Decomposing Firm-level Sales Variation*

Jakob R. Munch and Daniel X. Nguyen†
University of Copenhagen

First version: August, 2008
This version: October 2010

Abstract

Recently, much of the trade literature has been focused on using firm-specific productivity to explain export heterogeneity. This study provides evidence for the importance of incorporating firm-destination-specific effects such as demand shocks in theories of exporter heterogeneity. Our study estimates the proportion of firm-level sales variation within a product-destination market that can be explained by firm-specific effects such as productivity. We use a highly detailed dataset comprising firm–product-destination-specific exports and correct for truncation as modeled by recent trade theories. We find the contribution of firm-specific effects to be lower than previous estimates. The contribution varies greatly across products but is less than 31% for the majority of products. Within-destination sales variation is primarily explained by firm-destination-specific heterogeneity.

Keywords: Firm heterogeneity, firm-level export data, truncation correction.

JEL Codes: F12, C24

---

*We would like to thank participants at Purdue University, the University of Copenhagen, the Fourth Danish International Economics Workshop, Aarhus, the 2008 ETSG conference, Warsaw, the Fall 2008 Midwest Meetings International Economics Group Meetings, Ohio State University, 2008 EIIT, University of Colorado, the LMDG 2008 conference on 'Labor Economics and International Trade', Sandbjerg and NOITS 2009, Copenhagen for helpful suggestions and comments. In particular, we would like to acknowledge David Hummels, Chong Xiang, and Mette Ejrnæs for suggestions that have greatly improved the direction and exposition of this paper. Financial support from the Tuborg Foundation is gratefully acknowledged. All errors are our own.

†Address: Department of Economics, University of Copenhagen, Øster Farimagsgade 5, bygning 26, 1353 København K , Denmark, E-mails: jakob.roland.munch@econ.ku.dk, daniel.x.nguyen@econ.ku.dk.

1 Introduction

There is substantial variation in firm-level export sales. Recent theoretical works have attributed this sales variation to heterogeneity in firm productivities.[1] These theories were motivated by earlier empirical studies that identified differences between firms that do export and firms that do not[2]: on average, exporters produce more, hire more labor, pay higher wages, and exhibit higher productivities as measured by either total factor productivity or value added per worker. The contrasts between exporters and nonexporters supported the story that productivity and exporting status were linked.

This paper decomposes firm-level sales variation within a product-destination market. We estimate the proportions of sales variation that can be explained by firm specific effects versus firm-destination specific effects. The firm specific component comprises all firm characteristics that would affect firm sales, such as productivity, quality, or economies of scale. Since productivity is only one of many firm specific characteristics, the sales variation explained by the firm specific effects is an upper bound of that explained by productivity. The firm-destination specific component can be thought of as shocks that affect the firm differently in different markets. These could be demand shocks or firm-destination specific cost shocks.[3]

Our empirical approach is in part inspired by the product-level sales decomposition of Hummels and Klenow (2005). That study decomposes a country's export sales into the number of products and the sales per product. Hummels and Klenow (2005) show how much of a country's export sales can be explained by models concentrating on intensive trade margins (Armington 1969) rather than models concentrating on extensive trade margins (Krugman 1979). We also employ a decomposition method to discriminate between trade theories. We show how much within-destination sales variation can be

[1] Notable examples include Eaton and Kortum (2002), Bernard, Eaton, Jensen, and Kortum (2003), Melitz (2003), Bernard, Jensen, and Schott (2006), and Bernard, Jensen, Redding, and Schott (2007).

[2] Notable works include Aw, Chung and Roberts (2000), Bernard and Jensen (1995, 1999) and Clerides, Lach and Tybout (1998).

[3] See Das, Roberts and Tybout (2007), Eaton, Kortum and Kramarz (2010) or Nguyen (2010) for examples.

explained by a model concentrating on firm specific effects such as Melitz (2003) rather than firm-destination specific effects.

Our paper adds to a small but growing literature examining the destinations to which firms export. The lack of work in the area is due primarily to the dearth of firm-destination specific export observations. 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 supply domestically for several years before exporting, that they usually begin exporting to one destination country, and that many stop exporting activities soon after they begin[4]. Since productivity is realized before supply to any destination and applies to all destinations, these studies present empirical patterns unreconciled by the productivity heterogeneity models.

The current study adds to and extends this literature by utilizing a highly disaggregated and detailed dataset – we observe destination specific shipment values for the universe of Danish exporters in 2001 to 2003. This disaggregation level allows us to identify the firm specific component and a firm-destination specific component of an export. We can estimate the contribution of each using a structural model as basis. We show that a simple extension of the standard productivity heterogeneity model by Melitz (2003) with destination specific effects yields a decomposition of firm-level sales that may be taken to the data.

Three contemporary studies have goals related, but not identical, to our own. Eaton, Kortum and Kramarz (2010) structurally estimate the contribution of firm specific productivity to both the probability of entering a destination and the variance of sales conditional on entry. They find that the variance of firm specific effects can account for 57% of the variation of entry into a destination and 39% of the variation of sales in a destination conditional upon entry into that destination. Kee and Krishna (2008) examine Bangladeshi exports of textiles to the US and EU. They find that a textile firm's market

[4]For example, Eaton, Eslava, Kugler and Tybout (2007), Damijan, Kostevc and Polanec (2007), Alvarez, Faruq and Lopez (2007).

share in EU cannot predict its market share in the US: the correlation between the two is not statistically different from zero. Lawless and Whelan (2008) use firm-destination data from a survey of 676 Irish-owned exporters to explain to where and how much firms export. They use OLS estimates to find that firm-year specific effects account for 41% of firm-destination-year sales variation.

