of the 82th percentile. The use of the lognormal distribution for firm size has a long tradition in economics. See Aitchison and Brown (1969) for a complete treatment of the distribution.

We also assume that the latent variable t_i follows a lognormal distribution. This is mainly a choice of convenience: we want a flexible two-parameter distribution defined over positive numbers. We pick the mean μ_t and standard deviation σ_t such that the model reproduces

the share and size of exporters. That is, we solve for μ_t and σ_t such that equations

$$s_x = \int \Psi_t(r) d\Psi_r(r), \quad (4)$$

$$E \{\log(r_i) | r_i \geq t_i\} = \int \log(r) \Psi_t(r) d\Psi_r(r) / s_x, \quad (5)$$

reproduce the observed values for the share of exporters, $s_x = .18$, and the average log total sales of exporters, $E \{\log(r_i) | r_i \geq t_i\} = 5.66$.⁸ We use Ψ_t to denote the c.d.f. of the latent variable distribution.

The latent variable is very dispersed: we find that the parameter values matching moments (4)-(5) are $\mu_t = 13.73$ and $\sigma_t = 7.7$. These parameters imply that the coefficient of variation of the latent variable is many orders of magnitude larger than for total sales! We explored an array of parameters for the distribution of total sales and found always that we need a huge dispersion for the latent variable in order to reproduce the share and size of exporters. More precisely, we consider values μ_r in the range 5.5 – 7.5 and σ_r in the range 2 – 3. The resulting parameters for μ_t and σ_t were always above 10 and 5, respectively.

There also good reasons to view our findings as a lower bound on the dispersion of the latent variable. Firms often carry multiple product lines, and there are many foreign markets to serve. However, it only takes one product to be exported to one destination for a firm to be called an exporter. Hence, if we think each market/product offers an independent opportunity to export, we have to see the latent variable t_i as the *minimum* realization among the ensemble of product- and destination-specific thresholds. Thus, the underlying distribution from which the product- and destination-specific latent variables are drawn would feature much more dispersion.

The results are perhaps not that surprising: after all, we have to make up for a large gap between the model and the data in the size premium of exporters. In order to reduce the size of exporters we need the latent variable to take very large values with high probability, so some large firms do not export. Simultaneously, some other firms must draw a low realization of the latent variable and export independently of their size.⁹

⁸Here we use the log total sales instead of total sales, so the exporter size premium is now given by the difference in average log total sales between exporters and non-exporters. The change of units has virtually no implication for the parameters of the latent variable distribution—but it turns out to be very convenient for the calibration of the model later.

⁹It is also necessary to introduce latent heterogeneity in a Melitz model in order to match the *dynamic* facts. Entry and exit rates in foreign markets are relatively high, as documented in Bernard, Jensen and Schott (2007). The volatility of firm employment in the data is clearly too low to explain these high rates in a basic Melitz model: firms rarely grow or contract enough to start and stop exporting so frequently. See Atkeson and Burstein (2008) for a discussion.

3.2 The role of firm size

The huge dispersion of the latent variable indicates that size plays only a small role in the determination of the export status of a firm. We can illustrate this point with a simple exercise. Note that given a firm size r_i , the probability that firm i exports is

$$\begin{aligned} \Pr(i \text{ exports} | r_i) &= \Pr(r_i \geq t_i | r_i) \\ &= \Psi_t(r_i). \end{aligned}$$

Given our results for Ψ_t , a firm of median size exports with probability .167, very close to the unconditional probability of .18.¹⁰ In other words, a firm of median size could be taken as representative of the industry as a whole. In contrast, a firm with the median latent variable will export only with probability .0016. Table 4 repeats the exercise with the 25th and 75th percentiles for total sales and the latent variable.

| <i>Ordered</i> | Percentile | | |
|---------------------------|-------------|-------------|-------------|
| | <i>25th</i> | <i>50th</i> | <i>75th</i> |
| <i>By total sales</i> | .1179 | .1674 | .2286 |
| <i>By latent variable</i> | 10^{-7} | .0016 | .1881 |

Table 4: **Conditional probability of exporting**

In the next section we will explore the implications of the additional heterogeneity in the context of a structural model. Here we anticipate our main result with a simple exercise that illustrates how the predictions for the extensive margin crucially depend on whether we match the exporter size in the data.

The exercise is as follows: we increase the total sales of *all firms* by a constant percentage δ . We then look at the predictions for the growth rate of the number of exporters for both models: the strict sorting model given by (2) and the latent variable model given by (3). It must be emphasized that there is no change either in the threshold or the distribution of the latent variable.

Table 5 reports the findings for the growth rate in the number of exporters for revenue increases of 5%, 10% and 20%. The strict sorting model predicts that a relatively large fraction of firms change status from non-exporter to exporter. For a 10% increase in total sales the model predicts a 5.29% growth rate in exporters. The distribution of total sales is still pretty thick just below the 82nd percentile. Because size is the only determinant of the export status of a firm, all the firms whose total log sales were originally between $\log(t) - \log(\delta)$ and $\log(t)$ will export after the revenue increase.

In contrast the latent variable model predicts much smaller growth rates for the number of exporters. For a 10% increase in total revenue, only 2.23% of firms switch export status.

¹⁰The median size is given by $\exp(\mu_r)$ or approximately \$600,000.

| <i>Model</i> | Revenue increase δ | | |
|------------------------|---------------------------|-----------------|-----------------|
| | $\delta = 5\%$ | $\delta = 10\%$ | $\delta = 20\%$ |
| <i>Strict Sorting</i> | 2.69% | 5.29% | 10.26% |
| <i>Latent Variable</i> | 0.67% | 2.23% | 3.67% |

Table 5: **Growth rate of the number of exporters**

In Figure 1 we have plotted the distribution (PDF) of *normalized* log total sales: $\log(r_i) - \log(t_i)$ for both models. Thus, for the strict sorting model we just have the distribution of log total sales, centered such that the 82nd percentile is equal to zero. The resulting distribution is plotted as a dashed line. For the latent heterogeneity, normalized log total sales are way more dispersed, reflecting the huge variation in the latent variable. The distribution is plotted in a solid line.

The key observation in Figure 1 is that the mass of firms close to the threshold is much larger under strict sorting than in the latent heterogeneity model. Thus, any displacement of the distribution to the right brings more firms above the threshold in the strict sorting, as documented in Table 5. In the latent heterogeneity, the gain in size must be quite large in order to overcome the other determinants of the exporting decision.

4 A Simple Framework of Exports and Exporters

In this section we set up a small model with economies of scale and firm heterogeneity. The model is simpler than Melitz (2003) in that it is a partial equilibrium model, taking the wage rate as given. As a result the model abstracts from entry and exit in the domestic market.

We explore the implications of the additional heterogeneity for an episode of trade liberalization. To this end we consider two versions of the model. In the first, there is strict sorting, so the model does not match the size of exporters. In the second version, we introduce the needed variation as a random fixed cost. The rest of the parameters are common across models. Finally, we briefly sketch three simpler models for comparison purposes.

4.1 Framework

There is a set Ω of firms that produce and sell in the home country. Firms are heterogeneous in their productivity, denoted φ , and the fixed costs they face if they start exporting, denoted f . Productivity and fixed costs are independently distributed over \mathfrak{R}_+ with c.d.f. G and H , respectively. We summarize a firm by its type $\omega = \{\varphi, f\}$. The set of firms Ω (and their distribution) is taken as a given.

Each firm is the single producer of a differentiated good with technology

$$y(\omega) = \varphi(\omega) l(\omega)$$

where $l(\omega)$ is the labor demanded by firm ω . Consumers in the home country aggregate the differentiated goods according to

$$Y^d = \left[\int_{\Omega} (y^d(\omega))^{\rho} d\omega \right]^{1/\rho}$$

where $\rho \in (0, 1)$ and $y^d(\omega)$ denote the output of firm ω sold in the home country. The demand for each good ω is given by

$$y^d(\omega) = (p^d(\omega) / P^d)^{-\theta} Y^d$$

where $\theta = (1 - \rho)^{-1}$ is the price elasticity, $p^d(\omega)$ the price set by firm ω , and the price index P^d is given by

$$P^d = \left[\int_{\Omega} (p^d(\omega))^{1-\theta} d\omega \right]^{\frac{1}{1-\theta}}.$$

Firms are monopolistic competitors and internalize the downward sloping demand for their products. The profit-maximization problem leads to the familiar price equation

$$p^d(\omega) = \frac{1}{\rho} \frac{w}{\varphi(\omega)}$$

where w is the wage rate. We take the wage rate as exogenously given, so ours is a partial equilibrium model.

