

Decomposing Firm-level Sales Variation¹

Jakob R. Munch and Daniel X. Nguyen²

July 8, 2008

¹We would like to thank David Hummels and participants at Purdue University and the University of Aarhus for their helpful suggestions and comments. All errors are our own.

²Jakob.Roland.Munch@econ.ku.dk, Daniel.X.Nguyen@econ.ku.dk

Abstract

Much of the ongoing international trade literature uses productivity heterogeneity to explain the variation of firm-level sales in a market. We measure the proportion of that variation that can be accounted for by productivity or other firm-specific characteristics. We introduce several empirical tests with structural bases in a monopolistic competitive framework and measure the contribution of firm-specific factors to overall market sales variation for Danish exporters between 1999 and 2002. We find that firm-specific heterogeneity accounts for less than 15% of both the variation across destinations and across time. We show that our measures of this contribution are upwardly biased if productivity is positively correlated to firm-destination-specific sales shocks, or if firm-destination-specific shocks are positively correlated across destinations. We conclude that while firm-specific productivity can account for some of the variation, the majority must be explained by firm-destination-specific heterogeneity sources such as market-specific demand.

1 Introduction

Between the years 1999 and 2002, over 3600 Danish manufacturing firms shipped over 2400 products to more than 170 destination countries. The annual values of these shipments ranged from less than a dollar to more than \$300 million. Recent works in international trade have attributed this revenue variation to differences in firm productivities. This paper estimates the degree to which productivity can explain firm revenue variation within a market. When corrected for truncation bias, estimates of that degree range from 0 to 30%, suggesting that productivity among firms is not the primary driver of revenue variation.

The current set of international trade literature is dominated by models using idiosyncratic productivities to explain differences among firms that export (Eaton and Kortum 2002, Bernard, Eaton, Jensen, and Kortum 2003, Melitz 2003, Bernard, Jensen, and Schott 2006; Bernard, Jensen, Redding, and Schott 2007). These models were motivated by earlier empirical studies identifying differences between firms that export and firms that do not (Aw, Chung, and Roberts 2000; Clerides, Lach, and Tybout 1998; Bernard and Jensen 1995, 1999). The studies found that, on average, exporters produce more, hire more labor, pay the labor higher wages, and exhibit higher productivities as measured by either Total Factor Productivity (TFP) or Value Added per worker. These contrasts supported the story that firms with high productivities self-selected themselves into exporting. However, the differences between firms that do and do not export are not absolute. Nguyen (2008) shows that the productivities of the top half of non-exporters are higher than those of the bottom half of exporters. He also shows that the correlation between a firm's foreign and domestic sales rank is weak and hard to reconcile with firm-specific productivity models of exporting. This paper provides further evidence against the claim that a firm's productivity determines how much it exports.

This paper also adds to a small but growing literature examining the destinations to which firms export. This is due primarily to the dearth of firm-level destination data. Eaton, Kortum, and Kramarz (2004) find that most French firms export to only one destination (the mode being Belgium), and that the entry of French firms into a market accounts for two-thirds of the growth of the French share of market sales. Newer studies show that firms usually wait to export, that they usually begin exporting to one destination country, and that many stop exporting activities soon after they begin (Eaton, Eslava, Kugler, and Tybout 2007; Damijan, Kostevc, and Polanec 2007; Alvarez, Faruq, and Lopez 2007). These studies, and the current one, present empirical patterns that are not readily explained by the productivity models of exporting such as Melitz (2003).

In Melitz (2003), each firm has a unique productivity that maps directly to its revenues in all destinations. Therefore, a firm's revenue relative to its competitors is identical across destinations. In this paper, we test that prediction at highly disaggregated product

levels. For example, we find that the prediction holds for exports of plastic boxes¹ to Sweden and Germany², as the left panel of Figure 1 depicts. The scatter shows an approximate one-to-one correlation between a Danish firm's relative revenue in Sweden and its relative revenue in Germany. A Danish plastic box manufacturer's revenue in Germany can precisely forecast its revenue in Sweden. However, the same prediction does not hold for exports of cotton shirts³. As the right panel of Figure 1 depicts, Danish shirt manufacturers with high relative revenues in Germany do not necessarily have high relative revenues in Sweden, or vice versa.

Insert Figure 1 here

The plastic boxes and cotton T-shirt examples illustrate the basics of the empirical strategy of this paper: we compare relative revenues across destinations to estimate the contribution of firm-specific productivity. We ask whether the bulk of Danish exports follow a pattern akin to the left panel of Figure 1 or the right one. Using a structural model based on Melitz (2003) and Nguyen (2008), we decompose a firm's revenue relative to the destination average into a firm-specific (FS) component and a firm-destination-specific (FDS) component. We then construct two statistics that estimate the contribution of the FS component. The first statistic is the proportion of a firm's relative sales captured by its productivity. The second is the proportion of markets sales variance that can be explained by productivity.

A related study by Eaton, Kortum, and Kramarz (2007) estimates the contribution of firm-specific productivity to sales variance by calibrating a firm heterogeneity model to the destination-specific total exports of French firms. Their model incorporates firm-specific productivity shocks drawn from a Pareto distribution and firm-destination-specific taste and cost shocks drawn from lognormal distributions. They calibrate the parameters of their model to match approximately 1000 distributional moments⁴. Using this computationally intensive calibration, Eaton Kortum and Kramarz (2007) find that the variance of the firm-specific shock can account for approximately half of the

¹"Boxes, cases, crates and similar articles for the conveyance or packaging of goods, of plastics" (product code 39231000)

²which are the two most popular destinations for Danish exporters.

³"Men's or boy's shirts of cotton (excl. knitted or crocheted, nightshirts, singlets and other vests)" (product code 62052000)

⁴For each of the 113 destination countries, Eaton, Kortum, and Kramarz (2007) calibrate their model to match a) the number of French firms selling to that country, b) the number of French firms selling below the fiftieth, seventy-fifth, and ninety-fifth percentile in that country, c) the number of French firms selling in that country and selling below the fiftieth, seventy-fifth, and ninety-fifth percentile in France, d) the number of French firms selling below the fiftieth, seventy-fifth, and ninety-fifth percentile in the France, e) the number of French firms exporting to that country whose ratio of sales in that country to sales in France is below the seventy-fifth of firms selling to that country, and f) the number of simulated firms selling to each possible combination of the seven most popular export destinations.

sales variation in a market. Employing Ockham’s Razor, this paper’s empirical strategy is more simple. Since we have Danish firms’ exports to multiple destinations, we can separate a firm’s relative revenue in a market into a firm-specific component and a firm-destination-specific component. Since productivity is anchored to the plant location, the firm-specific component can be used to estimate the upper bound of the contribution of productivity to the relative revenue. Because we have product-level data not present in Eaton, Kortum, and Kramarz (2007), we can perform this exercise at economy, sector, and product levels, allowing us to compare the contribution of productivity for both broadly defined and narrowly defined markets.

We describe the data in the following section. In section 3, we outline the structural model motivating our empirical tests and construct our statistics of interest. We then consider possible truncation and econometric issues that may bias our estimates. Results follow. Finally, as our estimations amount to separating firm-specific effects from firm-destination-specific effects, we consider alternative interpretations of our results. We calculate the variances of measurement error or firm-destination-specific trade costs needed to reproduce our empirical results. The necessary variances are implausibly large. We conclude by advocating a firm-destination-specific demand shock story such as that in Nguyen (2008) to better explain market sales variation.

2 Danish Firm-level Data

The Danish External Trade Statistics provides product-level destination-specific export data for all Danish firms for the years 1999-2002. Exports are recorded according to the eight-digit Combined Nomenclature product code which encompasses approximately 10,000 different product categories. In addition to firm level exports by product and destination, the dataset contains information on firm size defined as the number of full-time employees, capital stock defined as the value of land, buildings, machines, equipment and inventory, and value added defined as revenues net of input costs. In the following analysis we restrict the sample to exports by manufacturing firms employing at least one employee. With this restriction, our dataset comprises 604934 firm-product-destination-year shipments of 6731 products to 242 countries by 6102 firms totaling 698 billion Danish Kroner (DKK). Table 1 lists some summary statistics for the number of products, firms and destinations in 2002 by Standard International Trade Classification (SITC1=5,6,7,8) industry.

Insert table 1 here

This fifth column of table 1 shows that the average firms per destination per product is less than 1.5: many products are only one firm and shipped to a single destination. Our regressions require firm-product-year observations that appear in at least two destinations and destination-product-year observations that contain at least two firms. These

requirements eliminate a third of the observations. We also exclude small trade flows by deleting the 5% smallest flows within each product-year to minimize measurement error or outlier effects. This last restriction does not affect our results significantly. These restrictions reduce our sample set to 305981 shipments of 2468 products to 179 countries by 3667 firms totaling 454 billion DKK, corresponding to 65% of the value of original sample.

While all trade flows with non-EU countries are recorded by customs authorities (and so the coverage rate in the data is close to complete), there is not a similar system in place for intra-EU trade. However, intra-EU trade is recorded through the Intrastat system, where firms are obliged to report trade data on a monthly basis. One source of inaccuracy in this system is that some firms appear not to report data to the system. Also, data on intra-EU trade is censored in a way such that only firms exporting goods with a total annual value exceeding a certain threshold⁵ are recorded in the files. No such data limitations exist for trade out of the EU. As a result the coverage rate in the Intrastat system is only in the range 85-90 percent. See Statistics Denmark (2003) for further details.

3 Theory

Our model is based on Melitz (2003) and Nguyen (2008). We employ two types of heterogeneity (firm-specific and firm-destination-specific) and consider multiple products and destinations. Ours is a partial equilibrium model identifying the sources of revenue heterogeneity; it does not attempt any general equilibrium closures such as market clearing or zero expected profits. Since Melitz (2003) and Nguyen (2008) are general equilibrium models, we invite the reader to review those papers for questions unrelated to this paper's empirical objective. Table 2 in Appendix lists the notation for ease of reference.

3.1 A Model of Revenue Variation

The small open economy of Denmark supplies products indexed by $n \in N$ to foreign destinations $j \in J$ in each time period t . In each product category n , there are $|n|$ firms indexed by $\omega \in n$ each producing a unique variety of n ⁶. We use n to denote both the product and the set of firms producing n . The utility u gained by consumers in each destination j by consuming quantities q of the varieties ω of the products n available to j at time t is

⁵For 2002, this threshold was 2.5 million Danish Kroner corresponding to 335,000 Euro .