In contrast to Kee and Krishna (2008) and Lawless and Whelan (2008) our empirical approach is derived directly from the broad class of CES trade models with firm heterogeneity. On the other hand our approach is much more empirical than Eaton, Kortum and Kramarz (2010) in the sense that we only use the theoretical framework to show how firm-level sales may be decomposed into its specific components. Other than that we do not impose much structure on the estimations. Another distinguishing feature of our empirical approach is that we are careful to account for truncation. Only the most productive firms will export and this must be accounted for to avoid biased estimates. Finally, our study uses the most detailed dataset of the related studies. The data cover the universe of Danish firms and uniquely identifies exports by destinations at the eight digit product level. This level of disaggregation at the product level is not available in the three other studies. It turns out to be very important as we document substantial cross product variation in the contribution of firm specific productivity to sales variation.[5]

In our main results, we estimate the contribution of firm specific heterogeneity to overall 2003 Danish export sales variance by HS6 product category. For half of Danish exported products, the contribution is lower than 31%. The mean firm specific contribution across our sample is only 33%, while firm-destination specific effects contribute the remaining 67%. In summary, firm specific effects play a nontrivial role, but firm-destination specific effects matter a great deal more in explaining firm-level sales variation. As robustness checks, we look at different product aggregation levels and different years. We also remove small trade flows and new trade flows. Our results consistently point towards firm-destination specific effects as the main driver of sales variation. This suggests that

[5]In Section 7 we offer a more thorough discussion of how our approach and results relate to the literature.

it is important to incorporate and work out the implications of firm-destination specific effects such as demand shocks in theories of exporter heterogeneity.

In the next section, we present an illustrative example of how firm specific effects may or may not drive firm sales across destinations. Next, we show how firm-level sales may be decomposed and how truncation biases standard estimation procedures. Section 4 outlines our strategy to overcome this bias. Section 5 descripes the Danish export data. Section 6 presents the estimation results. We conclude with a discussion of our results in contrast to previous estimates.

2 An illustrative example

To aid the reader in understanding the goal of this paper, we begin with an illustrative example. Suppose Denmark exports to only two destinations: Sweden and Germany. Models based based exclusively on productivity heterogeneity predict that firms that sell to both destinations should have relative revenues that are one-to-one correlated. If a Danish brick firm's sales to Germany are twice the average of all Danish brick firms selling to Germany, we can say that this firm is twice as productive as average. The same firm's Swedish sales should also be twice the average. The one-to-one correlation predicts that the variation in German relative revenues should completely explain the variation in Swedish relative revenues.

Figure 1 depicts firm (log) revenues to Sweden and Germany for building bricks and plastic boxes[6]. The revenues are relative to the mean Danish firm revenues of the respective product to the respective destination.

Insert Figure 1 here

The OLS results for boxes support a weaker interpretation of models based on pro-

[6] The two products are more precisely "Building bricks (excl. those of siliceous fossil meals or similar siliceous earths, and refractory bricks of heading 6902)" (CN8 product 69041000) and "Boxes, cases, crates and similar articles for the conveyance or packaging of goods, of plastics" (CN8 product 39231000). Sweden and Germany are the two most popular destinations for Danish exporters.

ductivity heterogeneity: that relative revenues are strongly and positively correlated, and close to one. The slope, although statistically different from one, is still high at 0.84. The variation in German relative revenues explains a little more than half of the variation in Swedish relative revenues.

In contrast, the OLS results for building bricks suggest a different story. The implied correlation is negative and not statistically different from zero. The $R^2 = 0.08$ shows that little of the Swedish variation is explained by the German variation.

Of course, this simple example ignores several estimation issues that will be carefully dealt with in the remainder of the paper. First, not all firms in the sample above export to both Sweden and Germany; this truncation may bias the results. Second, we will not rely on the estimated slope as a measure of the contribution of productivity to relative revenues. One reason is that the estimated slope could be negative, which is not interpretable. Another reason is that the correlation does not tell us the contribution of productivity variation to the total variation. For example, consider Figure 2 below, which presents three graphs of simulated relative revenues in two destinations. All three scatters have fitted slopes of 1, but different R^2 values. If productivity heterogeneity is the sole source of the variation, we should expect to see scatters similar to that in the upper left panel of Figure 2. As the contribution of firm-destination specific heterogeneity rises, the scatters begin to look more like the ones in the upper right and bottom left of Figure 2. Therefore, estimated slopes close to unity are misleading confirmations of models based on firm-level productivity differences. Instead we will focus on a statistic similar to R^2 as our measure of the contribution of firm productivity to sales variation.

Insert Figure 2 here

3 Theory and sales decomposition

Consider a small country exporting N products to foreign country destinations $j \in J$. For each product $n \in N$, there are M_n firms each producing a unique variety ω. Not all firms export to all destinations; only $m_{nj} < M_n$ firms supply to destination j. For the rest of this section, we focus our attention on a single product and therefore drop the subscript n without loss of generality. Each firm ω has a potential destination specific sales $r^*_{\omega j}$. We show in Appendix A that in a simple extension of the Melitz (2003) model the potential sales $r^*_{\omega j}$ can be decomposed into a destination specific component a_j, a firm specific component b_ω, and a firm-destination specific component $x_{\omega j}$:

$$\ln r^*_{\omega j} = a_j + b_\omega + x_{\omega j} \tag{1}$$

The destination specific component a_j captures all the country characteristics that make all firms have higher sales, such as the usual gravity terms. For instance, all Danish brick firms potentially have higher sales in the Germany than in Croatia because Germany contains more people, has higher per capita income, and is closer to Denmark.