It is clear that only the productivity parameter φ will determine domestic sales. We can thus write $p^d(\varphi)$ and $y^d(\varphi)$. The c.e.s. demand structure also implies that firm φ revenues from domestic sales can be expressed

$$r^d(\varphi) = \left(\frac{\varphi}{\tilde{\varphi}} \right)^{\theta-1} R^d \tag{6}$$

where R^d are total sales revenues in the domestic market, and

$$\tilde{\varphi} = \left[\int_0^{\infty} \varphi^{\theta-1} dG(\varphi) \right]^{\frac{1}{\theta-1}}$$

is the average productivity defined as in Hopenhayn (1992) and Melitz (2003). Since we take both the wage and the distribution of firms in the home country as given, total domestic sales R^d are also exogenously determined.¹¹ We will still make use of the relationship between productivity and domestic sales given by (6).

¹¹Briefly, $R^d = r^d(\tilde{\varphi})$ by (6), and $P^d = w / (\tilde{\varphi}\rho)$ by substituting the price for each good, $p^d(\varphi)$. We do not model the import decision of the domestic households.

We now move to the determination of exports and exporters. Not all firms export: let Ω_x denote the set of firms that do and M_x its measure. We normalize the measure of all firms to one, so M_x is also the share of exporters. Consumers in the foreign country aggregate the subset of exported goods according to

$$Y^f = \left[\int_{\Omega_x} (y^f(\omega))^\rho d\omega \right]^{1/\rho} \quad (7)$$

where $y^f(\omega)$ is the output of firm ω sold in the foreign country. The foreign demand for exporter good ω is given

$$y^f(\omega) = (p^f(\omega) / P^f)^{-\theta} Y^f$$

where

$$P^f = \left[\int_{\Omega_x} (p^f(\omega))^{1-\theta} d\omega \right]^{\frac{1}{1-\theta}}. \quad (8)$$

Finally, we assume there is an aggregate demand for exports, given by

$$Y^f = Y^* (P^f)^{-\nu}, \quad (9)$$

where Y^* is the (exogenously given) income of the foreign country, and ν is the price elasticity of aggregate exports of the home country. We assume that $\nu < \theta$; that is, exports of the home country are closer substitutes of each other than they are of a good produced elsewhere.

Let us first solve for the export revenues of a firm, taking as given the set of exporters Ω_x . Profit maximization implies that

$$p^f(\omega) = \frac{1}{\rho} \frac{\tau w}{\varphi(\omega)}$$

where $\tau > 1$ is an iceberg trade cost associated with exports. It is clear again that, conditional on exporting, only the productivity parameter φ determines sales.

The c.e.s demand system allows us to write a firm's export revenues as a function of the average export revenues within exporters,

$$r^f(\varphi) = \left(\frac{\varphi}{\tilde{\varphi}_x} \right)^{\theta-1} \frac{R^f}{M_x} \quad (10)$$

where R^f is total export sales and

$$\tilde{\varphi}_x = \left[\frac{1}{M_x} \int_{\Omega_x} (\varphi(\omega))^{\theta-1} d\omega \right]^{\frac{1}{\theta-1}}$$

is the average productivity among exporters. Note that the set of exporters Ω_x affects the export revenues of each firm (10), both through the share of exporters M_x and the productivity distribution within the set.

Last but most certainly not least, we get to the determination of the set of exporters Ω_x . A firm that exports incurs in a per period fixed cost. As a result, a firm ω will find it profitable to export only if its net income abroad would cover the fixed expenses,

$$\frac{1}{\theta} r^f(\varphi(\omega)) \geq f(\omega), \quad (11)$$

where export net income, that is, export revenues minus costs, is expressed as $r^f(\varphi(\omega))/\theta$. Thus the set of exporters Ω_x is the set of firms $\omega \in \Omega$ such that the entry condition (11) holds. However, export revenues depend themselves on the set of exporters, so in equilibrium exports and exporters are determined simultaneously.

4.2 Strict sorting and latent heterogeneity

We now consider two versions of the model above. In the first we stick to the basic Melitz model and shut down the heterogeneity in the fixed costs. As a result exporters and non-exporters are strictly sorted by size and the model inherits the inability to match the size of exporters as documented in Section 2. We name this first version of the model after the strict sorting property. In the second calibration we use the dispersion on fixed costs to reproduce the latent variable distribution we worked out in Section 3. By construction the model then matches the share and size of exporters. We label this calibration as the “latent heterogeneity” model.

It must be emphasized that the *only* difference between the two models is the distribution of fixed costs. The models will share the same parameter values for the elasticities, trade costs, and the distribution of productivity. We will also set the exogenous variables such that both models match the same facts in the data.

We start with the common parameters. We set the elasticity of substitution across exported goods to $\theta = 8$. This number is essentially in the middle of the range of estimates surveyed by Anderson and Van Wincoop (2004). For the aggregate demand for home exports (9) we set the price elasticity of $\nu = 6$. Our baseline trade costs are set at 50%, $\tau = 1.5$. This is the midpoint between the estimated trade costs in 1987 and 2002 reported by Alessandria and Choi (2008), and slightly below the findings in Anderson and Van Wincoop (2004). The levels of the foreign income and the wage rate are set to reproduce the ratio of total exports to total sales in the industry.

The last common parameter is the distribution of productivity, G . For the calibration we want to use the observed distribution for total sales. Unfortunately the mapping from the productivity distribution to the total sales distribution is not the same for both versions of the model.¹² It is thus not possible to have a common parametrization for G that matches the observed distribution of total sales in both models. However, we do not want any difference in the calibration of the models to govern the results—other than the size of exporters, that

¹²Both models share the mapping from G to the distribution of *domestic* total sales. The distribution of export sales, though, are different for each model.

is. We thus adopt the following compromise: we set the productivity distribution to capture *all* of the variation in total sales. More precisely, we take G to be a lognormal distribution with standard deviation equal to $\sigma_\varphi = \sigma_r / (\theta - 1)$. The location parameter can be set such that the average productivity $\tilde{\varphi}$ among all producers is normalized to one.

Table 6 summarizes all of the parameters common to both models.

| Parameter | | Value |
|----------------------------|------------------|-------|
| Elasticity of Substitution | θ | 8 |
| Price-Elasticity exports | ν | 6 |
| Trade Cost | τ | 1.5 |
| Std.Dev. Log-productivity | σ_φ | .37 |
| Mean Log-productivity | μ_φ | -0.48 |

Table 6: **Calibration - common parameters**

Finally we get to what makes the strict sorting and latent heterogeneity models different. For the strict sorting we shut down all of the variation in the fixed cost, so H is a degenerate distribution at some point f . The value of the fixed cost f is set such that the share of exporters is 18% as in the data. For the latent heterogeneity model we calibrate the distribution of fixed costs H such that we capture all the variation in the latent variable t_i . For this it is very convenient to use a log-normal distribution for fixed costs so the mapping is simple. By construction the latent heterogeneity model replicates the observed exporter size premium; by appropriate choice of the location parameter, we match the share of exporters as well. The resulting parameters are $\mu_f = .98$ and $\sigma_f = 7.69$.

4.3 Three analytical models

While our analysis is based on the two models just described, we find it useful to consider three simpler models for comparison purposes: a representative firm model, a homogeneous firm model, and a heterogeneous-productivity model with a Pareto distribution. In these three models we can derive the relationship between total exports and trade costs in closed form. This helps build an intuition about the mechanisms at play in the richer full model.

Representative firm model. There is a representative firm that exports under all circumstances. In terms of the framework above, we encompass the representative firm model by having a degenerate distribution of firm productivity and a fixed cost equal to zero. There are, thus, no changes in the measure of firms exporting or the average productivity of exporters. We express the price of exports (8) in terms of logs,

$$\log(P_f) = \log(\tau) + \text{const.}$$

where all terms that are constant are collected in *const*. Substituting in the aggregate demand for exports (9) we have the simplest model of exports,

$$\log(R_f) = (1 - \nu) \log(\tau) + \text{const.} \quad (12)$$

We maintain the same value for the price elasticity ν given in Table 6.

Homogeneous firm model. We introduce entry by assuming a positive fixed cost but abstract from all heterogeneity: all firms have identical labor productivity and face the same fixed cost of exporting. Some firms, though, export, while some others do not. This can be an equilibrium only if firms are actually indifferent between exporting or not, so the entry condition (11) holds with equality. Since there is no firm heterogeneity, the entry condition (11) effectively pins down the average export sales,

$$\log(R_f/M_x) = \text{const.}$$

Thus all of the adjustment in this model occurs through entry, since firms do not change their export intensity in response to trade costs.

Since all firms have identical productivity, there is no change in the average productivity of exporters either. The export price (8) is thus a function of the trade costs and the share of exporters,

$$\log(P_f) = \frac{-1}{\theta - 1} \log(M_x) + \log(\tau) + \text{const.}$$

We substitute the export price in the aggregate demand for exports (9) to obtain

$$\log(R_f) = \frac{\nu - 1}{\theta - 1} \log(M_x) + (1 - \nu) \log(\tau) + \text{const.}$$

Finally, we use that $\log(R_f) = \log(M_x) + \text{const}$ to obtain total export revenues as a function of the trade cost,

$$\log(R_f) = -\frac{(\theta - 1)(\nu - 1)}{\theta - \nu} \log(\tau) + \text{const.} \quad (13)$$

The assumption that $\theta > \nu$ guarantees that the export revenues increase with a fall in trade costs, as we would expect. We maintain the same value for parameters ν and θ given in Table 6.