⁶While there can exist a multi-product firm that produces a set of products $N_\omega \subseteq N$, we treat that firm as $|N_\omega|$ distinct firms, as that firm's $|N_\omega|$ products each competes in its own separate market.

$$u_{jt} = \prod_{n=1}^N \left(\int_{\omega \in n_{jt}} \exp\left(\frac{x_{\omega jt}}{\sigma}\right) (q_{\omega jt})^{\frac{\sigma-1}{\sigma}} \right)^{\alpha_{nt}} \quad (1)$$

where $\sigma > 1$ is a measure of the substitutability among varieties of a product and α_n is the Cobb-Douglas fraction of income spent on product n such that $\sum_{n=1}^N \alpha_{nt} = 1$. The Cobb-Douglas utility form results in price competition among firms of n in j at t , but not between firms of disparate products, destinations, or times. Therefore, a market is completely described by the triplet n_{jt} . $x_{\omega jt} \in (-\infty, \infty)$ is a market n_{jt} specific demand shifter for ω ⁷. Destination j 's demand for firm ω 's variety at time t is represented by:

$$q_{\omega jt} = \frac{p_{\omega jt}^{-\sigma}}{P_{n_{jt}}} (\alpha_{nt} Y_{jt}) \exp(x_{\omega jt}) \quad (2)$$

$$P_{n_{jt}} = \int_{\omega \in n_{jt}} \exp(x_{\omega jt}) p_{\omega jt}^{1-\sigma} \quad (3)$$

where $p_{\omega jt}$ is the price of variety ω and $\alpha_{nt} Y_{jt}$ is the total expenditure on all varieties in market n_{jt} . $P_{n_{jt}}$ is the corresponding price index and is unaffected by the actions of any single firm supplying to n_{jt} .

Firms produce their unique varieties using the single input labor. The labor costs l incurred when producing q units of output for destination j are

$$l_{\omega jt} = f_{n_{jt}} + \exp\left(\frac{\varphi_{\omega t}}{1-\sigma}\right) \tau_{n_{jt}} q_{\omega jt} \quad (4)$$

where f and τ are the fixed and marginal trade costs, respectively, of supplying to market n_{jt} . $\exp\left(\frac{\varphi_{\omega t}}{1-\sigma}\right)$ is the marginal production cost of firm ω that is constant across destinations. $\varphi_{\omega t}$ is a normalized measure of the firm's productivity: a higher φ translates to a lower marginal production cost. The profit-maximizing firm sets the price of ω as a constant mark-up ($= \frac{\sigma}{\sigma-1}$) over marginal cost and garners sales of $r_{\omega jt}$:

$$r_{\omega jt} = \frac{\alpha_{nt} Y_{jt} \tau_{n_{jt}}^{1-\sigma}}{P_{n_{jt}}} \exp(\varphi_{\omega t} + x_{\omega jt}) \quad (5)$$

Equation (5) characterizes the revenue for any firm selling in market n_{jt} . Each firm ω draws its productivity parameter $\varphi_{\omega t}$ from an exogenous product specific distribution with a mean of $\bar{\varphi}_{nt}$ and a variance of $s_{\varphi n}^2$:

$$E[\varphi_{\omega t}] = \bar{\varphi}_{nt} \quad (6)$$

$$E[(\varphi_{\omega t} - \bar{\varphi}_{nt})^2] = s_{\varphi n}^2 \quad (7)$$

⁷We consider other interpretations of $x_{j\omega}$ in section 7

For each market, the firm draws their demand shifter $x_{\omega jt}$ from an exogenous market specific distribution with mean \bar{x}_{njt} and standard deviation s_{xn}^2 :

$$E[x_{\omega jt}] = \bar{x}_{njt} \quad (8)$$

$$E[(x_{\omega jt} - \bar{x}_{njt})^2] = s_{xn}^2 \quad (9)$$

3.2 Decomposition

In this paper, we decompose sales variation in njt and identify the contribution of $\varphi_{\omega t}$ to it. To do so, we take logs of equation (5):

$$\ln r_{\omega jt} = \ln (\alpha_{nt} \tau_{jn}^{1-\sigma} Y_{jt} / P_{jnt}) + \varphi_{\omega t} + x_{\omega jt} \quad (10)$$

To eliminate market-specific effects, we define $E[\ln r_{\omega jt}]$ as the unconditional expectation of $\ln r_{\omega jt}$ for all Danish firms producing a variety of n , whether or not they supply to j :

$$E[\ln r_{\omega jt}] = \ln (\alpha_{nt} \tau_{jn}^{1-\sigma} Y_{jt} / P_{jnt}) + E[\varphi_{\omega t}] + E[x_{\omega jt}]$$

and define relative revenue $\tilde{r}_{\omega jt}$, relative productivity $\tilde{\varphi}_{\omega t}$, and relative demand shifter $\tilde{x}_{\omega jt}$:

$$\tilde{r}_{\omega jt} = \ln r_{\omega jt} - E[\ln r_{\omega jt}] \quad (11)$$

$$\tilde{\varphi}_{\omega t} = \varphi_{\omega t} - \bar{\varphi}_{nt} \quad (12)$$

$$\tilde{x}_{\omega jt} = x_{\omega jt} - \bar{x}_{njt} \quad (13)$$

so that equation (10) reduces to:

$$\tilde{r}_{\omega jt} = \tilde{\varphi}_{\omega t} + \tilde{x}_{\omega jt}. \quad (14)$$

$\tilde{r}_{\omega jt}$ is the market njt revenue of firm ω relative to the expected revenue of a Danish producer of n . $\tilde{\varphi}_{\omega t}$ is ω 's productivity relative to the average producer of n , which affects its relative revenue in all destinations. $\tilde{x}_{\omega jt}$ is ω 's demand shifter specific to j . Our goal is to estimate the relative importance of $\tilde{\varphi}_{\omega t}$ and $\tilde{x}_{\omega jt}$. To do so, we introduce two statistics. First, we define

$$\beta_n = \frac{\tilde{\varphi}_{\omega t}}{\tilde{\varphi}_{\omega t} + \tilde{x}_{\omega jt}} \quad (15)$$

as the proportion of variety ω 's relative revenue captured by $\tilde{\varphi}_{\omega t}$. If $\beta_n = 1$, relative productivity contributes to the entirety of relative sales of a firm producing n . If $\beta_n = 0$, firm-destination-specific demand shocks are entirely driving sales for that product. We estimate β_n for each product n .

Our second statistic considers the variances of the distributions of $\tilde{\varphi}$ and \tilde{x} . We define

$$Q_n^2 = \frac{\text{var}(\tilde{\varphi}_{\omega t})}{\text{var}(\tilde{r}_{\omega jt})} \quad (16)$$

as the proportion of the variance of $\tilde{r}_{\omega jt}$ that is accounted for by $\tilde{\varphi}_{\omega t}$. If $Q_n^2 = 1$, the firm's relative productivity completely characterizes its relative revenue. If $Q_n^2 = 0$, none of the variation in a firm's relative revenue can be explained by productivity. We estimate Q_n^2 for each product n .

4 Econometric considerations

To identify β_n and Q_n^2 , we estimate equation (14) with Danish firm-destination-product-year-specific export data from 1999-2002 using both TFP and firm-dummies to proxy for $\tilde{\varphi}_{\omega t}$. In this section, we discuss the implication of two major sources of bias in our regression. The first is truncation error due to self-selection in our unbalanced panel dataset. The second is error due to possible correlation between $\varphi_{\omega t}$ and $x_{\omega jt}$, or between $x_{\omega jt}$ and $x_{\omega it}$ for two destinations j and i .

4.1 Truncation

In equation (11), we construct $\tilde{r}_{\omega jt}$ using the unconditional expectation of $\ln r_{\omega jt}$, which can be estimated with

$$E[\ln r_{\omega jt}] = \frac{1}{|n|} \sum_{\omega \in n} \ln r_{\omega jt}. \quad (17)$$

The calculation of $E[\ln r_{\omega jt}]$ would be straightforward if all firms producing n supplied to njt . However, we have an unbalanced panel where not every firm ships to every destination every period. We showed in Table 1 that while, on average, 8.5 Danish firms export an 8-digit CN product n , there are only 1.5 firms per destination j per product n . So, on average, six of the firms producing n are not supplying to j . Since we cannot observe Danish firms that choose not to supply to j , we can only measure $E_{njt}[\ln r_{\omega jt}]$, the market njt observed mean conditional upon entry:

$$E_{njt}[\ln r_{\omega jt}] = \frac{1}{|n_{jt}|} \sum_{\omega \in n_{jt}} \ln r_{\omega jt}. \quad (18)$$

If the entry of firm ω to j is uncorrelated to its potential revenue $r_{\omega jt}$, then $E_{njt}[\ln r_{\omega jt}]$ serves as an unbiased measure of $E[\ln r_{\omega jt}]$ in the construction of $\tilde{r}_{\omega jt}$. However, if a firm's potential $r_{\omega jt}$ is correlated with its decision to enter a market, then $E_{njt}[\ln r_{\omega jt}]$, and subsequently $\tilde{r}_{\omega jt}$, is biased. Current literature suggests this latter scenario, so we consider its consequences:

Melitz (2003) models market entry as a function of expected profit. Profit $\pi_{\omega jt}$ for ω in njt is

$$\pi_{\omega jt} = \frac{r_{\omega jt}}{\sigma} - f_{njt} \quad (19)$$

and positive if $r_{\omega jt} > r_{njt}^*$, where

$$r_{njt}^* = \sigma f_{njt}. \quad (20)$$

If each firm ω self-select itself into market j on the condition that $r_{\omega jt} > r_{njt}^*$, as productivity models of exporting suggest, then the variance of revenues of observed firms in j is less than the true population variance (Greene 2003, pg. 759):

$$\text{var}(r_{\omega jt} | r_{\omega jt} > r_{njt}^*) < \text{var}(r_{\omega jt}) \quad (21)$$

which biases our estimate of Q_n^2 upwards. If we are to believe Melitz (2003), we need to correct for this truncation bias. To do so, we assume that log revenue $\ln r_{\omega jt}$ is normally distributed with mean $E[\ln r_{\omega jt}]$ and standard deviation s_n and all firms with $r_{\omega jt} > r_{njt}^*$ are observed in export market njt . This normality assumption is supported by the distribution of domestic revenues of Danish firms presented in Figure 2 and is consistent with previous studies of firm size distributions (Cabral and Mata, 2003) and export selection-bias (Helpman, Melitz, and Rubinstein, 2008).