The firm specific effect b_ω represent those firm characteristics that causes one firm to have higher sales than another. These include lower costs (Melitz (2003)), higher quality (Baldwin and Harrigan (2009), Kugler and Verhoogen (2010)), better capability (Johnson 2007), or any other firm specific characteristic discussed in the recent firm-heterogeneity trade literature. In the literature, firm specific effects are drawn from exogenous and independent distributions. In this model, b_ω is drawn from a normal distribution with product specific mean $\bar{b}$ and variance s_b^2.

Finally, the firm-destination specific effect $x_{\omega j}$ captures sales variation that cannot be explained by gravity-style country specific effects or firm-heterogeneity-type firm specific effects. It captures many types of firm-destination specific characteristics, such as perceived quality (Nguyen 2010) or other demand shocks (Das, Roberts, and Tybout 2006), or firm-destination specific trade costs (Eaton, Kortum and Kramarz (2010)). The firm-destination specific effects are drawn from exogenous independent normal distributions with product specific mean $\bar{x}_j$ variance s_x^2.

These normality assumptions are supported by the distribution of domestic revenues of Danish firms presented in Figure 3 and is consistent with previous studies of firm size distributions (Cabral and Mata (2003)) and export selection[7].

Insert Figure 3

The productivity draw and taste draws are constructed to be uncorrelated with one another. Given this assumption and equation (1), the variance of the within-destination sales distribution can be expressed as the sum of the variances of the firm specific effect and the firm-destination specific effect, $s_b^2 + s_x^2$. Our goal is to estimate the contribution of the firm specific effect to the variance of potential sales for firms within a destination, controlling for destination specific effects. That is, we estimate the statistic

$$Q^2 = \frac{s_b^2}{s_b^2 + s_x^2}, \tag{2}$$

for each Danish product exported in 2003.

The statistic Q^2 is similar to $\bar{R}^2$, the adjusted coefficient of determination, if b and x are not correlated. If they were correlated, then $\bar{R}^2$ attributes all of the correlation to the firm specific productivity draw. This would bias the estimation of firm specific effects upwards, see Appendix B for a formal proof.

3.1 Truncation issues

If sales were observed for every firm-destination pair, a simple ANOVA of $\ln r^*_{\omega j}$ on destination and firm specific effects would consistently decompose the variance into that explained by a_j and that explained by b_ω, with the residual being attributed to $x_{\omega j}$. However, our dataset is an unbalanced panel where not every firm sells to every destination. The firm-destination sales $r^*_{\omega j}$ is truncated, with the truncation endogenously correlated

[7] A growing theoretical literature approximates firm productivity with the Pareto distribution. This is to some extent driven by the analytical tractability of this distribution, see e.g. Chaney (2008), Helpman, Melitz and Yeaple (2004) and Eaton, Kortum and Kramarz (2010) as well as some empirical literature (Axtell 2001). The latter two use log-normal error terms to better fit the data. In these studies, the sales distributions deviate from Pareto in a way that is indicative of a truncated log normal distribution: first, the curvature is concave, not convex like Pareto, and second, the mass of firms with very low sales is too large to fit the Pareto distribution.

with a_j, b_ω, and $x_{\omega j}$. In this section, we show how we correct for these sources of bias.

Melitz (2003) suggests that the presence of a firm in a destination is tied to its potential profit in that market. Following Krugman (1980) and Melitz (2003), we assume firm ω's profits π gained from supplying to j are

$$\pi_{\omega j} = \frac{1}{\sigma} r^*_{\omega j} - f_j. \tag{3}$$

where $\frac{1}{\sigma} > 0$ is the portion of firm ω's sales over its variable cost of supply and f is the fixed cost of supply to destination j. Profits are positive when $r^*_{\omega j} > c_j$, where $c_j = \sigma f_j$ and is unknown to the econometrician. Therefore, we cannot observe all $r^*_{\omega j}$. We only observe $r_{\omega j}$, where

$$r_{\omega j} = \begin{cases} r^*_{\omega j} & \text{for } r^*_{\omega j} \geq c_j \\ 0 & \text{for } r^*_{\omega j} < c_j. \end{cases} \tag{4}$$

Equation (4) given (1) is the standard Type 1 Tobit Model with latent effects described in Honoré and Kyriazidou (2000). In an earlier work, Honoré (1992) shows that if the latent effects[8] are correlated with the probability of truncation, then the Heckman (1979) two-step procedure is biased. Honoré's solution to this problem treats the latent effects as nuisance variables and differences them out. This method renders the effects immeasurable. In our study, a_j is a nuisance variable, but b_ω is a parameter of interest, so we cannot use Honoré's approach. Instead, we treat a_j as a fixed effect, b_ω as a random effect, and $x_{\omega j}$ as a residual. We then estimate s_b^2 and s_x^2 using a Monte Carlo Expectation-Maximization Maximum Likelihood Estimation (MCEM) proposed by Walker (1996) and used widely in biometrics research. Kuhn and Lavielle (2005) show under very general conditions that the MCEM procedure obtains consistent estimates for nonlinear mixed-effects models. Using simulated datasets that resemble our actual dataset, we verify that our MCEM procedure estimates s_b^2 and s_x^2 consistently.

[8] In our case the latent effects correspond to a_j and b_ω.