Heterogeneous-productivity firms with a Pareto distribution. The last of our models features heterogeneous firms and an identical fixed cost of exporting. The only difference with the strict sorting model discussed above is that productivity follows a Pareto distribution. This allows us to solve for the model analytically but has stark implications for the margin of adjustment, as will be clear very soon.

We now have to simultaneously solve for the export demand (9) and the entry condition (11). We start with the latter. Assuming that there is a subset of firms that do not export, the

entry condition must hold with equality.¹³ The firm's export revenues are strictly increasing in productivity, φ , so we can characterize the set of exporters with a simple threshold rule. Let $\hat{\varphi}_x$ hold

$$\frac{1}{\theta} r^f(\hat{\varphi}_x) = f.$$

The set of exporters Ω_x is then given by the set of firms $\varphi \geq \hat{\varphi}_x$ and thus $M_x = 1 - G(\hat{\varphi}_x)$. Using the c.d.f. for the Pareto distribution we have that

$$\log(M_x) = -k \log(\hat{\varphi}_x) \quad (14)$$

where k is the slope parameter associated with the Pareto distribution.

We now get to use the key property of the Pareto distribution that allows for an analytical solution. It is easy to show that if φ is distributed according to a Pareto distribution with slope k , then we have a linear relationship between $\tilde{\varphi}_x$ and $\hat{\varphi}_x$,

$$\tilde{\varphi}_x = \left[\int_{\varphi \geq \hat{\varphi}_x} \varphi^{\theta-1} \frac{dG(\varphi)}{1 - G(\hat{\varphi}_x)} \right]^{\frac{1}{\theta-1}} = \left[\frac{k}{k+1-\theta} \right] \hat{\varphi}_x.$$

We write the export revenues of firm $\hat{\varphi}_x$ in terms of the average export revenues using (10),

$$\left(\frac{\hat{\varphi}_x}{\tilde{\varphi}_x} \right)^{\theta-1} \frac{R_f}{M_x} = \theta f$$

and it is immediate that the average export sales are pinned down by the entry condition,

$$\log(R_f/M_x) = \text{const.}$$

Thus average export sales do not respond to trade costs, so the growth rate in total revenues equals the growth rate in the number of exporters.¹⁴ This stark implication does not hold for other productivity distributions.

We are set to solve the model. The export price (8) is

$$\log(P_f) = \frac{-1}{\theta-1} \log(M_x) + \log(\tau) - \log(\tilde{\varphi}_x) + \text{const.}$$

Using the relationships between $\hat{\varphi}_x$ and $\tilde{\varphi}_x$ as well as (14),

$$\log(P_f) = \left[-\frac{1}{\theta-1} + \frac{1}{k} \right] \log(M_x) + \log(\tau) + \text{const.}$$

Now we substitute in the aggregate demand for exports to obtain

$$\log(R_f) = - \left[\frac{1}{\nu-1} - \frac{1}{\theta-1} + \frac{1}{k} \right]^{-1} \log(\tau) + \text{const} \quad (15)$$

where the average exports are constant.

¹³Since the support for the Pareto distribution is unbounded, there is always a firm that finds it profitable to export.

¹⁴Note that the export sales per firm increase with a fall in trade costs; however, new exporters are smaller than incumbents and drive the average down.

5 Trade Liberalization

We compare the two models' predictions for export growth in response to a fall in trade costs. The exercise is labeled as a “trade liberalization,” but there is no distinction in our model as to whether it is a tariff reduction or an improvement in the shipping technology.¹⁵ We consider a range of trade liberalizations up to a halving of the trade costs—a fall of 25 percentage points.

5.1 Aggregate response

We find that both models have virtually identical implications for trade volume. Figure 2 plots the growth rate of different variables (as a percentage rate) as a function of the fall in trade costs (in percentage points). The solid and dashed lines correspond to the strict sorting and latent heterogeneity models, respectively. The top left panel displays the growth rate of exports. Both models predict that trade will approximately triple once trade costs are cut in half. The models' predictions are so similar that the lines are on top of each other for most of the range of trade costs. The top right panel in Figure 2 displays the growth rate in the export price index, as given by (8). Again there are no differences between both models, and the price index falls with trade costs at the same rate for both models. Hence the profile for export revenues is very similar (bottom left panel). Finally, the bottom right panel in Figure 2 displays total employment in the industry. The trade liberalization results in a sizeable expansion of employment, about 10% for the largest fall in trade costs considered.

Why do the two models predict such similar aggregate trade patterns? The reason is that, in this respect, both models are very similar to the representative firm model presented in Section 4.3. This is quite obvious in Figure 3, where we have added the representative firm model, indicated with a dotted line. We plot export growth (top) and the export price growth (bottom) as a function of the fall in trade costs. Clearly, all three models are very similar in both prices and quantities.

It may be puzzling at first that firm-level heterogeneity and the extensive margin do not make a difference for aggregate variables. The intuition for this is that firm-level heterogeneity essentially cancels the boost in trade due to the extensive margin. We can illustrate this by comparing first the representative firm model with the homogeneous firm model and then with the heterogenous-productivity firm model with a Pareto distribution.

Endogenous entry amplifies the response of export revenues to a fall in trade costs. A simple comparison of the representative firm model (12) and the homogeneous firm model (13) backs the claim: the elasticity of export revenues with respect to trade costs is augmented by $(\theta - 1)/(\theta - \nu)$, which is always greater than one. Everything else constant, more exporters bring the aggregate price of exports down through a love-of-variety effect. As long as there is entry, the price of exports will then fall by more than one-to-one with trade costs and sustain additional demand for exports. For the choice of parameters reported in Table

¹⁵Or, for that matter, an exogenous change in the real exchange rate.

6, the response of trade volume is more than twice as large in the homogeneous firm model than in the representative firm model.

We now look at the heterogeneous-productivity firm model with a Pareto distribution to see how the trade response changes with firm heterogeneity. Comparing (15) and (13) it is clear that the response in trade in the heterogeneous firm model is between those in the representative firm model and the homogeneous firm model. Which model the heterogeneous firm model resembles the most depends on the slope parameter of the Pareto distribution, k . As k grows large, the heterogeneous model approaches the homogeneous model and the export growth is the largest. This is not surprising, since the Pareto distribution becomes degenerate, with all of the mass at the location parameter.

For low slope parameters k the Pareto distribution features a thick tail. In this case the response of exports to a fall in trade costs is quite muted. The firms that start exporting in response to the fall in trade costs are less productive than incumbent exporters. Thus the average productivity among exporters falls rapidly, which in turn drives the aggregate price of exports. The more skewed is the distribution of productivity, the faster the average productivity drops with entry.

Quantitatively, we find that the heterogeneity cancels virtually *all* of the amplification introduced by entry, rendering the heterogeneous and the representative firm models very similar in their implications for export growth. We base this assertion on the skewness of the empirical distribution of total sales. In the model the right tail of total sales is distributed according to a Pareto with slope parameter $k/(\theta - 1)$. The right tail in the observed distribution of total sales is well approximated by a Pareto with a slope parameter ξ in the range 1.02 – 1.06. Setting $k = \xi(\theta - 1)$ we can rewrite the response of exports in the heterogeneous firm model (15) as

$$\log(R_f) = - \left[\frac{1}{\nu - 1} - \frac{1}{\theta - 1} \frac{\xi - 1}{\xi} \right]^{-1} \log(\tau) + \text{const.}$$

From the above expression it is clear that if ξ is close to one, then the term $(\xi - 1)/\xi$ is approximately zero and the heterogeneous firm model is very close to the representative firm model (12).

5.2 Margins of adjustment

Behind the similarities in the response of aggregate trade, the models display marked differences in the adjustment in the number of exporters. The top panel in Figure 4 plots the growth rate in the number of exporters for each model. In the strict sorting model the number of exporters grows very fast. Entry drives up the number of exporters by almost 60% when trade costs fall by 25 percentage points. Even with a small drop in trade costs, such as 5 percentage points, the number of exporters grows by more than 10%, suggesting that even at relatively high frequencies entry could play an important role.

In contrast, there is not much entry in the foreign market in the latent heterogeneity model. The growth rate of the number of exporters barely gets over 15%, one fourth of the growth in the strict sorting model. For smaller trade cost reductions such as 5%, the number of exporters is very close to flat. The similarities for total trade volume only make the differences in entry even more striking.