[Insert Figure 2 here]

For the normally distributed truncated variable $\ln r_{\omega jt}$, Greene (2003, pg. 759) shows that the truncated expectation $E[\ln r_{\omega jt} | r_{\omega jt} > r_{njt}^*]$ is greater than the unconditional mean $E[\ln r_{\omega jt}]$ by:

$$E[\ln r_{\omega jt} | r_{\omega jt} > r_{njt}^*] = E[\ln r_{\omega jt}] + s_n \lambda(\gamma_{njt}^*) \quad (22)$$

$$\gamma_{njt}^* = \frac{\ln r_{njt}^* - E[\ln r_{\omega jt}]}{s_n} \quad (23)$$

where $\lambda(\gamma) = \frac{\phi(\gamma)}{1 - \Phi(\gamma)}$ is the inverse Mills ratio in which $\phi(\gamma)$ and $\Phi(\gamma)$ are the standard normal pdf and cdf's, respectively, and γ_{njt}^* is the cutoff revenue normalized to a standard normal distribution. Following Greene (2003, pg 760), we approximate the percentage of firms greater than the cutoff as

$$1 - \Phi(\gamma_{njt}) = P(r_{\omega jt} > r_{njt}^*) = \frac{|n_{jt}|}{|n_t|}, \quad (24)$$

where $|n_{jt}|$ is the observed number of firms in market njt and $|n_t|$ is the observed total number of firms exporting the product at time t . $\frac{|n_{jt}|}{|n_t|}$ is the sample proportion of firms that can profit in j . We can then use properties of Φ to estimate γ_{njt}^* with $\hat{\gamma}_{njt}$:

$$\hat{\gamma}_{njt} = \Phi^{-1} \left(1 - \frac{|n_{jt}|}{|n_t|} \right). \quad (25)$$

By proxying r_{njt}^* with the lowest firm sales in market njt , and estimating the truncated mean $E[\ln r_{\omega jt} | r_{\omega jt} > r_{njt}^*]$ with the observed mean $E_{njt}[\ln r_{\omega jt}]$, we can manipulate equations (11) and (22) to derive $\widehat{E[\ln r_{\omega jt}]}$, an approximation for $E[\ln r_{\omega jt}]$:

$$\widehat{E[\ln r_{\omega jt}]} = E_{njt}[\ln r_{\omega jt}] - s_n \lambda (\hat{\gamma}_{njt}) \quad (26)$$

$$= \frac{\phi(\hat{\gamma}_{njt}) \ln r_{njt}^* - (n_{jt}/n_t) \hat{\gamma}_{njt} E_{njt}[\ln r_{\omega jt}]}{\phi(\hat{\gamma}_{njt}) - (n_{jt}/n_t) \hat{\gamma}_{njt}} \quad (27)$$

and construct an untruncated measure of $\tilde{r}_{\omega jt} = \ln r_{\omega jt} - \widehat{E[\ln r_{\omega jt}]}$. To be thorough, we report results using revenues relative to both the corrected population mean ($\tilde{r}_{\omega jt} = \ln r_{\omega jt} - \widehat{E[\ln r_{\omega jt}]}$) and the truncated observed mean ($\tilde{r}_{\omega jt} = \ln r_{\omega jt} - E_{njt}[\ln r_{\omega jt}]$).

4.2 Coefficient of Determination Bias

The reader might notice that Q_n^2 resembles the adjusted coefficient of determination $\bar{R}^2 = 1 - \frac{\text{var}(\tilde{x}_{\omega jt})}{\text{var}(\tilde{r}_{\omega jt})}$ obtained when regressing $\tilde{r}_{\omega jt}$ on $\tilde{\varphi}_{\omega t}$ with residuals $\tilde{x}_{\omega jt}$ ⁸. We did too and use this familiar statistic. However, if our measure of $\tilde{\varphi}_{\omega jt}$ and the error $\tilde{x}_{\omega jt}$ are correlated, the variation in $\tilde{r}_{\omega jt}$ caused by the covariance between $\tilde{\varphi}_{\omega jt}$ and $\tilde{x}_{\omega jt}$ is attributed entirely to $\tilde{x}_{\omega jt}$. The correlation makes \bar{R}^2 an upwards biased estimate of Q_n^2 , no matter the sign of the correlation⁹. Firm specific and firm-destination specific shocks would be correlated if φ denoted the firm's quality. Firms with higher qualities would draw higher demand shocks in all destinations. The Apple iPod is a contemporary example of a variety of the product "mp3 player" whose costs of production are higher than most of its competitors and whose success is driven primarily by demand, which is highly correlated across markets. The correlation between demand shocks would be attributed to the firm-specific quality. Conversely, $\varphi_{\omega t}$ and $x_{\omega jt}$ could be negatively correlated if $x_{\omega jt}$ denoted non-iceberg trade costs. For example, suppose pencil firm A is more productive than firm B such that A can price its pencils at half the price of B's, which results in three times the quantity demanded. If both faced the same shipping

⁸We use \bar{R}^2 instead of the unadjusted R^2 because many of our regressions at the 8-digit CN product level contain a small number of observations. Cramer (1987) shows that R^2 is heavily biased for small samples. For example, if we estimate a model with a true population $R^2 = 0.30$ with a sample of 10 observations and 3 explanatory variables, we would expect to find a sample $R^2 = 0.45$: an upward bias of 50%. Even at $m = 20$ observations, our expected sample $R^2 = 0.37$. Cramer (1987) shows that \bar{R}^2 is much more accurate in small samples: in the same example with 10 observations, the expected sample $\bar{R}^2 = 0.29$. Even with only 5 observations, the expected sample $\bar{R}^2 = 0.28$. The \bar{R}^2 bias is only slightly downwards and never greater than a couple of percentage points. However, since the marginal cost of doing so is minimal, we also include R^2 in our results.

⁹Proof that $\bar{R}^2 \geq Q^2$ is in Section A.

costs per pencil, firm A would garner 50% more revenue, but pay three times as much in shipping costs¹⁰.

To circumvent this problem, we regress $\tilde{r}_{\omega jt}$ on only firm dummies and obtain the firm-specific coefficients λ_{it} . By construction, \bar{R}^2 is then the variation explained by all firm specific effects. All characteristics of a firm ω that increases or decreases the values of all its exports relative to other producers of n are captured by $\lambda_{\omega t}$, including price, quality, plant location, brand awareness, etc. Since productivity is only one of the many factors that distinguishes a firm, the resulting \bar{R}^2 would then be an upper bound of the contribution of productivity to the variation of $\tilde{r}_{\omega jt}$. As an added exercise, we find the variation of $\lambda_{\omega t}$ that is explained by usual measures of productivity TFP and value added per worker.

5 Estimating Q_n^2 with \bar{R}^2

With the caveats discussed, we measure estimate Q_n^2 with \bar{R}^2 from regressions of $\tilde{r}_{\omega jt}$ on $\tilde{\varphi}_{\omega t}$. Our measure of $\varphi_{\omega t}$ is Total Factor Productivity (TFP). To calculate TFP, we regress the log forms of the value added for firm ω at time t on its capital K and labor L inputs

$$\log(VA_{\omega t}) = \alpha_t + \gamma_K \log(K_{\omega t}) + \gamma_L \log(L_{\omega t}) + TFP_{\omega t}, \quad (28)$$

where $TFP_{\omega t}$ are the residuals. The regression of equation (28) yields coefficients to capital (0.12) and labor (0.86) that are roughly in line with the literature¹¹. For detailed results, see Table 3. TFP has historically been a rough measure of firm level productivity and correlates positively with firm revenues, which we show in Figure 3. As the figure caption indicates, a regression of TFP on the log of total revenues results in a positive coefficient of 0.05 that is statistically significant at the 1% level. The results do not vary greatly when we regress by sectors, by years, or restrict the sample to only firms with greater than 50 workers.

[Insert Table 3]

[Insert Figure 3]

Although TFP¹² is imprecise, Figure 3 shows that it is positively correlated to firm revenues, which makes it a possible proxy for $\varphi_{\omega t}$. We estimate equation (14) using

¹⁰For a detailed discussion of the effects of non-iceberg trade costs, see Hummels & Skiba (2004)

¹¹Olley and Pakes (1996) find OLS coefficients of 0.17 and 0.85 for labor and capital, respectively, for a balanced panel of telecommunications firms. Pacvnik (1999) splits labor into unskilled and skilled and adds materials as a regressors. Across all 8 broadly defined industries, she finds the sum of the OLS coefficients on skilled and unskilled labor to be approximately 7 times greater than that on capital, in line with ours and Olley and Pakes (1996).

¹²Recent literature has explored possible causes behind the imprecision of using TFP as a productivity measure. The coefficients to the inputs may be plagued by endogeneity bias which may affect the

$\varphi_{\omega t} = TFP_{\omega t}$ and $\tilde{r}_{\omega jt} = \ln r_{\omega jt} - E[\widehat{\ln r_{\omega jt}}]$ to correct to truncation bias:

$$\left(\ln r_{\omega jt} - E[\widehat{\ln r_{\omega jt}}] \right) = -0.02 * TFP_{\omega t} + \varepsilon_{\omega jt} \quad (29)$$

The -0.02 coefficient is the statistically insignificant estimate of the marginal effect of $TFP_{\omega t}$ from the regression considering all Danish exports as a single product. Results do not change when the data are separated by SITC Rev. 3 top level sectors, as summarized in Table 4. TFP does not have any impact on relative revenues: none of the coefficients are significantly different from zero. All $R^2 = 0$: relative TFP does not explain any of the variation in relative revenues.

Insert Table 4 here

We also estimate equation (30) for each of 2302 8-digit CN products. When we define our markets at this granulation, we drastically reduce the number of observations per regression - 264 regressions were discarded because they had only 2 observations. We find that $TFP_{\omega t}$ hardly explains the variation of relative revenue. Our median \bar{R}^2 is less than 1% and most of the values are close to zero. Figure 4 shows how tightly the results are gathered around zero.