4 Estimation strategy

This paper estimates via MCEM the portion of sales variance contributed by firm specific effects as measured by our statistic in equation (2). Using our model given by equation (4) given (1), we can characterize the distribution of $r_{\omega j}$ given a_j, b_ω, c_j, and s_x^2. The density of $\ln r_{\omega j}$ conditional on observing positive revenues is

$$\begin{aligned}\Pr\left(\ln r_{\omega j} = r | a_j, b_\omega, c_j, s_x^2, I = 1\right) &= \Pr\left(a_j + b_\omega + x_{\omega j} = r\right) \\ &= \frac{1}{s_x}\varphi\left(\frac{r - a_j - b_\omega}{s_x}\right), \end{aligned} \quad (5)$$

where $\varphi(\cdot)$ is the standard normal pdf, and where I denotes the indicator function that takes the value 0 if observed sales $r_{\omega j} = 0$ and 1 if $r_{\omega j} > 0$. The probability of a truncated observation is given by

$$\begin{aligned}\Pr\left(r_{\omega j} = 0 | a_j, b_\omega, c_j, s_x^2, I = 0\right) &= \Pr\left(a_j + b_\omega + x_{\omega j} < c_j\right) \\ &= \Phi\left(\frac{c_j - a_j - b_\omega}{s_x}\right), \end{aligned} \quad (6)$$

where $\Phi(\cdot)$ is the standard normal cdf. Combining (5) and (6) with the indicator function I, we derive the unconditional probability of $\ln r_{\omega j} = r$:

$$\Pr\left(\ln r_{\omega j} = r | a_j, b_\omega, c, s_x^2\right) = \frac{1}{s_x}\varphi\left(\frac{r - a_j - b_\omega}{s_x}\right) I + \Phi\left(\frac{c_j - a_j - b_\omega}{s_x}\right)(1 - I). \quad (7)$$

Following Wooldridge (2002), we find L_ω, the joint density of $\vec{r}_\omega = (r_{\omega 1}, r_{\omega 2}, ..., r_{\omega J})$ given b_ω, c, s_x^2 and the vector $\vec{a}_j = (a_1, a_2, ..., a_J)$:

$$L_\omega\left(\vec{r}_\omega | \vec{a}_j, c_j, b_\omega, s_x^2\right) = \prod_{j=1}^{J}\left(\frac{1}{s_x}\varphi\left(\frac{r_{\omega j} - a_j - b_\omega}{s_x}\right) I + \Phi\left(\frac{c_j - a_j - b_\omega}{s_x}\right)(1 - I)\right). \quad (8)$$

We treat b_ω as a random effect drawn from a normal distribution with mean zero and variance s_b^2. Given s_b^2, we can integrate out L_ω's dependence on b_ω. Finally, we sum this integral over all m firms to arrive at our log-likelihood l of the unknown parameters given observations $\vec{r} = \{r_{\omega j}|\omega = 1, ..., W; j = 1, ..., J\}$:

$$l\left(\vec{a}_j, c_j, s_x^2, s_b^2|\vec{r}\right) = \sum_{\omega=1}^{W} \ln\left(\int_{-\infty}^{\infty} L_\omega\left(\vec{r}_\omega|\vec{a}_j, b_\omega, c_j, s_x^2\right) \frac{1}{s_b}\varphi\left(\frac{b}{s_b}\right) db\right). \tag{9}$$

We treat the destination specific effect a_j as a fixed effect. We remain ambivalent as to the underlying distribution of a_j, since country specific effects are not of interest. We obtain our estimates via MCEM. In each iteration, the b_ω's are integrated out using a forty-point Gaussian Quadrature (the E-step). The parameters are then obtained by maximizing l (the M-step) using the `MAXLIK` procedure in Gauss. The steps are repeated until the squared sum of the gradient of estimated coefficents was less than $1e-5$.

The sample space of $\vec{r}_\omega$ is a function of the unknown parameter c_j. Zuehkle (2003) suggests estimating c_j with the minimum order statistic of the untruncated $r_{\omega j}$:

$$\hat{c}_j = \min\{r_{\omega j}|r_{\omega j} > 0; \omega = 1, ..., W; j = 1, ..., J\}. \tag{10}$$

Carson and Sun (2007) proves that $\hat{c}_j$ converges to c_j at the rate of $1/W$. They also show that MCEM estimates of the remaining coefficients are asymptotically normal with asymptotic variances identical to the case when c_j is known. We follow their lead and use $c_j = \hat{c}_j$. We then estimate the other parameters in (9) via MCEM, as previously described.

4.1 Monte Carlo simulation

We verify our procedure's ability to accurately estimate Q^2 under various conditions. We simulate 90 datasets of $(W, J) = (100, 100)$ possible firms and destinations[9] and compare the estimated $\hat{Q}^2$ with the known true $Q^2 = \tilde{Q}^2$. The simulation procedure is outlined in the appendix.

The results of our simulations are summarized in Figure 4 below. Our estimated $\hat{Q}^2$

[9] We also try $(W, J) = (100, 50)$ wth similar results.

tracks the true value $\tilde{Q}^2$ well, with none of the median estimates more than one standard deviation from the true value. We use median values because there were a handful of outliers that resulted from the MCEM not converging.

Insert Figure 4 here

5 Danish firm-level data

The Danish External Trade Statistics provides product-level destination specific export data for the universe of Danish firms. Exports are recorded according to the eight-digit Combined Nomenclature (CN) product code which encompasses approximately 10,000 different product categories. While all trade flows with non-EU countries are recorded by customs authorities (and so the coverage rate in the data is higher than 95 percent), 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 predominantly small 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 threshold[10] 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 lower but still in the range 85-90 percent. See Statistics Denmark (2003) for further details.