We seek to quantify further the role of export entry in both models. We decompose the growth rate of exports in the change of export intensity and participation,

$$\Delta \log(Y^f) = \Delta \log(r^f) + \Delta \log(M_x). \quad (16)$$

The first term on the right-hand side is the intensive margin, that is, the growth rate on the average export revenues per exporter; and the second is the extensive margin, or the growth in the number of exporters. We normalize export growth by total sales and express each margin as a percentage of the total.¹⁶ Table 7 collects the results for trade cost reductions of 5, 15, and 25 percentage points. In the strict sorting model the extensive margin is more than 60% of the growth rate in exports. However, entry plays a much more minor role in the latent heterogeneity model, just below 20%. These numbers are very similar across the range of trade cost cuts. They are also quite robust to alternative parameterizations for elasticities. This leads us to the conclusion that the Melitz model is at odds with a large role for the extensive margin, once it is augmented to account for the share and size of exporters.

| <i>Model</i> | Trade Cost Decrease | | |
|-----------------------------|---------------------|-------|-------|
| | 5 | 15 | 25 |
| <i>Strict Sorting</i> | 63.3% | 63.3% | 63.5% |
| <i>Latent Heterogeneity</i> | 19.2% | 19.8% | 20.4% |

Table 7: **The role of the extensive margin.** *Trade Costs decrease in percentage points.*

The reader may be wondering how the two models can have such different entry rates and yet very similar export growth. The bottom panel of Figure 4 clues us in. Despite having much lower entry rates, total employment by exporters actually grows much faster in the latent heterogeneity model. The reason is that new exporters are very different in the two models. In the strict sorting model new exporters are very small in comparison with the incumbent exporters, about 10 times smaller than the average exporter size prior to trade liberalization.¹⁷ This is a direct consequence of the strict sorting: since exporters are almost 100 times larger than non-exporters, the firm at the threshold is still quite large compared with non-exporters but, more important, it is very small compared with exporters. Returning to Table 2 in Section 2, the *smallest* exporter has about 5 times more total sales than the average non-exporter, but 18 times less total sales than the average exporter.

¹⁶We are following Alessandria and Choi (2008) closely, although we do not distinguish between export intensity and premium as they do.

¹⁷Because previous exporters grow rapidly with the trade liberalization, the new exporters are 50 times smaller when compared with the average size of the incumbent exporter after trade costs decrease.

In contrast, new exporters in the latent heterogeneity model are still smaller than existing exporters, but not by much. For a fall in trade costs of 5 percentage points, new exporters are just 40% smaller than the average exporter size prior to the trade liberalization. If we go all the way to a halving of trade costs, new exporters are actually 30% larger than the average exporter size prior to the trade liberalization. In short, new exporters are not very different from the average firm in the industry.

We now evaluate the main mechanism in Melitz (2003), namely, the gains in average productivity in the industry due to the trade liberalization. We follow Melitz (2003) and define

$$\hat{\varphi} = \left(\tilde{\varphi}^{\theta-1} + M_x (\tilde{\varphi}_x/\tau)^{\theta-1} \right)^{\frac{1}{\theta-1}}$$

as the aggregate productivity.¹⁸ It must be noted that our model can only do a partial evaluation. Since we take the wage rate as given, the set of non-exporters does not change so there are no productivity gains from the exit of the least productive firms as in Melitz (2003).

We find that aggregate productivity grows significantly less in the latent heterogeneity model than in the strict sorting model. Figure 5 plots the results. Aggregate productivity growth in the latent heterogeneity model is about half the one in the strict sorting model—a ratio that is approximately constant across the range of trade cost decreases.

That productivity gains are smaller is not surprising given our previous results. Given that we abstract from exit in the domestic market, it can be said that the mechanism in Melitz (2003) works through the selection and entry of exporters. First, exporters are more productive than the average firm, so anything that expands the total employment of exporters will lead to productivity gains due to composition effects. Second, new exporters experience a large jump on their output, and since they are still more productive than most firms in the economy, also induce gains in aggregate productivity.

Both selection and entry are much weaker in the latent heterogeneity model than in the strict sorting model. We have seen that strict sorting greatly overpredicts the size of exporters, that is, their productivity. The latent heterogeneity reconciles the model and data: exporters are still more productive than non-exporters, but only modestly so. Moreover, there is much less entry in the latent heterogeneity model, so the second source of productivity gains is weaker too.

We conclude our analysis with a look at how the average exporter changes with trade liberalization. Figure 6 displays the average among exporters for export revenues, export output, average productivity, and total employment. It must be emphasized that the set of exporters changes as we cut the trade costs. First, we note that export revenues and output grow much faster in the latent heterogeneity model. This is, of course, the flip side of the results on the extensive margin documented in Table 7. Second, the average productivity for exporters falls in both models but by different amounts, as shown in the bottom left panel in

¹⁸Melitz (2003) refers to this as combined average productivity and shows that it completely summarizes the effects of the distribution of productivity levels on the aggregate outcome.

Figure 6. In both models new exporters are less productive than incumbent exporters but in the strict sorting model, they are much less so. As a result the exporter average productivity drops by a staggering 6% but only a modest 1% in the latent heterogeneity model.

The differences in the average output and productivity of exporters combine for opposite predictions with respect to total employment for exporter. In the strict sorting model exporters, on average, employ fewer workers as trade costs fall. New exporters do employ more workers than they did before entering the foreign market. However, they are so small compared with the incumbent exporters that they bring the average down by an astounding 25%. Recall that the number of exporters grows by more than 50% so entry has a big impact on averages. In contrast, total employment for exporters grows in the latent heterogeneity model, as the weaker selection and entry effects cannot overturn the employment gains due to overall expansion of exports.

5.3 Robustness

We briefly discuss here alternative specifications for the common structure of both models. We start with our choice of the lognormal distribution for firm productivity. In particular, theory work has favored instead the Pareto distribution for its tractability.¹⁹ It has been argued that the Pareto distribution is a very good approximation of the tail of the firm-size distribution.²⁰ However, the Pareto distribution proves to be a very restrictive choice for our purposes. As we discussed in Section 4.3, a Pareto distribution for firm productivity implies that average export sales per exporter are constant. Thus, as trade costs fall, all of the adjustment must occur through the extensive margin, that is, the number of exporters.²¹ This exclusive role of the extensive margin is clearly at odds with the data. Instead, we favor a more flexible specification that allows for both margins to be active.

| | Baseline | | $\theta = 12$ | | $\nu = 3$ | |
|---------------------|----------|------|---------------|------|-----------|------|
| | SS | LH | SS | LH | SS | LH |
| Total Exports | 52.8 | 55.8 | 52.3 | 54.1 | 23.3 | 23.7 |
| Number of Exporters | 20.3 | 6.1 | 20.1 | 6.0 | 7.7 | 2.4 |
| Agg. Productivity | 0.5 | 0.2 | 0.1 | 0 | 0.5 | 0.2 |

Table 8: **Robustness analysis.** *All variables are growth percentage rates after a fall of 10 percentage points in trade.*

Finally, we also explore some alternative parameter specifications. Table 8 reports the growth rate for exports, the number of exporters, and aggregate productivity in the event of

¹⁹See Chaney (2008), for example.

²⁰See Axtell (2001) and Luttmer (2007).

²¹It is important to recall that we have abstracted from general equilibrium effects on the wage rate, which, in turn, could have affected the fixed costs associated with exporting. This channel is operative in Melitz (2003) and leads to adjustment along the intensive margin.

a trade cost reduction of 10 percentage points for both the strict sorting (SS) and the latent heterogeneity (LH) model. In addition to the baseline calibration, we consider two alternative parameterizations. In the first, we set the elasticity of substitution between exports to 12. There is no significant change for both models' predictions for the growth rate of exports and exporters, but there is now virtually no aggregate productivity growth. This shows that the “love of variety” effect is the main driver of aggregate productivity. We also look at a second parametrization with a very low price elasticity of aggregate exports. Naturally, total trade growth is smaller as demand does not react to the fall in costs. Interestingly aggregate productivity displays similar gains.

6 Some Theories of Heterogeneity

6.1 Small exporters

Bernard, Jensen and Schott (2007) and Eaton, Kortum, and Kramarz (2004) report that a large fraction of exporters actually sell very little abroad. These small exporters are at odds with the basic Melitz model. As we discussed, the smallest exporter should be larger than more than 80% of the firms. Assuming that foreign sales are proportional to total sales among exporters, we would find that the smallest exporter should have foreign sales totalling approximately \$1 million.²² Instead, Bernard, Jensen and Schott (2007) report that more than 75% of all exporters sell less than \$1 million abroad.

Arkolakis (2008) develops a theory of market penetration that effectively introduces convex market access costs for foreign sales. Exporters with relatively low productivity sell very little abroad as the cost of reaching more consumers exceeds marginal revenue. Larger firms, capable of more net revenue per sale, serve most of the consumers abroad and thus behave (approximately) as in the standard Melitz model. Thus, export intensity is increasing with size. We note that while the model in Arkolakis (2008) can explain the presence of exporters with small exports, it does not explain why some small firms export and some large firms do not. In other words, the strict sorting in total sales between exporters and non-exporters is preserved.

Our first observation is that, as long as the strict property is preserved, variation in export intensity cannot reconcile the model with the data. The simplest way to see this is to compute the exporter size premium in terms of domestic sales under strict sorting, following the same procedure we did with total sales. We report the results in Table 9. Excluding foreign sales does not reduce the disparity between the data and model. This is perhaps not surprising given that foreign sales are a small fraction of total sales for the U.S.