Insert Table 5 here

Insert Figure 4 here

We also regress equation (14) without correcting for truncation. Accordingly, we use the observed productivity in market njt conditional on entry:

$$(\ln r_{\omega jt} - E_{njt}[\ln r_{\omega t}]) = 0.33 * (TFP_{\omega t} - E_{njt}[TFP_{\omega t}]) + \varepsilon_{\omega jt} \quad (30)$$

The 0.33 coefficient on $(TFP_{\omega t} - E_{njt}[TFP_{\omega t}])$ in equation (30) corresponds to the regression considering all Danish exports as varieties of a single product. It shows a positive and much stronger correlation between relative productivity and relative revenue, but these results are biased upwards if the firms self-select into markets. Similarly, \bar{R}^2 is still basically zero: TFP explains less than 0% of the variation in relative revenues for the entire economy, and 0% for 3 of the 4 sectors. Complete results of regressions of all observations and separated by SITC Rev. 3 top level sectors are summarized in Table 6:

Insert Table 6 here

residuals (see e.g. Akerberg, Benkard, Berry & Pakes (2007)). In addition, many firm and industry characteristics are not captured by equation (28).

This pattern continues as we disaggregate down to the CN 8-digit level. Table 7 summarizes the results.

Insert Table 7 here

At this lowest level of aggregation and without correction for truncation, TFP can only explain on average 16% of the variation of relative sales in a destination, and explains less than 4% of the variation for half of the products. The difference between the mean and median is partially driven by high \bar{R}^2 's for regressions with few observations; if we restrict our regressions to only those with greater than 5 observations, the mean \bar{R}^2 drops to 0.13. Increasing the restriction to >10 observations further reduces it to 11%. The distribution of \bar{R}^2 presented in Figure 5 shows that most of the \bar{R}^2 are very close to 0. As a measure of firm-specific productivity, TFP cannot explain much of the variation of firm sales in a destination.

Insert Figure 5 here

5.1 Results: Dummy Variables

Turning to our second approach, we use firm-specific dummy variables $\lambda_{\omega t}$ as regressors. Since our model assumes single product firms, we treat multiple products from one firm as separate products from separate firms. Our dummy variable actually captures firm-product variation, not just firm variation.

First, we regress the truncated corrected relative revenues on firm-product dummies:

$$\left(\ln r_{\omega jt} - E[\widehat{\ln r_{\omega jt}}] \right) = \lambda_{\omega t} + x_{\omega jt} \quad (31)$$

for the entire sample (i.e. for the all Danish firms across all product categories across all four years) and by SITC1 industries without correcting for truncation. Results are summarized in Table 8. There is clearly variation among firms: the reported F-test statistic for the hypothesis that all firms are equal correspond to a probability of 0.0000. However, the variation in firm dummy coefficients explains only 17% of the sales variation within a destination if we consider all Danish exports as varieties of a single product. Regressing by top-level sectors, we find \bar{R}^2 between 13% for Machinery and 30% for Miscellaneous Manufactures. We also perform these regressions at the 8-digit CN product category, as seen in Figure 6. At this level of disaggregation, the mean and median \bar{R}^2 are 13% and 6%, respectively. As Figure 6 shows, there are some products that exhibit a very high \bar{R}^2 , but most exhibit an \bar{R}^2 less than 25%. There is no correlation between \bar{R}^2 and the number of firms in the regression or the number of observations. We also find no correlation between the firm-sales weighted capital/labor ratio for a product and the product's \bar{R}^2 .

Insert Table 8 here

Insert Figure 6 here

We also report results from estimations that do not correct for truncation. These results are obtained from regressing the ωjt revenue relative to the observed mean in market njt on firm-product dummies:

$$(\ln r_{\omega jt} - E_{njt}[\ln r_{\omega jt}]) = \lambda_{\omega t} + x_{\omega jt} \quad (32)$$

\bar{R}^2 for the uncorrected regressions at the entire economy (0.12) and top sector levels (0.14 – 0.23) are similar to those obtained with corrected regressions. Table 9 summarizes.

Insert Table 9 here

Calculations of \bar{R}^2 at the 8-digit CN product are much higher for the uncorrected regressions than for the corrected ones, reflecting the truncation bias. We obtain a mean and median \bar{R}^2 of 39%, and a standard deviation of 29%. Figure 7 shows the distribution of \bar{R}^2 across all products within the four SITC1 industries.

Insert Figure 7 here

Because of the construction of the statistic, approximately 8% of \bar{R}^2 are lower than zero. We cannot interpret these within the confines of our model. If we round these errant statistics to zero, we increase our average \bar{R}^2 to 41%. For completeness, we also calculate unadjusted R^2 for the same regressions and find a mean and median of 49-50% with a standard deviation of 23%. This result is in line with Eaton, Kortum and Kramarz’s (2007) result that firm-specific heterogeneity accounts for approximately half of total variation. However, as Garner (1987) shows, R^2 has a significant upward bias in small samples. Predictably, if we restrict our results to R^2 resulting from regressions with greater than ten observations, the mean drops to 39%.

5.2 The contribution of $\tilde{\varphi}_{\omega t}$ to the variation in $\lambda_{\omega t}$

Our firm dummy regressions separate firm-specific effects and firm-destination-specific effects. The resulting dummy coefficient $\lambda_{\omega t}$ captured the contributions of all firm-specific characteristics, including productivity. Since $TFP_{\omega t}$ is a proxy for productivity, we can use it to estimate how much firm-specific variation is explained by productivity. To do so, we regress $\lambda_{\omega t}$ on $TFP_{\omega t}$

$$\lambda_{\omega t} = b * TFP_{\omega t} + \varepsilon_{\omega t} \quad (33)$$

with b as the marginal effect of $TFP_{\omega t}$ and $\varepsilon_{\omega t}$ as the residual. The results, collected in Table 10, do not favor the interpretation of $TFP_{\omega t}$ as productivity. At the economy and sector level, the estimated marginal effects were indistinguishable from zero for every sector except Chemicals. \bar{R}^2 calculations were all less than a percent. This pattern continues at the product level, as shown in Figure 8. While $TFP_{\omega t}$ explains over 30% of the variation of firm-specific effects for 10% of products, most products exhibit \bar{R}^2 below 1%.

[Insert Figure 8]

Since our results from regressing $\lambda_{\omega t}$ on $TFP_{\omega t}$ did not endorse TFP as an explanatory variable, we question whether $TFP_{\omega t}$ is a proper measure of $\varphi_{\omega t}$. We also considered Value Added per worker (VApw) as a measure of productivity. We regress $\lambda_{\omega t}$ on $VApw_{\omega t}$:

$$\lambda_{\omega t} = b * VApw_{\omega t} + \varepsilon_{\omega t} \quad (34)$$

with b as the marginal effect of $VApw_{\omega t}$ and $\varepsilon_{\omega t}$ as the residual and report the results at high levels of aggregation in Table 11 and by product in Figure 9. The estimated b 's are significant and positive for $VApw_{\omega t}$, suggesting that $VApw_{\omega t}$ is a better predictor of the average relative firm revenue across destinations than $TFP_{\omega t}$ is. However, the barely-different-from zero \bar{R}^2 calculations suggests that it too cannot explain the variation in average Danish firm exports.

[Insert Table 11]

[Insert Figure 9]

These results suggest that either that productivity is not a significant contributor to firm-specific effects, or that both $TFP_{\omega t}$ and $VApw_{\omega t}$ are poor proxies for productivity.

6 Estimating β_n with revenue changes

Since $TFP_{\omega t}$, and to a lesser extent, $VApw_{\omega t}$, cannot explain even the variation in firm-specific fixed effects, we loathe to use these variables to estimate β_n . However, we can still estimate β_n by exploiting serial correlations in $\tilde{\varphi}_{\omega t}$ and $\tilde{x}_{\omega jt}$. Suppose $\tilde{\varphi}_{\omega t}$ and $\tilde{x}_{\omega jt}$ both follow random walks:

$$\tilde{\varphi}_{\omega t} = \tilde{\varphi}_{\omega t-1} + \varepsilon_{\omega t} \quad (35)$$

$$\tilde{x}_{\omega jt} = \tilde{x}_{\omega jt-1} + \varepsilon_{j\omega t} \quad (36)$$

where $\varepsilon_{\omega t}$ and $\varepsilon_{j\omega t}$ are annual disturbances and $E[\varepsilon_{\omega t}\varepsilon_{j\omega t}] = 0$, $E[\varepsilon_{\omega t}\varepsilon_{\omega s}] = 0$ for $s \neq t$, and $E[\varepsilon_{i\omega t}\varepsilon_{j\omega t}] = 0$ for $i \neq j$. The random walk specifications indicate that the

best estimate of this period's productivity of (destination-specific demand shock for) firm ω 's variety is the previous year's productivity (demand shock). Defining $\Delta\kappa_t = \kappa_t - \kappa_{t-1}$ for $\kappa \in \{\tilde{r}_{\omega j}, \tilde{\varphi}_\omega, \tilde{x}_{\omega j}\}$ and combining equations (14), (35), (36) and the definition of β_n from (15), we can construct equations relating the change in relative revenues for in market j for ω over time:

$$\beta_n \Delta \tilde{r}_{\omega jt} = \varepsilon_{\omega t} \quad (37)$$

$$\Delta \tilde{r}_{\omega jt} = \varepsilon_{\omega t} + \varepsilon_{j\omega t} \quad (38)$$

As equation (37) indicates, a change in firm-specific productivity $\Delta\tilde{\varphi}_\omega = \varepsilon_{\omega t}$ results in a proportional β_n change in relative revenues of a firm in any supplied market j . Since this productivity change pervades to all destinations, we can combine equation (37) for destination i and (38) for destination j to arrive at our regression equation:

$$\Delta \tilde{r}_{\omega jt} = \beta_n \Delta \tilde{r}_{\omega it} + \varepsilon_{j\omega t} \quad (39)$$

Equation (39) states that the contribution of firm-specific productivity to relative market sales can be measured by the correlation between the contemporary changes in relative sales in any two markets.

6.1 Regression results

We estimate β_n in equation (39) by regressing the annual change in the relative revenues for a firm in its largest export destination on the change in relative revenues for the same firm in its second largest export destination. We restrict our observations to product-destination-markets in which greater than 10 firms operate each year and correct for truncation.