This study examines Danish manufacturing exports in 2003, but for robustness checks we also use data from 2001 and 2002. We select all manufacturing firms with positive inputs of labor and capital and with positive export sales. Also, we consider only manufacturing products by selecting products in one-digit SITC categories 5, 6, 7 and 8. With these restrictions our 2003 dataset comprises 155,426 firm-destination-product sales observations by 4,304 firms in 5,339 eight-digit CN8 products to 223 destination countries,

[10] For the years considered, this threshold was DKK 2.5 million corresponding to approximately USD 500,000.

see Table 1. The aggregate value of all these trade flows totals 182 billion Danish kroner (DKK), which in 2003 roughly correspond to USD 28 billion.

Insert Table 1 here

Table 1 shows some similarities between exporters in Denmark and those in bigger economies. The median number of destinations for firm-product exports is 1, which is in line with the findings for the US (Bernard and Jensen, 1995) and France (Eaton, Kortum, Kramarz, 2004). Clearly, some firms ship their products to many destinations – the mean number of destinations is 3.4 and the maximum number is 138.

At the disaggregated eight-digit product level most destinations do not have many Danish firms present. The median number of firms is 1 and the mean is 2.2. At the slightly more aggregated six-digit Harmonized System (HS6) level the mean number of firms is 2.6, but still more than half of the product-destinations have only one firm. This presents a problem for our empirical strategy, because it cannot identify the destination specific effect with only a single firm. Therefore, in the following, we will only consider sufficiently important products by imposing some restrictions on the data. First, we disregard product-destinations with less than five firms and, second, products with less than 25 firm-destinations in total are deleted. Third, we also need cross destination variation to estimate the firm specific effects, so products that are shipped to fewer than three destinations (by any firm) are omitted. With these restrictions we end up with just 491 CN8 products or 480 HS6 products. However, they constitute more than a third or a half of the overall trade volume respectively, see Table 1.

For the restricted samples there is not much difference between the HS6 and CN8 levels. In the following we focus on the HS6 level as it covers the largest fraction of the total Danish export volume, but we report results for the CN8 level as well.

6 Estimation results

We use the MCEM procedure to obtain an estimate, Q^2_{MCEM}, for the contribution of firm specific effects for each Danish export in 2003. We do this at the HS6 product level. We also perform OLS dummy regressions of destination-mean-differenced observed revenues $\left(\ln r_{\omega j} - \sum_{\omega=1}^{W_j} \ln r_{\omega j}\right)$ on firm fixed-effects. From the OLS regressions, we retrieve the adjusted coefficient of determination $\bar{R}^2$ and estimate Q^2_{OLS} by[11]:

$$Q^2_{OLS} = \max\left\{0, \bar{R}^2\right\}. \tag{11}$$

Q^2 is defined as a positive number, so we treat negative $\bar{R}^2$ values as estimates of 0 for Q^2. In the following, we compare the MCEM estimates Q^2_{MCEM} to the OLS estimates Q^2_{OLS}. Our main results are derived at the detailed product level discussed in the next section, which is followed by a number of robustness checks.

6.1 Product level

As noted in section 3, there are many HS6 products that contain few firms selling to few destinations. We drop products containing fewer than 25 firm-destination observations, products that are exported to less than three destinations, and product-destinations categories containing fewer than five firms. A total 66,488 firm-destination-product observations remained, spanning 3,790 firms in 480 products to 84 countries and totalling DKK 91 billion. We estimate s_b^2, s_x^2, and consequently Q^2 for each of these 480 products.

Our estimation procedure resulted in mean and median values of 33% and 31% for Q^2_{MCEM} across the 480 HS6 products. This is considerably lower than OLS estimates, which resulted in mean and median values of 43% and 46% for Q^2_{OLS}. For comparison, Lawless and Whelan (2008) obtain an $\bar{R}^2$ of 41% across their sample of Irish exporters. It should be noted that the product level dimension of the data is important, as the estimated Q^2_{MCEM}'s exhibit substantial variation across products. Histograms for the

[11] We use the adjusted coefficient of determination to avoid small sample bias. Cramer (1987) shows that the unadjusted R^2 is heavily biased upwards for small samples.

MCEM and OLS estimates are presented in Figure 5 below:

Insert Figure 5 here

The histogram of the MCEM estimates in Figure 5 is systematically to the left of that of the OLS estimates. To understand why, we compare the difference between Q^2_{OLS} and Q^2_{MCEM} for each product. Figure 6 shows that Q^2_{OLS} generally overshoots Q^2_{MCEM}. The overshooting is exacerbated at low values of Q^2_{MCEM}. For products with Q^2_{MCEM} between 10% and 20%, Q^2_{OLS} averages 31%. For products with Q^2_{MCEM} between 20% and 30%, Q^2_{OLS} averages around 41%. This upwards bias shifts the histogram for Q^2_{OLS} to the right. Q^2_{OLS} actually undershoots Q^2_{MCEM} at values of Q^2_{MCEM} greater than 60%. Our simulations showed that Q^2_{MCEM} slightly undershoots the true Q^2 at high values, so Q^2_{OLS}'s downward bias is even worse. That is, Q^2_{MCEM} is a more accurate estimator for Q^2 than Q^2_{OLS} across the entire range.

Insert Figure 6 here

To sum up we have found that OLS estimates of the contribution of firm specific effects are generally biased. Our results show that the direction of bias is dependent on the degree to which firm specific effects affect sales variation. In HS6 products where firm specific effects do not contribute much to the overal sales variation, an OLS dummy regression overestimates the contribution of the firm specific effect. In products where firm specific effects play a large role, the OLS regression underestimates the true contribution.