We also note that small exporters arise naturally once we introduce latent heterogeneity. A small firm may export despite its low productivity if it draws a small fixed cost. Such a

²²The export intensity, that is, the ratio of foreign sales to total sales per firm, is on average 14% among exporters. Under strict sorting, the smallest exporter's total sales is just below \$5 million.

| | Data | Model |
|--------------------------|----------------|----------------|
| Ave. Domestic Sales - NX | \$8.1 million | \$740,000 |
| Ave. Domestic Sales - X | \$31.3 million | \$60.3 million |
| Exporter Size Premium | 3.8 | 81.5 |

Table 9: **Exporter size premium in terms of domestic sales.**

firm will export little, since it will not be competitive abroad. We use the fitted distribution in Section 3 as well as a constant export intensity of 14% in order to generate a distribution of export sales. We find that, under latent heterogeneity, over a quarter of exporters sell less than \$50,000 abroad. Recall that under the strict sorting model the *smallest* exporter would have about \$1 million in foreign sales. Thus latent heterogeneity goes a long way in explaining the presence of small exporters. That said, many exporters sell a very small amount in the data—e.g., less than \$20,000. For the latent heterogeneity to match this, we would need even more variation in the latent variable. However, such a small amount suggests that a fraction of exporters are firms that are not actively selling in a foreign market but rather they just fulfill the occasional order coming from abroad.

6.2 Industry-level heterogeneity

Here we explore whether the latent heterogeneity stems from between-industry variation. Manufacturing includes goods as diverse as tobacco products and machinery. So it is quite reasonable to think that sectors face very different trade costs, both fixed and variable. If the required heterogeneity is present at a very aggregate level, then we may be able to capture it with a simple two-sector specification, with a tradable and a non-tradable sector.²³

More generally, we compute the size premium implied by strict sorting for each three-digit NAICS industry code. The procedure is the same as the one we used in Section 2 for manufacturing as a whole. Bernard et al. (2007) report the share of firms exporting in each sector for 2002. As noted by Bernard et al. (2007), there is a large variation in the share of exporters. In Printing only 5% of the firms exports, while in Computers and Electronic Products almost 40% of the firms do. We also have the summary of the distribution of total sales for each sector—as in Table 1—provided by the Census.

Table 10 reports the size premium as predicted by strict sorting for each three-digit NAICS code. We compute a lower bound by assuming that all firms within any given size bin are identical; we also report a point estimate based on a fitted Pareto distribution. Both deliver the same message: *for each sector* the predicted size premium is very large. The reason is that the firm size distribution within a sector remains very skewed, so any strict sorting exercise is bound to return large size premia. The differences on the share of exporters does create a lot of variation in the implied size premiums across sectors. However,

²³This is often done in open macroeconomic models in order to replicate home bias or deviations from the law of one price.

| <i>Industry</i> | <i>NAICS code</i> | <i>Share Exporters</i> | Size Premium Prediction | |
|--------------------------|-------------------|------------------------|--------------------------------|-----------------|
| | | | <i>Lower Bound</i> | <i>Estimate</i> |
| Food | 311 | .12 | 115.3 | 118 |
| Beverage and tobacco | 312 | .23 | 193.7 | 246 |
| Textile mills | 313 | .23 | 49.6 | 69 |
| Textile product mills | 314 | .12 | 57.6 | 68 |
| Apparel | 315 | .08 | 55.8 | 56 |
| Leather | 316 | .24 | 38.8 | 43 |
| Wood product | 321 | .08 | 30.2 | 32 |
| Paper | 322 | .24 | 48.7 | 53 |
| Printing | 323 | .05 | 41.0 | 43 |
| Petroleum and coal | 324 | .18 | 164.4 | 165 |
| Chemical | 325 | .36 | 100.7 | 176 |
| Plastics | 326 | .28 | 30.9 | 31 |
| Nonmetallic mineral | 327 | .09 | 36.2 | 38 |
| Primary metal | 331 | .30 | 69.9 | 70 |
| Fabricated metal | 332 | .14 | 28.0 | 45 |
| Machinery | 333 | .33 | 33.0 | 43 |
| Computer and electronic | 334 | .38 | 72.4 | 97 |
| Electrical equipment | 335 | .38 | 48.3 | 67 |
| Transportation equipment | 336 | .28 | 190.5 | 298 |
| Furniture | 337 | .07 | 48.4 | 50 |
| Miscellaneous | 339 | .02 | 87.8 | 88 |
| Aggregate Manufacturing | 31-33 | .18 | 81.2 | 95 |

Table 10: **Lower bound for size premium predictions, by industry.**

it does not get them in the range observed in the data. As we saw in Table 3 the skewness in total sales is such that even if exporters were way more common than actually observed, the implied size premium would remain very large.

6.3 Sunk Costs

So far we have considered sources of variation in the export decision in a static framework. However, it is well known that the data display a fair amount of export hysteresis, which provides support for the presence of sunk costs.²⁴ While we are not concerned with the dynamics of exporting per se, export hysteresis has important implications for the cross-

²⁴Baldwin (1988) and Baldwin and Krugman (1989) are credited with developing the first theories of export hysteresis. Roberts and Tybout (1997) evaluate empirically the role of sunk costs through reduced-form models. More recent work has estimated structural models with sunk costs; see Das et al. (2007) and Ruhl and Willis (2008).

section of exporters. In the presence of sunk costs, the exporter status of a firm is not determined by its current productivity alone. An exporter may have no reason to stop exporting for productivity level φ , since it can still cover the per period fixed costs. The same productivity level, though, may not be high enough to convince a firm to start exporting as it will not cover the sunk costs. Indeed, there seems to be a basis for this possibility. Alessandria and Choi (2007b, 2008) calibrate models of sunk export costs that match the size and share of exporters, as well as the distribution of firm size in manufacturing. In these papers, though, there are many other sources of heterogeneity, so the quantitative role of export hysteresis in explaining the cross-sectional facts is not clear.

It is thus possible that the history of each firm provides the necessary heterogeneity to match the exporter size premium. A full analysis of this hypothesis is beyond the scope of this paper. We include here a simple exercise that suggests that sunk costs can explain the share and size of exporters but only by shifting the size “puzzle” from exporters to *new* exporters.

Consider the following variation of the reduced-form model in Section 3. Firm i 's total sales at date d , denoted r_{id} , are an *iid* random variable, with distribution Ψ_r . Since we do not evaluate the model at any frequency, the lack of any persistence is not particularly worrisome and allows for an easy characterization. There are two thresholds, t_0 and t_1 with $t_0 \leq t_1$, that determine the transition in and out of the foreign market as follows:

- An exporter at date $d - 1$ exports at date d if $r_{id} \geq t_0$.
- A non-exporter at date $d - 1$ exports at date d if $r_{id} \geq t_1$.

All readers familiar with sunk cost models will recognize the thresholds t_0 and t_1 as the stopper and starter points.

We now briefly show how we map the observations on the share and size of exporters, given the distribution of total sales, to pin down the stopper and starter points. The first equation to use is the steady-state condition on the share of exporters. Clearly, a fraction $\Psi_r(t_0)$ of previous exporters exit the exporter market, while a fraction $1 - \Psi_r(t_1)$ of the previously non-exporters start exporting now. If the share of exporters is constant, we must have that

$$s_x = (1 - \Psi_r(t_0)) s_x + (1 - \Psi_r(t_1)) (1 - s_x). \quad (17)$$

With some manipulation we have that equation (17) gives the share of exporters as a function of the starter and stopper points,

$$s_x = \frac{1 - \Psi_r(t_1)}{1 - \Psi_r(t_1) + \Psi_r(t_0)}. \quad (18)$$

Note the measure of exporters between t_0 and t_1 is simply $(\Psi_r(t_1) - \Psi_r(t_0))s_x$, so their proportion among exporters themselves is just $\Psi_r(t_1) - \Psi_r(t_0)$. We can then compute the average total sales by exporters as

$$\mu_x = (\Psi_r(t_1) - \Psi_r(t_0)) E\{r|r \in [t_0, t_1]\} + (1 - \Psi_r(t_1) + \Psi_r(t_0)) E\{r|r \geq t_1\}. \quad (19)$$

We then solve for the two equations (17) and (19), with t_0 and t_1 as the two unknowns.

We have no problem finding values for t_0 and t_1 that match the share and size of exporters. As simple as the set-up is, we can actually compute the entry and exit rates. We find that both are quite small (2.5% and 11.5%, respectively) indicating that there is a lot of persistence in the export status of the firm. Unfortunately, it is difficult to compare the numbers with the data, since we did not specify at what frequency the model is operating. Given our assumption that total sales are *iid*, we should take a period to be at very least five years. At that frequency the exit rate is very low compared with the data, as Bernard and Jensen (1999) report an annual stopper rate of 17%.²⁵

While the sunk cost model breaks the strict sorting between exporters and non-exporters, it now features strict sorting between *new* exporters and non-exporters. Because the entry rate is so small, the threshold for entry is way deep in the tail of the distribution (the 98th percentile to be precise). The size premium between new exporters and non-exporters is thus even larger, in the neighborhood of 200. All empirical evidence points to new exporters being significantly smaller than existing exporters—so it looks like to reconcile the sunk cost model with the size of new exporters we will need, again, a lot of independent variation in the entry decision.