Table 12 shows that at the economy and sectoral level, β_n is consistently estimated to be approximately 0.10, implying that 10% of the annual change in a firm's relative market revenue can be attributed to destination-invariant changes within the firm. Figure 10 paints a worse picture for productivity: the distribution of β_n at the 8-digit product level is centered approximately about zero. The mean and median 8 digit CN product β_n is -0.45 and -0.08, respectively. However, approximately 25% of the 1278 betas are results of regressions that do not have enough observations to produce a standard error. When restricting the regressions to only those with enough observations to produce standard errors, the mean and median becomes -0.09 and -0.10, respectively. Restricting the regressions to those with greater than 10 observations does not change the results. A negative β_n indicates that, on average, an increase in a firm's revenue in its largest market actually dictates a decrease in its second largest. Productivity models cannot produce this prediction; perhaps destination-specific demand shocks, or production constraints can explain the negative β_n 's.

Insert Table 12
 Insert Figure 10

We also report results from regressions using uncorrected relative revenues $\tilde{r}_{\omega jt} = \ln r_{\omega jt} - E_{njt}[\ln r_{\omega jt}]$. Table 13 summarizes the results at the economy and sector level. The results, as expected, are closer to 1 due to the bias, but still remarkably low. β_n is estimated to be approximately 0.20 implying that annual firm-specific changes account for 20% of the change in revenues in the top two export destinations. The results are much less favorable at the 8 digit product level. Figure 11 shows the distribution of β_n for 8-digit products for each sector. The distributions are all centered around zero, with a mean of -0.32 and a median of 0.17. If we eliminate those regressions without standard errors, the mean and median β'_n s are both 0.16. Restricting the regressions to those with greater than 10 observations does not change the mean and median β'_n s significantly (0.16 and 0.15, respectively). This suggests that, even if firms randomly enter markets, productivity changes account for less than a fifth of a firm's revenue change in any market.

Insert Table 13
 Insert Figure 11

7 Alternative explanations

Our empirical results show that productivity cannot explain much of the variation of firm sales in a market. Nguyen (2008) argues that firm-destination-specific demand shocks can explain the variation. However, our empirical strategy consists of separating firm-specific effects from firm-destination-specific effects, the interpretations of these effects are up for debate. A destination specific demand shock story is only one of several possible explanations. One could argue that the differences in relative revenues among markets arise from measurement error or firm-destination-specific trade costs. In this section, we calculate the variation in measurement error or firm-destination-specific trade costs needed in order to explain the total variation.

7.1 Measurement Error

Suppose that measurement error is causing our low estimates of Q^2 . That is, suppose $x_{j\omega t}$ in equation (5) is measurement error in $r_{\omega jt}$. We can calculate what degree of mismeasurement is required to achieve our results. Since \bar{R}^2 is defined as

$$\bar{R}^2 = 1 - \frac{\text{var}(x_{j\omega t})}{\text{var}(\tilde{r}_{\omega jt})} \quad (40)$$

we calculate the variation in $x_{j\omega t}$ necessary to achieve calculations of $\bar{R}^2 = \hat{R}^2$:

$$var(x_{j\omega t}) = (1 - \hat{R}^2) var(\tilde{r}_{\omega jt}) \quad (41)$$

Using truncated corrected relative revenues, and firm dummies we estimate an average $\hat{R}^2 = 0.13$ across 8-digit products. The variance of log export values for Danish exporters is $var(\tilde{r}_{\omega jt}) = 2.8$, which results in a required $s_x^2 \geq 2.4$. The average recorded revenue for a Danish export is $r_{\omega jt} = \$9,800$. To obtain $var(\tilde{r}_{\omega jt}) = 2.8$, that \$9,800 must be mismeasured anywhere from \$2,080 to \$46,000 given a conservative one standard deviation error bandwidth. If we use uncorrected revenues, we estimate an average $\bar{R}^2 = 0.37$ across 8-digit products. Using this more conservative metric, the same \$9,800 export still has a one standard deviation bandwidth of \$2,700 to \$36,000.

We apply the same principles to our estimates of β_n . If we assume $\tilde{x}_{\omega jt}$ is due to measurement error and therefore not serially correlated, the resulting regression equation (39) becomes

$$\Delta \tilde{r}_{\omega jt} = \beta_n \Delta \tilde{r}_{\omega it} + \tilde{x}_{\omega jt} - \tilde{x}_{\omega jt-1} \quad (42)$$

where $E[\tilde{x}_{\omega jt}\tilde{x}_{\omega jt-1}] = 0$. Taking the variance of both sides and rearranging, we find that the variance of $\tilde{x}_{\omega jt}$ required to achieve an estimated value of $\beta_n = \hat{\beta}_n$ given that the true $\beta_n = 1$ is

$$var(\tilde{x}_{\omega jt}) = (1 - \hat{\beta}_n)^2 var(\tilde{r}_{\omega jt}) \quad (43)$$

Given our average sample $var(\tilde{r}_{\omega jt}) = 2.8$, $\hat{\beta}_n = 0.15$, a $var(\tilde{x}_{\omega jt}) = 2.0$ is required to obtain our results. That variance implies that our \$9,800 average recorded export value has a one standard deviation error bandwidth of \$2,400 to \$39,000.

Our data is customs trade data from which tariff revenues are calculated. It does not seem plausible to have errors that large. However, if Danish customs agents are somehow mismeasuring the revenues by these proportions, this should be reflected in the variance of market revenues for a firm over time. As a check, we use the annual difference in a firm's export to estimate measurement error. Suppose that a firm's log revenue $\ln r_{\omega jt}$ is constant over all four years and that any deviation is due entirely to measurement error $x_{\omega jt}$, which is uncorrelated over time:

$$\ln r_{\omega jt} = \ln r_{\omega j} + x_{\omega jt} \quad (44)$$

Therefore, the variance of the changes in log destination revenues for a variety in destination j over the years is twice the variance of measurement error:

$$var(\ln r_{\omega jt} - \ln r_{\omega jt-1}) = var(x_{\omega jt} - x_{\omega jt-1}) \quad (45)$$

$$= 2var(x_{\omega jt}) \quad (46)$$

since $x_{\omega jt}$ and $x_{\omega jt-1}$ share the same variance. Figure 12 shows the distribution of the standard deviations for 43921 variety-destination pairs which existed for all four years of our data. The mean and median standard deviations of change are 1.03 and 0.74, respectively, which correspond to mean and median $var(\tilde{x}_{\omega jt})$ of 0.55 and 0.27, respectively. These values are much lower than the $var(\tilde{x}_{\omega jt}) \geq 2$ needed to obtain our results.

Insert Figure 12

7.2 Trade Costs

Suppose our firm-destination-specific heterogeneity is primarily due to firm-destination specific trade costs, and that variation is the cause of our low estimates. That is, suppose that the revenue equation (5) is of the form

$$r_{\omega jt} = \frac{\alpha_{nt}^{1-\sigma} (\tau_{j\omega t})^{\sigma-1} Y_{jt}}{P_{jnt}} \exp \varphi_{\omega t} \quad (47)$$

where $\ln \tau_{j\omega t}$ has a mean of τ_{njt} and a variance of s_τ^2 . The regression equation corresponding to equation (14) is

$$\tilde{r}_{\omega jt} = \tilde{\varphi}_{\omega t} + (\sigma - 1) \ln \tilde{\tau}_{\omega jt} \quad (48)$$

where $(\sigma - 1) \ln \tilde{\tau}_{\omega jt}$ is the residual error when regressing $\tilde{r}_{\omega jt}$ on firm-specific dummies. As in the previous section, we can calculate the variance of $\ln \tilde{\tau}_{\omega jt}$ necessary to achieve our regression results. To obtain an $\bar{R}^2 = \hat{R}^2$ in our estimations, $var(\ln \tilde{\tau}_{\omega jt})$ must be:

$$var(\ln \tilde{\tau}_{\omega jt}) = \frac{(1 - \hat{R}^2) var(\tilde{r}_{\omega jt})}{(\sigma - 1)^2} \quad (49)$$

The average variance of relative firm revenues is $var(\tilde{r}_{\omega jt}) = 2.8$. Using our mean estimate of $\hat{R}^2 = 0.13$ and the value of $\sigma = 3$ taken from the recent literature (Hummels 2001, Anderson and Van Wincoop 2003), we find $var(\ln \tilde{\tau}_{\omega jt}) = 0.41$ to account for the low \hat{R}^2 . A $var(\ln \tilde{\tau}_{\omega jt})$ of 0.41 implies a typical firm faces ad valorem trade costs that are 18 times greater or lower than the average trade costs across all firms exporting the same product to the same destination. If an average firm faces trade costs that increase the price of their good by 50%, as Anderson and Van Wincoop (2004) suggests, a firm facing trade costs that are a single standard deviation from that mean experiences ad valorem rates of 227%. It seems implausible that a firm that faces 227% trade costs could compete with firms selling the same product to the same destination with average trade costs of 50%. If we use an average $\tilde{\tau}_{\omega jt} = 1.1$ instead of $\tilde{\tau}_{\omega jt} = 1.5$, firms facing

trade costs one standard deviation higher than the mean still face trade costs greater than 140% - 14 times the

Similar to the analysis for measurement error, firm-destination-specific trade costs can bias β_n downwards to $\hat{\beta}_n$ if

$$var(\ln \tilde{\tau}_{\omega jt}) = \frac{(1 - \hat{\beta}_n)^2 var(\tilde{r}_{\omega jt})}{(\sigma - 1)^2} \quad (50)$$

Given our average sample $var(\tilde{r}_{\omega jt}) = 2.8$, $\hat{\beta}_n = 0.15$, and our estimate of $\sigma = 3$, a $var(\ln \tilde{\tau}_{\omega jt}) = 0.4$. Again, this variance results in an implausible spread of trade costs.

Our results for $var(\ln \tilde{\tau}_{\omega jt})$ depend on the value of σ , which we take from the literature to equal three. However, another we can also determine the value of σ necessary to obtain plausible values of $var(\ln \tilde{\tau}_{\omega jt})$. To aid exposition, we present Figure 13, which plots the value of τ that is one standard deviation greater than a mean $\tau = 1.5$.

[Insert Figure 13]

The spread in trade costs remain implausibly high for $\sigma < 10$. At a high estimate of $\sigma = 13$, firms one standard deviation higher than the mean face $\tau = 1.7$, which is still 40% greater than average. If we believe that $x_{\omega jt}$ are firm specific trade costs, our results suggest that σ is higher than previous estimates (Hummels 2001).