Since our Q^2_{MCEM}'s are product specific, we investigate whether there are any patterns in the Q^2_{MCEM}'s across products. Our theoretical model is stylized and does not give us any predictions about how Q^2_{MCEM} varies with product characteristics, so a priori we do not have any expectations about any relationships. However, we regressed our estimates are several product-level characteristics to investigate any possible relationship.

We find no relationship between the contribution of firm specific effects and a number of product specific characteristics. We regressed Q^2_{MCEM} on the mean and variance of the capital labor ratio within the HS6 product code, the mean and variance of the value added per worker for firms within the HS6 product code, and the mean and variance of the total HS6 output. We found no significant correlation.

We also did not find any correlation between Q^2_{MCEM} and previously estimated measures of product differentiation. We regressed the Q^2_{MCEM}'s on import demand elasticities for the U.S. estimated by Broda and Weinstein (2006) and import demand elasticities for Denmark estimated by Broda, Greenfield and Weinstein (2006). Again we did not find any correlation. Finally, we partitioned our results by the Rauch (1999) classification of product differentiation. Approximately 90% of our products are classified as 'differentiated' while most of the remaining 10% are classified as 'reference priced.'[12]. The 'differentiated' products had a median Q^2_{MCEM} of 33% while the 'reference priced' products had a median Q^2_{MCEM} of 20%. However, there were less than 50 estimated 'reference priced' products, so we refrain from speculating about any true differences.

6.2 Measurement error

Firm-destination specific effects contributes over two-thirds of the sales variation in a product-destination market for over half of Danish HS6 exports. Our theory suggests that this variation is due to firm–destination specific demand variation. However, if the export sales data are riddled by measurement error, then that error could be a possible source of variation that reduces the relative contribution from firm specific effects.

As a first robustness check we have calculated Q^2_{MCEM} for a similar sample but where small trade flows are excluded by deleting observations with a value less than DKK 1,000. This is to ensure economically unimportant and perhaps noisy observations do not affect our results. Those results are similar to the results presented above, with mean/median of 34%/33% for the MCEM procedure and 41%/38% for the OLS procedure.

Second, suppose that measurement error is the sole cause of the firm-destination spe-

[12] Rauch (1999) also classifies products according to whether they are traded on organized exchanges, but we had only a handful of products of this type in our sample. This is because our dataset contains only manufacturing products.

cific variation. Our sample has a log sales mean of 10.6 and sample log sales variance of 8.4. If measurement error is the cause of two-thirds of that variance, that would imply that an average Danish export recorded at a value of DKK 40,000 has a 68% (1 standard deviation) confidence interval of DKK 4,000 to 420,000. Our data is customs trade data from which tariff revenues are calculated, and it does not seem plausible to have measurement errors that large.

6.3 Aggregation

We also estimate Q^2_{OLS} and Q^2_{MCEM} at the CN8 product level, the most disaggregated level available to us. The results are similar to estimates performed at the HS6 level. We obtain a mean and median of 31% for Q^2_{MCEM}, and 46% and 44% for Q^2_{OLS}. The histograms at the CN8 level are presented in Figure 7 below:

Insert Figure 7

The histograms in Figure 7 resembles those in Figure 5; OLS estimates are systematically higher than MCEM estimates.

As Table 1 shows, our data restrictions reduce the sample size to about a third of the trade volume at the eight digit level. This reduction did not result in a substantial gain in the number of products: only 491 CN8 products passed the estimation restrictions, compared to 480 HS6 products. By disaggregating to CN8, we threw away observations without gaining much in return. We use the HS6 product classification for our robustness checks below, as that sample comprises a higher total export volume.

As briefly mentioned in the introduction, our dataset contains product code information that are not contained in Eaton, Kortum and Kramarz (2010) or Lawless and Whelan (2008). To better compare our results to theirs, we aggregate our data to broader industries. We estimate Q^2_{MCEM} and Q^2_{OLS} at the HS2 industry level.[13] As before, we restrict

[13]To compare more directly with existing studies we should estimate one Q^2_{MCEM} and one Q^2_{OLS} for firm level sales without any distinction between different products. However, that proved infeasible as Gauss was unable to handle the size of the dataset.

our analysis to industries containing at least 25 firm-destination observations, industries that are exported to at least three destinations, and industry-destinations categories containing at least five firms. With these restrictions, we have 77,411 observations spanning 4,276 firms in 54 industries exporting to 161 countries, totalling DKK 178 billion.

For the 54 HS2 industries, we obtain median estimates of 32% for Q^2_{MCEM} and 38% for Q^2_{OLS}. This result is in line with our previous estimates at the HS6 product level. For over half of Danish exporting industries, firm specific effects explain less than a third of total sales variation.

We obtain a mean of 43% for Q^2_{MCEM}, which is higher than the 35% mean obtained for Q^2_{OLS}. This is due to 16 industries having estimates of Q^2_{MCEM} greater than 80%. Figure 8 show this case.

Insert Figure 8 here

There was no obvious pattern why these industries exhibited higher contributions of firm specific effects. These results suggests that, if anything, the estimated contribution of firm specific effects rises with the level of aggregation.

6.4 Consistency over time

To see if our results are consistent over time, we repeat the exercise for the year 2001, with similar estimates for Q^2. For the 401 HS6 products that fit our restrictions in 2001, we obtain median estimates of $Q^2_{MCEM} = 34\%$ versus $Q^2_{OLS} = 43\%$.