6.4 Relaxing independence

We return to our fitting exercise from Section 3 and drop the assumption that the latent variable t_i is independent of firm size. Some firm characteristics—e.g., managerial quality, geographical location—may simultaneously affect total sales and the costs associated with exporting. We do not pursue any particular hypothesis. Instead, we will explore what kind of correlation between total sales and the latent variable would reduce the size-independent variation needed to match the data.

We now allow for total sales and the latent variable $\{r_i, t_i\}$ to be *jointly* distributed according to a log-normal distribution with mean μ and variance-covariance matrix Σ . We parameterize the joint distribution in terms of the latent variable:

$$\log(t_i) = \gamma \log(r_i) + \log(\epsilon_i)$$

where ϵ_i is independent of total sales r_i and distributed according to a log-normal distribution with mean μ_ϵ and standard deviation σ_ϵ . The parameter γ governs the correlation between total sales r_i and the latent variable t_i .

It is easy to rewrite the model in terms of total sales r_i and independent variation ϵ_i . Firm i exports if

$$r_i \geq t_i$$

or, in logs and in terms of ϵ_i ,

$$(1 - \gamma) \log(r_i) \geq \log(\epsilon_i).$$

²⁵The rate is for U.S. establishments in the period 1984-1992.

Given a value for γ we can proceed as we did in Section 3 to find values for μ_ϵ and σ_ϵ such that the model matches the share and size of exporters observed in the data.²⁶

We find that total sales and the latent variable have to be positively correlated in order to reduce the amount of variation independent of firm size, that is, $\gamma > 0$. The reason is quite straightforward. Strict sorting overestimates the size of exporters. Having larger firms clear, on average, a higher hurdle to start exporting reduces the number of large firms that export. There is then less need for the size-independent variable ϵ_i to take large values with high probability.

A positive correlation between firm-size and the latent variable questions the hypothesis that firms face *fixed* costs associated with exporting. A more natural explanation is that domestic and export production both require some input that is on a fixed supply at the firm level. For example, each firm may be endowed with one manager with a limited span of control, as in Lucas (1978). For firms with higher labor productivity, the opportunity cost of the manager's control is higher.

It should be made clear that the positive correlation does not rule out non-convexities but rather suggests that they require a more nuanced approach. For example, Cooper and Haltiwanger (1993) assume that firms face some down time when replacing "machines" or, more generally, when pursuing some investment such as upgrading the production process. The non-convexity arises because it is not possible to replace, say, a fraction of a machine with a fraction of the down time. The down time is more expensive for more productive firms, so while the investment technology is identical across firms, the *cost* of investment is higher for larger firms.

Another possibility is that large firms face different market conditions than small firms. For example, the fair wage literature links the wages paid with the firm's profitability.²⁷ If the overhead costs of exporting are in terms of in-firm labor, then more productive firms would face a higher cost of exporting.

²⁶Let $\log(h_i) = (1-\gamma)\log(r_i)$. Given γ and the distribution of total sales Ψ_r , we can derive the distribution Ψ_ϵ . The share of exporters is the share of firms with $\log(h_i) \geq \log(\epsilon_i)$, so

$$s_x = \int \Psi_\epsilon(h) d\Psi_h(h).$$

Similarly, the size premium is related to the log difference in h_i between exporters and non-exporters:

$$(1-\gamma)E\{\log(r_i) | r_i \geq t_i\} = \int \log(h) \Psi_\epsilon(h) d\Psi_h(h) / s_x.$$

²⁷The literature dates back to Akerloff (1982). It has been recently introduced in international trade. See Amiti and Davis (2008), for example.

6.5 Heterogeneity in trade costs and foreign demand

So far we have explored some hypothesis for the large amount of size-independent variation needed to match the data on the size and share of exporters. We have not, though, explored the implications of these hypotheses in the event of a trade liberalization. Perhaps our results in Section 5 are tied to the assumption that fixed costs explain all of the variation in the data. An obvious alternative is that firms have heterogeneous trade costs or there is firm-market demand heterogeneity, since perhaps foreign consumers do not value quality as domestic consumers do.²⁸ Since foreign sales are a small fraction of total sales for U.S. firms, heterogeneity in trade costs or foreign demand would be essentially independent of size.

We find that the model implications are virtually unchanged if the necessary size-independent variation is modeled as heterogeneous fixed costs, trade costs, or foreign demand.²⁹ It is easy to see why this is the case. First, if the additional heterogeneity enters the exporter revenues multiplicatively, then the resulting model is actually isomorphic to heterogeneity in fixed costs. The reason is that the model is essentially loglinear. Consider this simple specification: on top of the trade costs τ , the marginal cost of firm's ω exports is $(1 + \eta(\omega))$ times more than the marginal cost for output sold at the domestic market. That is, firms differ on the *variable* costs associated with exporting.

We can rewrite (10) in this model as

$$r^f(\omega) = \left(\frac{\varphi(\omega)}{\tilde{\varphi}_x} \right)^{\theta-1} (1 + \eta(\omega))^{-\theta} \frac{R^f}{M_x}.$$

The entry condition (11) would now be

$$\left(\frac{\varphi(\omega)}{\tilde{\varphi}_x} \right)^{\theta-1} (1 + \eta(\omega))^{-\theta} \frac{R^f}{M_x} \geq \theta f(\omega).$$

We could, though, simply define $\tilde{f}(\omega) = f(\omega)(1 + \eta(\omega))^\theta$ and the equilibrium condition is identical.³⁰ Quantitatively we would calibrate the distribution of both η and f to reproduce the variation in the latent variable—in other words, we would calibrate the variable $\tilde{f}(\omega)$ as we did with $f(\omega)$ in Section 4. So as long η and f remain independent of φ , the results would be unchanged.³¹

We do not expect big differences even if the conditions for the isomorphism do not apply. Say we consider firm variation in the trade costs *per unit*, that is, the marginal cost of a firm is now $\eta(\omega) + \tau\omega/\varphi(\omega)$. In this case the entry condition (11) is not loglinear. We can,

²⁸Munch and Nguyen (2008) find that firm-specific factors such as productivity explain a very small fraction of the sales variation across destinations for Danish exporters.

²⁹As long, of course, as the full amount of latent heterogeneity is captured by the specification of choice.

³⁰This requires, of course, that the additional heterogeneity is normalized such that it does not change the baseline calibration or the policy experiment.

³¹The source of heterogeneity will make a difference in a general equilibrium model or in the computation of the transition dynamics. We are skeptical, though, that the differences will be economically significant.

though, think of a n th order log-approximation to (10) around the representative firm type, $\tilde{\omega}$. The higher order terms may be important for large reductions of trade costs: they are, though, only on the measure that the variation η interacts with the productivity parameter, φ . This interaction is not captured by the latent variable in Section 3—so non-multiplicative factors affect results only on the measure that they induce departures from the assumption of linear independence.

7 The Extensive Margin in the Data

The empirical literature has not come to a consensus on the quantitative importance of the extensive margin for aggregate trade patterns. This reflects, in part, that there is no unique concept of the extensive margin: one can define entry and exit at the level of the firm, plant, or product.

Two papers are well known for arguing that the extensive margin is an important dimension of aggregate data. Hummels and Klenow (2005) find that the extensive margin accounts for 60% of the cross-country differences in trade. Broda and Weinstein (2006) find large welfare gains associated with the expansion in import variety for the U.S over the last three decades. Both papers build upon the analysis in Feenstra (1994) and share a focus on the long run.

In our exercise we have used the measure of the extensive margin in Alessandria and Choi (2008) so we take some time to discuss their results. Using the census of manufacturers, Alessandria and Choi (2008) look at the increase in export participation by plants and find that it accounts for half of the export growth in the U.S. from 1987 to 2002. Interestingly, this is quite close to the prediction of the standard Melitz model for the extensive margin. We cannot ignore, though, that the standard version of the model allocates an important role to the extensive margin only by greatly exaggerating the economies of scale.

Recently, several papers have taken the position that the extensive margin contributes little to aggregate trade patterns. Besedes and Prusa (2008) argue that new exporting relationships have little impact on long-run export growth because they tend to be very short-lived. Arkolakis et al. (2008) document a sizeable increase in variety in Costa Rica from 1986 to 1992. They argue, though, that the increase did not translate into large welfare gains because new varieties were imported in small quantities. Armenter and Koren (2008) show that several well-known facts about the extensive margin are the result of the sparsity of the data.

Summarizing, there is no definitive evidence of the quantitative importance of the extensive margin in the determination of aggregate trade patterns. While we may record much activity along the extensive margin—say, many new products are exported in a given year—this does not necessarily translate into aggregate patterns—either because these new products are dropped the next year or they constitute a very small fraction of total exports.