8 Conclusions

In this paper, we estimate the contribution of firm productivity to the heterogeneity of revenues in a market and find it to be remarkably low. When controlling for self-selection truncation bias, we find that firm-specific factors contribute from 0 to 30% across our different estimates. When using firm dummies as the broadest interpretation of productivity, we find that firm-specific effects can only explain 13-18% relative revenue variation at the economy, sector, and product levels. We also show that the unexplained firm-destination variation cannot plausibly be interpreted as caused by measurement error or firm-destination-specific trade costs. Our results suggest that firm-specific productivity is not capturing the majority of heterogeneity and is not the primary driver of variation in a market. Our results should not be interpreted as a condemnation of firm-specific productivity models such as Melitz(2003); the model deftly explains variation between exporters and nonexporters. However, Melitz (2003) is limited to firm-specific differences, and we show that the majority of variation is firm-destination specific. Nguyen (2008) shows how this variation can be generated with a single mechanism involving demand heterogeneity. In it, he presents a model in which firms test destinations and

receive firm-destination-specific *perceived quality* draws. Higher *perceived qualities* result in higher sales. Since demands are firm-destination-specific, a firm can have high relative sales in one destination but low relative sales in another. Productivity heterogeneity models cannot generate this sales ranking inversion. I reconcile higher average domestic sales for exporters than for nonexporters by correlating a firm's *perceived qualities* with a firm-specific but unknown-to-the-firm *latent quality*. Since demand is thought of as more uncertain than costs, models incorporating firm-destination-specific demand heterogeneity may better explain the variation of sales.

A Appendix: Proof that $\bar{R}^2 \geq Q^2$

Consider a sample of M observations with each observation m having the relation

$$y_m = \beta x_m + \varepsilon_m \quad (51)$$

where y_m and x_m are observed and ε_m unobserved and possibly correlated with x_m . An OLS of y on x results in parameter estimate $\hat{\beta}$ for β with residuals

$$\hat{\varepsilon}_m = y_m - \hat{\beta}x_m \quad (52)$$

We decompose the sum of the squared residuals:

$$\hat{\varepsilon}'\hat{\varepsilon} = \sum_{m=1}^M \hat{\varepsilon}_m^2 \quad (53)$$

$$= \sum_{m=1}^M \left(\beta x_m + \varepsilon_m - \hat{\beta}x_m \right)^2 \quad (54)$$

$$= \sum_{m=1}^M \varepsilon_m^2 + \left(\beta - \hat{\beta} \right)^2 \sum_{m=1}^M x_m^2 + 2 \left(\beta - \hat{\beta} \right) \sum_{m=1}^M \varepsilon_m x_m \quad (55)$$

By the construction of ordinary least squares,

$$\beta - \hat{\beta} = - \frac{\sum_{m=1}^M \varepsilon_m x_m}{\sum_{m=1}^M x_m^2} \quad (56)$$

We can insert the difference between β and $\hat{\beta}$ from equation (56) into equation (55):

$$\begin{aligned} \hat{\varepsilon}'\hat{\varepsilon} &= \sum_{m=1}^M \varepsilon_m^2 + \left(\frac{\sum_{m=1}^M \varepsilon_m x_m}{\sum_{m=1}^M x_m^2} \right)^2 \sum_{m=1}^M x_m^2 - 2 \left(\frac{\sum_{m=1}^M \varepsilon_m x_m}{\sum_{m=1}^M x_m^2} \right) \sum_{m=1}^M \varepsilon_m x_m \\ \hat{\varepsilon}'\hat{\varepsilon} &= \varepsilon'\varepsilon - \frac{\left(\sum_{m=1}^M \varepsilon_m x_m \right)^2}{\sum_{m=1}^M x_m^2} \end{aligned} \quad (57)$$

Since all terms on the right hand side are squares or a sum of squares, $\hat{\varepsilon}'\hat{\varepsilon} \leq \varepsilon'\varepsilon$. The terms are equal when ε and x are uncorrelated. The adjusted coefficient of correlation \bar{R}^2 :

$$\bar{R}^2 = 1 - \frac{\hat{\varepsilon}'\hat{\varepsilon}}{y'y} \frac{M-1}{M-K} \quad (58)$$

can now be assigned a bias when used as an estimate of

$$Q^2 = 1 - \frac{var(\varepsilon)}{var(y)} = 1 - \frac{\varepsilon' \varepsilon}{y' y} \frac{M-1}{M-K} \quad (59)$$

where Q^2 is the variation in y explained by x . Since $\hat{\varepsilon}' \hat{\varepsilon} \leq \varepsilon' \varepsilon$, $\bar{R}^2 \geq Q^2$: \bar{R}^2 is an upper bound of Q^2 .

B Appendix: Tables

Table 1: Products, destinations and firms by industry, 2002

SITC1 industry	Number of CN8 products	Destinations per product	Firms per product	Destinations per firm per product	Firms per destination per product
5 Chemicals	835	11.7	5.1	3.4	1.2
6 Manuf. Materials	1923	9.4	6.5	2.5	1.4
7 Machinery	1483	16.8	11.0	3.3	1.5
8 Misc. Manuf.	1068	16.7	11.3	3.3	1.6
Total	5309	13.2	8.5		

Note: Sample excludes firms with zero employees

Table 2: Notation

Notation	Description
i, j	A destination country. $i, j \in J$
s, t	time period (year) $s, t \in (1999, 2000, 2001, 2002)$
n	8 digit product category. $n \in N$
$ \ast $	The cardinality (number of elements) of set \ast
ω	The unique variety of a firm producing product n
u	utility
σ	elasticity of substitution between varieties of product n
α_n	Cobb-Douglas proportion of income spent on product n
$q_{\omega jt}$	The quantity of variety ω supplied to j in t
$p_{\omega jt}$	The price of variety ω supplied to j in t
τ_{jn}	The iceberg trade cost
P_{njt}	The price index for product n in j at t
$r_{\omega jt}$	The revenue of variety ω in j at t
$\tilde{r}_{\omega jt}$	The relative revenue of variety ω in j at t
Y_{jt}	The total expenditure of j
φ_ω	The firm-specific productivity of ω
$x_{j\omega}$	The firm-destination-specific demand shock for variety ω in j
s_φ^2, s_x^2	The variances of φ_ω and $x_{j\omega}$, respectively
$\bar{\varphi}_\omega, \bar{x}_{j\omega}$	The means of φ_ω and $x_{j\omega}$, respectively
β_n	The contribution of firm-specific productivity to the firm's relative revenues
Q_n^2	The theoretical proportion of total variance explained by productivity
R^2	The coefficient of correlation
\bar{R}^2	The adjusted coefficient of correlation

Table 3: TFP regressions

	All firms		Firms with >50 workers	
	Coeff.	Std. err.	Coeff.	Std.err.
Log(capital)	0.1210	0.0036	0.1290	0.0060
Log(labor)	0.8614	0.0048	0.9018	0.0095
Year 2000	0.0353	0.0109	0.0264	0.0156
Year 2001	0.0680	0.0109	0.0500	0.0157
Year 2002	0.0980	0.0109	0.0876	0.0158
Constant	-1.5456	0.0214	-1.8144	0.0437
No. of obs.	10440		3827	
R^2	0.9278		0.8749	

Table 4: Cross Sectional Regression results using TFP with truncation correction

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	327613	24817	54195	125879	122722
Number of firms	3706	638	1521	1870	2062
Coef. on TFP	-0.02	0.22	0.02	0.07	-0.18
Std. Err.	0.09	0.21	0.25	0.16	0.17
R^2	0.00	0.00	0.00	0.00	0.00
\bar{R}^2	0.00	0.00	0.00	0.00	0.00

Table 5: Cross Sectional Regression results using TFP with truncation correction at the 8-digit product level

Number of Regressions w/ observations > 2	2302
mean observations / regression	142
median observations / regression	40
mean R^2	0.08
median R^2	0.01
mean \bar{R}^2	0.02
median \bar{R}^2	0.00

Table 6: Cross Sectional Regression results using TFP without truncation correction

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	327613	24817	54195	125879	122722
Number of firms	3706	638	1521	1870	2062
Coef. on TFP	0.33	0.65	0.16	0.24	0.35
Std. Err.	0.01	0.03	0.04	0.02	0.02
R^2	0.00	0.02	0.00	0.00	0.00
\bar{R}^2	0.00	0.02	0.00	0.00	0.00

Table 7: Cross Sectional Regression results using TFP without truncation correction at the 8-digit product level

Number of Regressions w/ observations > 2	2302
mean observations / regression	142
median observations / regression	40
mean R^2	0.20
median R^2	0.08
mean \bar{R}^2	0.16
median \bar{R}^2	0.04

Table 8: Cross Sectional Regression firm dummy results with correction for truncation

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	330018	25197	54387	127086	123348
Number of firms	3706	638	1521	1870	2062
F-test statistic	19.01	9.77	13.57	11.00	26.14
R^2	0.18	0.20	0.28	0.14	0.31
\bar{R}^2	0.17	0.18	0.26	0.13	0.30

Table 9: Cross Sectional Regression with dummies, without truncation correction, results

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	330018	25197	54387	127086	123348
Number of firms	3706	638	1521	1870	2062
F-test statistic	13.53	12.85	11.20	12.02	13.70
R^2	0.13	0.25	0.24	0.15	0.19
\bar{R}^2	0.12	0.23	0.22	0.14	0.18

Table 10: Regression of firm-specific effects on TFP, corrected for truncation

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	52296	4226	11635	17681	18754
Coef	0.02	0.20	0.02	0.00	-0.01
pvalue	0.53	0.02	0.83	1.00	0.87
R^2	0.00	0.00	0.00	0.00	0.00
\bar{R}^2	0.00	0.00	0.00	0.00	0.00

Table 11: Regression of firm-specific effects on Value Added per worker, corrected for truncation

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	52458	4234	11655	17755	18814
Coef	0.11	0.30	0.12	0.09	0.08
pvalue	0.00	0.00	0.11	0.12	0.13
R^2	0.00	0.00	0.00	0.00	0.00
\bar{R}^2	0.00	0.00	0.00	0.00	0.00

Table 12: Equation 39 regression, with truncation correction

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	16207	1398	3429	5055	6325
β_n	0.11	0.08	0.10	0.10	0.12
Std. Err.	0.01	0.02	0.02	0.01	0.01
R^2	0.02	0.02	0.02	0.02	0.02
\bar{R}^2	0.02	0.02	0.02	0.02	0.02