The Q^2 estimates are not only correlated in the aggregate, but at the individual product level. There were 350 HS6 products that passed our estimation restrictions in both 2001 and 2003. We regressed $Q^2_{MCEM}$for 2003 on that for 2001 for these 350 products. Our estimated marginal effect was 0.76 with a standard error of 0.03. That is, a 10% increase in $Q^2_{MCEM,2001}$ corresponded to a 7.6% increase in $Q^2_{MCEM,2003}$.

The correlation is almost one-to-one when we restrict our regression constant to zero. That regression results in an estimated marginal effect of 0.92 with a standard error of 0.02. Figure 9 presents the point estimates for the two years:

Insert Figure 9 here

The strong correlation between $Q^2_{MCEM,2003}$ and $Q^2_{MCEM,2001}$ contrasts with the lack of correlation between OLS estimates $Q^2_{OLS,2003}$ and $Q^2_{OLS,2001}$. A regression of the two OLS estimates resulted in no significant correlation between the two. Figure 10 shows this lack of consistency across years:

Insert Figure 10 here

This exercise gives further evidence that our procedure accurately identifies the contribution of firm specific effects, while OLS estimates do not.

6.5 Established exports

Nguyen (2010) suggests that much of the export sales variation is due to firms testing destinations in order to determine whether they can be successful exporting to that destination. Therefore, firm–destination specific effects should play a larger role in the first year of exporting. To test that, we restrict our sample to only those firm-product-destination observations in 2003 that were also positive in 2002. That is, only 2003 exports by those firms that exported the same product to the same destination in both 2002 and 2003 were considered. This restriction leaves us with 31,242 observations spanning 303 HS6 products, 2,491 firms, and 64 countries and totalling DKK 69 billion.

The predictions from Nguyen (2010) are supported by the data. For the 303 established exports in 2003, we obtain mean and median values of 39% and 40% for Q^2_{MCEM} and 49% and 50% for Q^2_{OLS}. These values are $20-25\%$ higher than those estimates estimated for the sample which included first time exports. Therefore, firm specific effects are more important for these established exports. Contrastly, firm-destination specific effects are more important for the first year of exporting than for established exports. The histogram of results for the established exports is displayed in Figure 11.

Insert Figure 11 here

6.6 Core products

Firms typically export multiple products, and for such firms the within-firm output distribution across products is known to be highly skewed with typically one core product accounting for a major part of firm sales, see e.g. Bernard, Redding, and Schott (2010). Until now we have treated each firm-product combination as independent units of observations, but within-firm correlation across products in export markets may arise if non-core products are more likely to be sold in destinations where fixed costs related to sales of the core product already have been incurred.

Therefore, as a robustness check, we repeat our exercise for only the core product of each firm. We define firm $\omega's$ core product as the HS6 category constituting the highest export sales for firm ω. We drop all other products exported by ω. With this and the forementioned restrictions, we are left with 6,686 observations spanning 73 HS6 products, 1,342 firms, and 61 countries, totalling DKK 3 billion.

The MCEM estimates for core products are similar to those for all products. We obtain a median $Q^2_{MCEM,CORE} = 40\%$ for the 73 HS6 categories comprising only core products. For these same 73 HS6 categories, we estimate a $Q^2_{MCEM,ALL} = 37\%$ when we include all products.

The OLS estimates, however, dropped significantly when we look only at core products. We obtain a median $Q^2_{OLS,CORE} = 25\%$ for the core products compared to $Q^2_{MCEM,ALL} = 43\%$ for the sample with all products. Figure 12 compares the point estimates of Q^2 using both data restrictions and estimation techniques:

Insert Figure 12 here

Q^2 estimates using just the core products can predict that using all products. A simple regression of $Q^2_{MCEM,ALL}$ on $Q^2_{MCEM,CORE}$ results in a positive and significant coefficient of 0.57 (standard error of 0.11). This estimate increases to 0.97 when we restrict the constant to zero. For Q^2_{OLS}, the same exercise results in a coefficient of 0.31 with a standard error of 0.08. We estimate a coefficient of 1.12 when we restrict the constant to zero.

This exercise shows that product scope does not adversely affect the measurement of the contribution of firm specific effects to sales within a product category. Our estimates for Q^2_{MCEM} using just core products are on par with our results using all products. Core-product estimates can track all-product estimates well, and the marginal relationship is not significantly different from unity when we restrict the constant to zero.

7 Discussion and conclusion

We use a highly detailed dataset for Danish exporters to estimate the contributions of firm specific and firm-destination specific effects to the variation of sales within a product-destination market. We find that the contribution of firm specific effects varies greatly across products, and that it explains less than 31 percent of the variation for over half of Danish HS6 products. Our results suggest that firm-destination specific heterogeneity, rather than a firm specific effect such as productivity, captures the majority of heterogeneity and is the primary driver of variation in a market.

Our results are ten to fifteen percentage points below Eaton, Kortum, and Kramarz (2010) and Lawless and Whelan (2008). There are several systematic differences between our approaches that could lead to our disparate results.

Our approach is different from Eaton, Kortum and Kramarz (2010) in several ways. The most significant difference is that our study is much more empirically focused. Our study decomposes sales variation and interprets the estimated heterogeneities in a very broad sense: we want to see what proportion of sales is firm specific and what is firm-destination specific. Eaton, Kortum and Kramarz (2010) introduce a model with stricter interpretations of the sources of heterogeneity and then calibrates their model to match the data. In other words, Eaton, Kortum and Kramarz (2010) interpret the results within the structure of their specific model, while we can interpret our results within a broader class of CES trade models with firm heterogeneity.