8 Conclusions

Since Melitz (2003), models with economies of scale and firm-level heterogeneity have become very popular in international trade. These models have been used to explain the characteristics of multi-product exporters and exporters-importers, among others. Most researchers are content to show that the qualitative predictions are as documented in the data.

In this paper we have shown that a simple Melitz model cannot match *quantitatively* the basic cross-sectional facts about exporters, namely, their frequency and size. The model can be easily reconciled with the data by introducing enough additional sources of variation. However, this is not without implications. First, we need a large amount of independent variation, suggesting that firm size is not the main determinant of the export status of a firm. Second, the augmented model has a minor role for the extensive margin in the event of a trade liberalization, and the productivity gains due to reallocation of resources are smaller than in the standard model.

Given the attention the literature has given to the extensive margin, it is quite surprising that the calibrated version of the Melitz model features only small changes in export participation which, in addition, contribute little to overall trade growth. This “puzzle” appears quite robust: the source of variation at work does not seem important. In order to generate relatively small exporters, one needs to assume that larger firms face significantly larger costs of exporting, which seems at odds with the concept of fixed costs.

We find the ideas in Melitz (2003) compelling enough to ask what can reconcile the data with a model of economies of scale in exporting. In our opinion such a model would need to move beyond the assumption of fixed costs and take a more nuanced look at how the economies of scale arise. We should also ask whether we are measuring export participation correctly. One particular hypothesis of interest is that the Melitz model fits well the behavior and characteristics of large exporters, but the facts on the extensive margin are driven by a set of small exporters. Since most exports are made by large exporters, Melitz (2003) would be the right framework from a macroeconomic perspective.

References

- [1] Aitchison, J. and Brown, J.A.C.: 1969 The Lognormal Distribution. The University Press, Cambridge.
- [2] Akerloff, G.: 1982 Labor Contracts as a Partial Gift Exchange. *Quarterly Journal of Economics* **97**(4), 543-569.
- [3] Alessandria, G. and Choi, H.: 2007a, Do Sunk Costs of Exporting Matter for Net Export Dynamics? *Quarterly Journal of Economics* **122**(1), 289-336.

- [4] Alessandria, G. and Choi, H.: 2007b, Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization. Working Paper 07-17, Federal Reserve Bank of Philadelphia.
- [5] Alessandria, G. and Choi, H.: 2008, Do Falling Iceberg Costs Account for Recent US Export Growth? Manuscript, Federal Reserve Bank of Philadelphia.
- [6] Amiti, M. and Davis, D.: 2008, Trade, Firms, and Wages: Theory and Evidence. NBER Working Paper 14106.
- [7] Anderson, J.E. and van Wincoop, E.: 2004, Trade Costs. *Journal of Economic Literature*, 42(2), 691-751.
- [8] Arkolakis, C. : 2008 Market Penetration Costs and the New Consumers Margin in International Trade. Working paper, Yale University.
- [9] Arkolakis, C., Demidova, S., Klenow, P.J. and Rodriguez-Clare, Andres: 2008, Endogenous Variety and the Gains from Trade. *American Economic Review* **98**(4), 444-450.
- [10] Armenter, R. and Koren, M.: 2008, A Balls-and-Bins Model of Trade. Working Paper. Federal Reserve Bank of New York.
- [11] Atkeson, A. and Burstein, A.: 2008, Innovation, Firm Dynamics, and International Trade. NBER Working paper 13326.
- [12] Axtell, R. L.: 2001, Zipf Distribution of U.S. Firm Sizes. *Science* **293**(5536), 1818–1820.
- [13] Baldwin, R.: 1998, Hysteresis in Import Prices: The Beachhead Effect. *American Economic Review* **78**(4), 773–85.
- [14] Baldwin, R., and Krugman, P.: 1989 Persistent Trade Effects of Large Exchange Rate Shocks. *Quarterly Journal of Economics* **104**(4), 635-54.
- [15] Baldwin, R. and Harrigan, J.: 2007, Zeros, Quality and Space: Trade Theory and Trade Evidence. NBER Working Paper No. 13214.
- [16] Bernard, A. B., Eaton, J., Jensen, J. B. and Kortum, S.: 2003, Plants and Productivity in International Trade. *American Economic Review* **93**(4), 1268–1290.
- [17] Bernard, A. B. and Jensen, J. B: 1999, Exceptional Exporter Performance: Cause, Effect, or Both? *Journal of International Economics* **47**(1), 1–25.
- [18] Bernard, A. B., Jensen, J. B., Redding, S. J. and Schott, P. K.: 2007, Firms in International Trade. *Journal of Economic Perspectives* **21**(3), 105–130.

- [19] Bernard, A. B., Jensen, J. B. and Schott, P. K.: 2007, Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods, *in* Dunne, J., J.B. Jensen and M.J. Roberts (eds.), *Producer Dynamics: New Evidence from Micro Data*.
- [20] Besedes, T. and Prusa, T.: 2008 The Role of Extensive and Intensive Margin and Export Growth. Working Paper, Rutgers University.
- [21] Breinlich, H. and Cuñat, A.: 2009 Trade Liberalization and Heterogenous Firm Models: An Evaluation Using the Canada-U.S. Free Trade Agreement. Mimeograph, LSE and University of Essex.
- [22] Broda, C. and Weinstein, D.: 2006 Globalization and the Gains from Variety. *Quarterly Journal of Economics* **121** (2).
- [23] Chaney, T.: 2008, Distorted Gravity: the Intensive and Extensive Margins of International Trade. *American Economic Review* **98**(4), 1707-1721.
- [24] Cooper, R., and Haltiwanger, J. : 1993, The Aggregate Implications of Machine Replacement: Theory and Evidence. *American Economic Review* **83**(3), 360-382.
- [25] Das, S., Roberts, M. and Tybout, J.: 2007 Market Entry Costs, Producer Heterogeneity, and Export Dynamics. *Econometrica* **75**(3), 837-873.
- [26] Eaton, J., Eslava, M., Kugler, M. and Tybout, J.: 2007, Export Dynamics in Colombia: Firm-Level Evidence. NBER Working Paper No. 13531.
- [27] Eaton, J., Kortum, S. and Kramarz, F.: 2004, Dissecting Trade: Firms, Industries, and Export Destinations. *American Economic Review* **94**(2), 150–154.
- [28] Eaton, J., Kortum, S. and Kramarz, F.: 2008, An Anatomy of International Trade: Evidence from French Firms. Working Paper.
- [29] Feenstra, R.:1994 New Product Varieties and the Measurement of International Prices. *American Economic Review* **84**(1), 157-177.
- [30] Ghironi, F. and Melitz, M.: 2005 International Trade and Macroeconomic Dynamics with Heterogeneous Firms. *Quarterly Journal of Economics* **120**, 865-915.
- [31] Ghironi, F. and Melitz, M.: 2007 Trade Flow Dynamics with Heterogeneous Firms. *American Economic Review* **97**, 356-361.
- [32] Hallak, J.C. and Sivadasan, J.: 2008 Productivity, Quality and Exporting Behavior Under Minimum Quality Requirements. Working paper.
- [33] Halpern, L., Koren, M. and Szeidl, A.: 2007 Imports and Productivity. Working Paper.

- [34] Helpman, E., Melitz, M. and Rubinstein, Y.: 2007, Estimating Trade Flows: Trading Partners and Trading Volumes. *Quarterly Journal of Economics*, forthcoming.
- [35] Hopenhayn, H: 1992, Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, **60**, 1127-1150.
- [36] Hummels, D., Klenow, P. J.: 2005, The Variety and Quality of a Nation's Exports. *American Economic Review* **95**(3), 704–723.
- [37] Johnson, Rob: 2007 Trade and Prices with Heterogeneous Firms. Working Paper.
- [38] Johnson, N. L., Keppm, A. W., and Kotz, S.: 2005, *Univariate Discrete Distributions*, John Wiley & Sons.
- [39] Kugler, M. and Verhoogen, E.: 2008 The Quality-Complementarity Hypothesis: Theory and Evidence from Colombia, *NBER Working Paper* 14418.
- [40] Lucas, R.: 1978 On the size distribution of business firms. *Bell Journal of Economics* **9**, 508-523.
- [41] Luttmer, E.: 2007 Selection, Growth, and the Size Distribution of Firms. *Quarterly Journal of Economics*, **122** (3), 1103-1144.
- [42] Mayer, T. and Ottaviano, G.: 2007, The Happy Few: The Internationalization of European Firms. Bruegel Blueprint Series. Volume III.
- [43] Melitz, M. J.: 2003, The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity. *Econometrica* **71**(6), 1695–1725.
- [44] Munch, Jakob and Nguyen, D.: 2008, Decomposing Firm-level Sales Variation. Working paper.
- [45] Roberts, M. and Tybout, J.: 1997 The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs. *American Economic Review* **87**(4), 545-564.
- [46] Ruhl, K.: 2003 The International Elasticity Puzzle. UT Austin Working Paper.
- [47] Ruhl, K. and Willis, J.: 2008 New Export Dynamics. NYU Working Paper.