Table 13: Equation 39 regression, without truncation correction

	All Sectors	Chemicals (SITC1=5)	Manuf. Materials (SITC1=6)	Machinery (SITC1=7)	Misc. Manuf. (SITC1=8)
Number of obs	16207	1398	3429	5055	6325
β_n	0.21	0.14	0.15	0.17	0.27
Std. Err.	0.01	0.05	0.02	0.01	0.01
R^2	0.05	0.04	0.02	0.05	0.08
\bar{R}^2	0.05	0.03	0.02	0.05	0.08

C Appendix: Figures

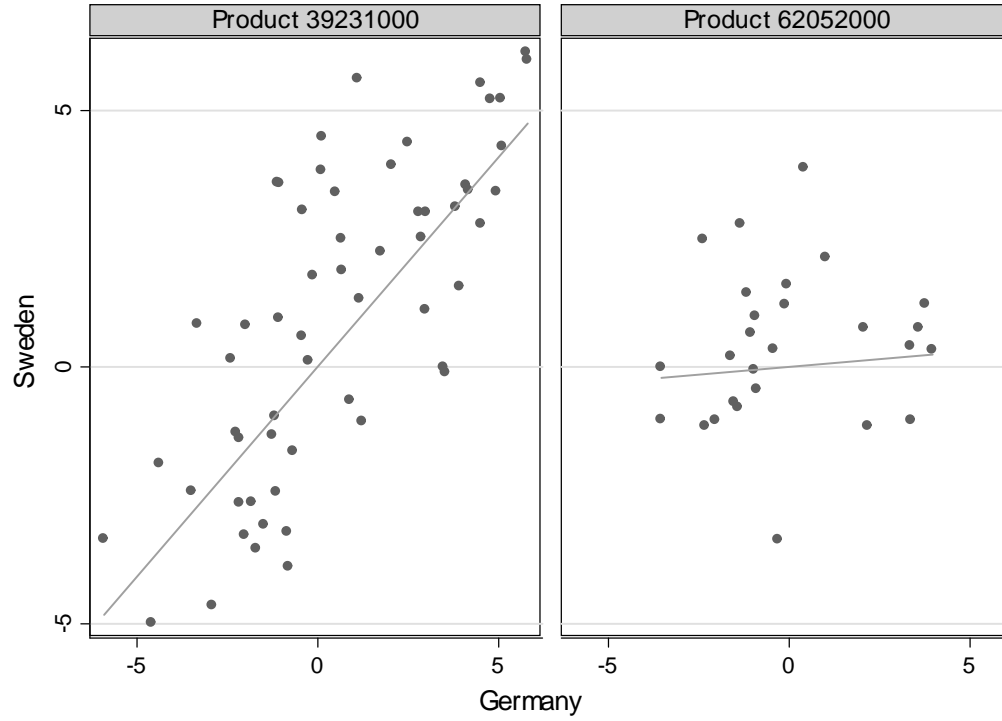


Figure 1: Sales relative to other Danish firms in Sweden and Germany for Danish exporters of plastic boxes (left panel) and Men's cotton shirts (right panel). Statistics for the lines with fitted values: Left panel: slope = 0.82, std.err. = 0.09, $R^2 = 0.57$. Right panel: slope = 0.06, std.err. = 0.14, $R^2 = 0.01$.




Figure 2: The distribution of log domestic revenues for Danish firms (1999-2002).




Figure 3: Scatter of a Danish exporter's TFP versus its total revenues across all destinations. Slope = 0.05, ($t - stat = 0.002$), $R^2 = 0.04$

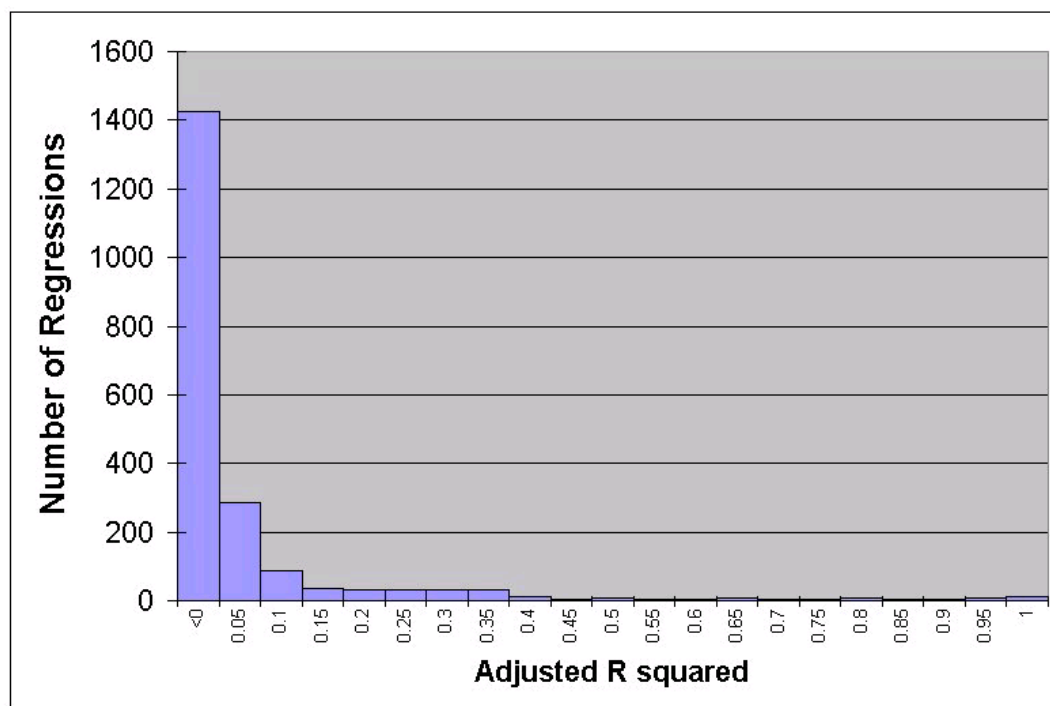


Figure 4: Histogram of \bar{R}^2 for TFP regressions at the 8-digit product level, corrected for truncation .

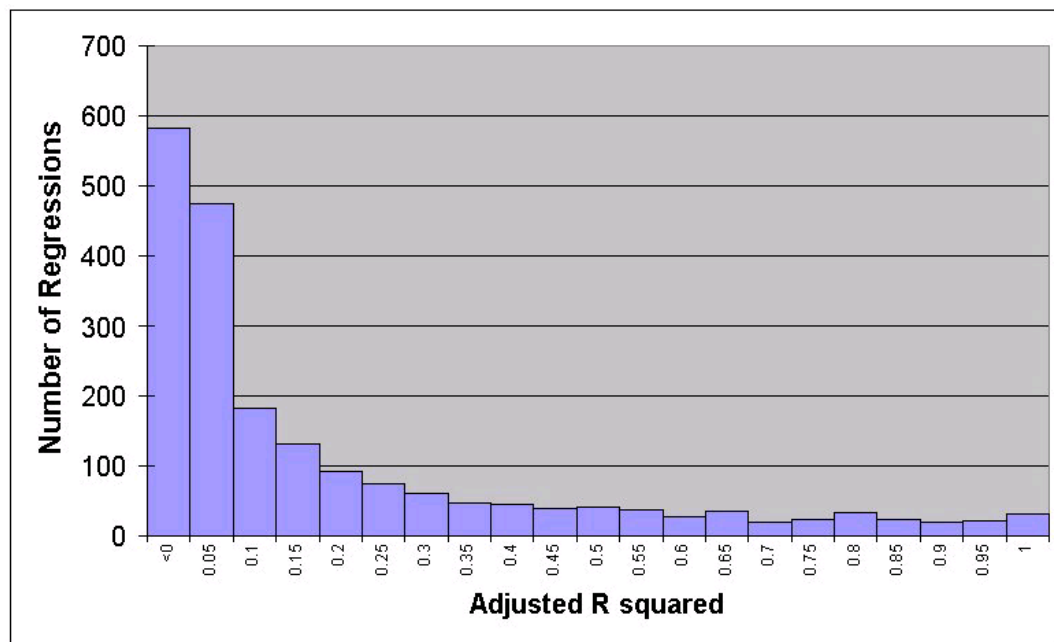


Figure 5: Histogram of \bar{R}^2 for TFP regressions at the 8-digit product level, without correction for truncation .




Figure 6: Histogram of \bar{R}^2 for firm dummy regressions at the 8-digit product level with correction for truncation.




Figure 7: Histogram of \bar{R}^2 for firm dummy regressions at the 8-digit product level, without correction for truncation,




Figure 8: Distributions of \bar{R}^2 for regressions of firm-specific fixed effects on total factor productivity at the 8-digit CN product level.

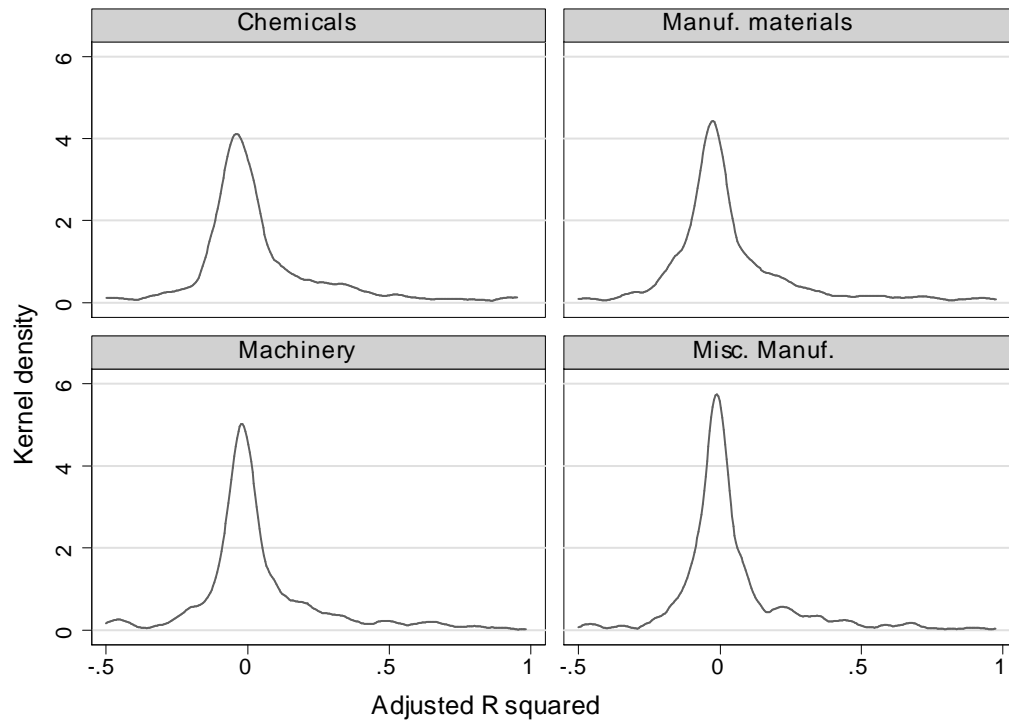


Figure 9: Distributions of \bar{R}^2 for product level regressions of firm-specific fixed levels on Value Added per worker