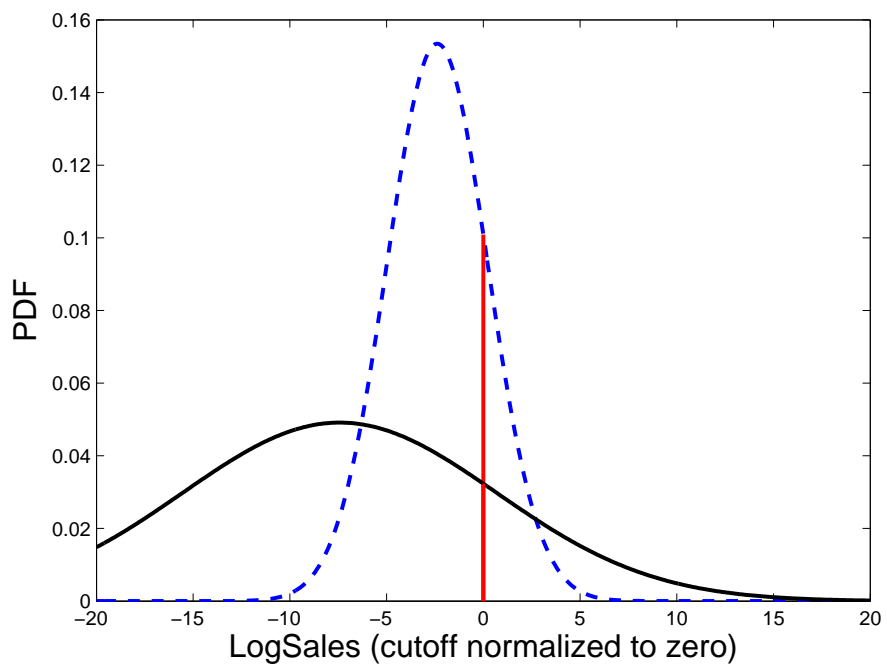


Figure 1: The extensive margin

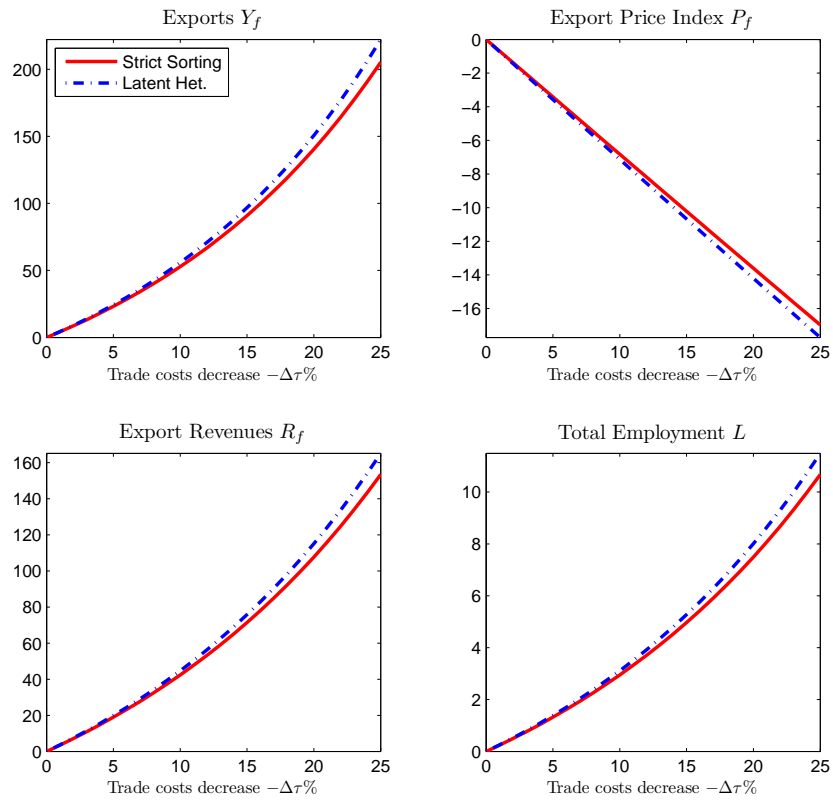


Figure 2: Aggregate Exports *Trade costs decrease in percentage points. All variables in growth rates*

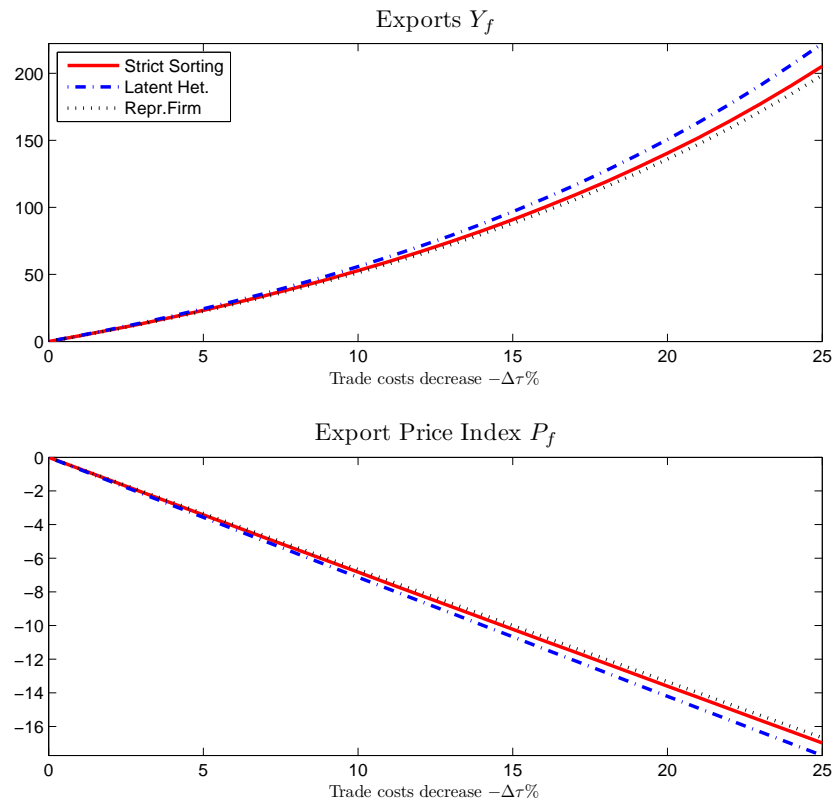


Figure 3: Aggregate Exports - Comparison with the Representative Firm Model *Trade costs decrease in percentage points. All variables in growth rates*

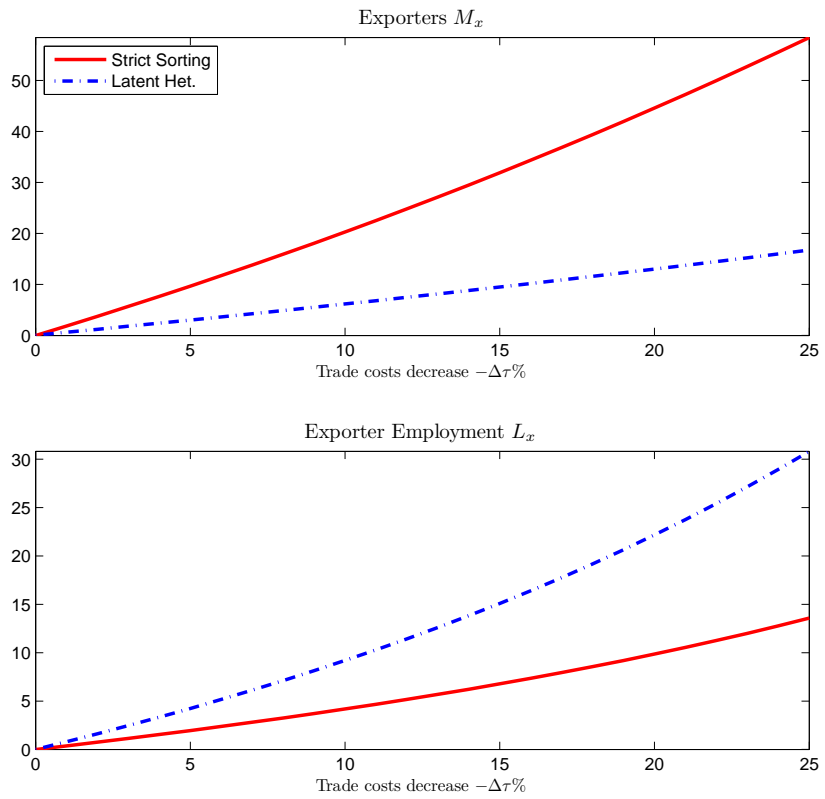


Figure 4: Entry and Employment Growth by Exporters *Trade costs decrease in percentage points. All variables in growth rates*




Figure 5: Aggregate Productivity *Trade costs decrease in percentage points. All variables in growth rates*




Figure 6: Exporters *Trade costs decrease in percentage points. All variables in growth rates*