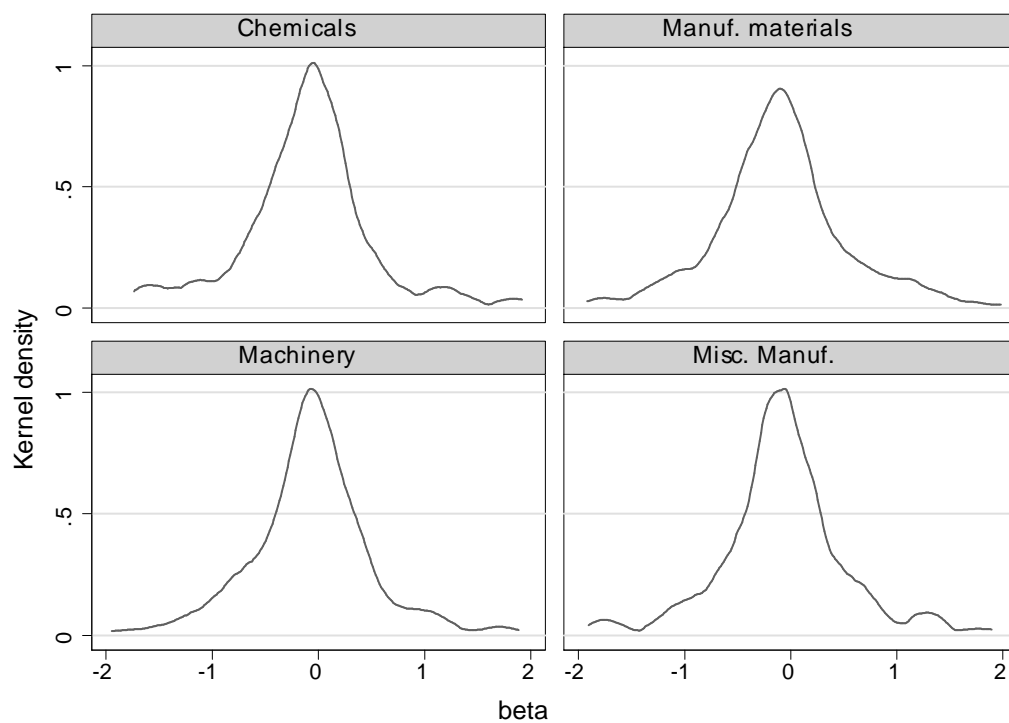


Figure 10: β_n estimates at the 8-digit product level, with truncation correction.




Figure 11: β_n estimates at the 8-digit product level, without truncation correction




Figure 12: The distribution of the standard deviation of changes in destination specific log revenues for Danish firms exporting to that destination for four consecutive years (1999-2002)

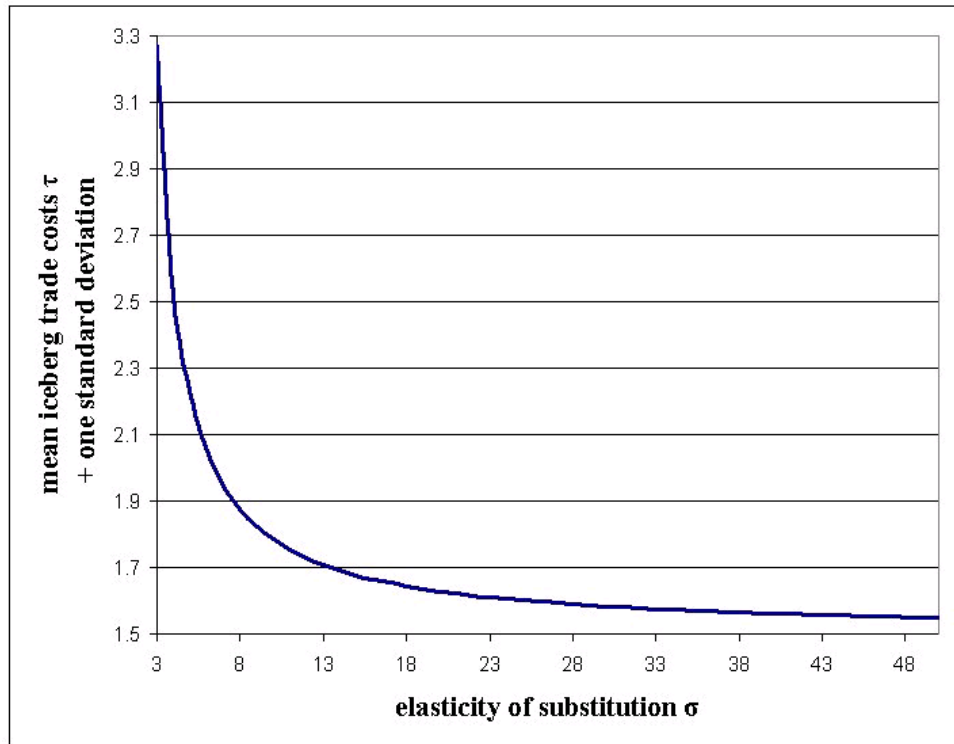


Figure 13: The τ faced by a firm with trade costs one standard deviation greater than the mean $\tau = 1.5$, given varying σ .

References

- [1] Akerberg, D., Benkard, C. D., Berry, S., Pakes, A., 2007, Econometric Tools for Analyzing Market Outcomes, in J. Heckman & E. Leamer, eds, 'Handbook of Econometrics Vol. 6A', North Holland, Amsterdam.
- [2] Alvarez, R., Faruq, H., 2007. Lopez, R. New Products in Export Markets: Learning from Experience and Learning from Others. Mimeo.
- [3] Anderson, James E., van Wincoop, E., 2003. Gravity with Gravitas: A Solution to the Border Puzzle. *American Economic Review* 93, 170–192.
- [4] Anderson, James E., van Wincoop, E., 2004. "Trade Costs," *Journal of Economic Literature*, American Economic Association, vol. 42(3), pages 691-751, September.
- [5] Aw, B.E., Chung, S., Roberts, M.J., 2000. Productivity and Turnover Patterns in the Export Market: Firm-Level Evidence from Taiwan and South Korea. *World Bank Economic Review* 14, 65-90.
- [6] Bernard, A., Eaton, J. Jensen, J.B., Kortum, S., 2003. Plants and Productivity in International Trade. *American Economic Review*, 93, 1268-1290.
- [7] Bernard, A., Jensen, J.B., 1995. Exporters, Jobs and Wages in US Manufacturing: 1976-1987. *Brookings Papers on Economic Activity, Microeconomics*, 67-119.
- [8] Bernard, A., Jensen, J.B., 1999. Exceptional Exporter Performance: Cause, Effect, or Both?. *Journal of International Economics* 47, 1-25.
- [9] Bernard, B., Jensen, J.B., Redding, S., Schott, P., 2007. Firms in International Trade. *Journal of Economic Perspectives* 21, 105-130.
- [10] Bernard, B., Jensen, J.B., Schott, P., 2006. Trade Costs, Firms, and Productivity. *Journal of Monetary Economics* 53, 917-937.
- [11] Cabral, L., Mata, J., 2003. On the Evolution of the Firm Size Distribution: Facts and Theory. *American Economic Review* 93, 1075-1090.
- [12] Clerides, S., Lach, S., Tybout, J., 1998. Is Learning by Exporting Important? Microdynamic Evidence from Columbia, Mexico, and Morocco. *Quarterly Journal of Economics* 113, 903-48.
- [13] Cramer, J.S., 1987. Mean and Variance of R^2 in small and moderate samples. *Journal of Econometrics* 35, 253-266.

- [14] Damijan, Kostevc, Polanec, 2007. Exporters' Strategies: Geographic and Product Diversification Patterns. Mimeo University of Ljubljana.
- [15] Eaton, J., Eslava, Kugler, Tybout, 2007. Export Dynamics in Colombia: Firm-Level Evidence. NBER Working Paper No. 13531.
- [16] Eaton, J. and Kortum, S, 2002. Technology, Geography, and Trade. *Econometrica* 70, 1741-1779.
- [17] Eaton, J., Kortum, S., Kramarz, F., 2004. Dissecting Trade: Firms, Industries, and Export Destinations. *American Economic Review, Papers and Proceedings*, 94, 150-154.
- [18] Eaton, J., Kortum, S., Kramarz, F., 2007. An Anatomy of International Trade: Evidence from French Firms. Mimeo, CREST.
- [19] Greene, W., 2003. *Econometric Analysis*. Fifth Edition. MacMillan.
- [20] Helpman, E., Melitz, M., Rubinstein, Y., 2008. Estimating Trade Flows: Trading Partners and Trading Volumes. *Quarterly Journal of Economics* 123, 441-487.
- [21] Hummels, D., 2001. Toward a Geography of Trade Costs, Mimeo. Purdue University.
- [22] Hummels, D. and A. Skiba, 2003 Shipping the Good Apples Out? An Empirical confirmation of the Alchian-Allen Conjecture *Journal of Political Economy* 112 1384-1402
- [23] Melitz, M. J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica* 71, 1695-1725.
- [24] Nguyen, D., 2008. Demand Uncertainty and Trade Failures. Mimeo.
- [25] Olley, S., Pakes, A., 1996. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica* 64, 1263-1298.
- [26] Pavcnik, N., 1999. Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. Mimeo, Department of Economics, Dartmouth College.
- [27] Statistics Denmark (2003), Danish External Trade Statistics, Copenhagen, available at <http://www.dst.dk/Vejviser/Portal/Udenrigshandel/METODE/omudenrigshandelsstatistik.aspx>