Eaton, Kortum and Kramarz (2010) separate their results into that concerning entry variation and that concerning sales variation conditional upon entry. They find that firm specific effects account for 57% of the variation in entry and 39% of the variation in sales. Estimating two different percentages is correct if the exporting decision is akin to a type II Tobit, or Heckman (1979) two-stage model. The theoretical firm heterogeneity trade literature (e.g. Melitz (2003) and Bernard et al. (2003)) do not separate the two, however. Current theory suggests that entry is driven by the unconditional sales via a type 1 Tobit model. This paper follows the theoretical literature and estimates the contribution of firm specific effects to the unconditional sales variation. The disparity between our results and theirs could arise from this difference: they measure the firm specific effect's contribution to conditional sales variation, while we measure its contribution to unconditional sales variation.

Finally, the disparity between Eaton, Kortum and Kramarz's (2010) and our results may depend on the different productivity distributions assumed.. They draw their firm specific effects from a Pareto distribution and their firm-destination specific demand and entry shocks from lognormal distributions. This is to some extent driven by the analytical tractability of this distribution (see e.g. Chaney (2008)). We use lognormal distributions for both our firm specific and firm-destination specific distributions. We use the lognormal distribution because the Pareto distribution does not completely characterize the distribution of sales. Even in studies arguing for the Pareto distribution (Axtell 2001), the sales distributions deviate from Pareto in a way that is indicative of a truncated log normal distribution: first, the curvature is concave, not convex like Pareto, and second, the mass of firms with very low sales is too large to fit the Pareto distribution. To account for these discrepancies, studies such as Helpman, Melitz, and Rubinstein (2008) and Eaton, Kortum, and Kramarz (2010) use log-normal error terms to better fit the data.

Our study improves on the methods pioneered by two other studies looking at the importance of firm specific effects. Kee and Krishna (2008) examine Bangladeshi exports of textiles to the US and EU. They find that a textile firm's market share in EU cannot predict its market share in the US: the correlation between the two is not statistically different from zero. Lawless and Whelan (2008) use firm-destination data from a survey of 676 Irish-owned exporters to explain to where and how much firms export. Using OLS regressions with fixed-effects, they find the variation in firm-year and country specific effects accounts for 57 percent of the total variation. By itself, the country specific effects explain 16 percent of the variation, leaving 41 percent of the variation explained by firm-year specific effects.[14] Our Q^2_{OLS} estimates resemble those of Lawless and Whelan's. In addition, we show that truncation issues bias OLS results and must be accounted for. Honoré and Kyriazadou (2000) discuss that the Heckman (1979) two-step procedure cannot correct for this truncation bias when entry into a destination is related to the firm-destination specific demand draws. Instead, this paper uses a monte carlo estimation maximization procedure to consistently account for truncation and the unobserved effects.

The Melitz (2003) model uses firm specific productivity to deftly explain variation between exporters and nonexporters. Much of the literature has built upon this idea, and it does indeed offer a tractable way of modelling firm heterogeneity. We would like to stress that our results should not be taken as a refutation of the Melitz (2003) model. After all, our results show that firm specific productivity plays an important role in explaining firm heterogeneity for many products. However, the majority of variation is firm-destination specific, which suggests a new direction based on firm-destination specific effects may better reconcile trade patterns.

We show that OLS estimates tend to overestimate low contributions and underestimate high contributions. Since the contribution of firm specific effects are low in most products, OLS regressions tend to overestimate in general. To consistently estimate firm specific effects, we employ a Monte Carlo Estimation-Maximization strategy used mainly in the Biometrics literature. We argue that this method can be employed fruitfully in studies of firm-level exporting with truncation issues.

[14] In other specifications they estimate the explanatory power of observed firm characteristics such as value added per employee and sector dummies instead of firm fixed effects.

References

[1] Alvarez, R., H. Faruq, and R. Lopez. 2007. "New Products in Export Markets: Learning from Experience and Learning from Others." Mimeo.

[2] Armington, Paul (1969), "A Theory of Demand for Products Distinguished by Place of Production." *IMF Staff Papers* 16, 159-176.

[3] Aw, B.E., S. Chung, and M. J. Roberts. 2000. "Productivity and Turnover Patterns in the Export Market: Firm-Level Evidence from Taiwan and South Korea." *World Bank Economic Review* 14: 65-90.

[4] Axtell, Robert, 2001. "Zipf distribution of U.S. Firm Sizes." *Science* 293: 1818 - 1820.

[5] Baldwin, R. and J. Harrigan (2009). "Zeros, Quality and Space: Trade Theory and Trade Evidence." *The American Economic Journal: Microeconomics*, forthcoming.

[6] Bernard, A., J. Eaton, J. B. Jensen, and S. Kortum. 2003. "Plants and Productivity in International Trade." *American Economic Review* 93: 1268-1290.

[7] Bernard, A., and J. B. Jensen. 1995. "Exporters, Jobs and Wages in US Manufacturing: 1976-1987." *Brookings Papers on Economic Activity, Microeconomics*, 67-119.

[8] Bernard, A., and J. B. Jensen. 1999. "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics* 47: 1-25.

[9] Bernard, A., J. B. Jensen, S. Redding, and P. Schott. 2007. "Firms in International Trade." *Journal of Economic Perspectives* 21: 105-130.

[10] Bernard, A., J. B. Jensen, and P. Schott. 2006. "Trade Costs, Firms, and Productivity." *Journal of Monetary Economics* 53: 917-937.

[11] Bernard, A., S. Redding, and P. Schott. 2010. "Multi-Product Firms and Product Switching." *American Economic Review* 100: 70-97.

[12] Broda, C., J. Greenfield and D. Weinstein. 2006. "From Groundnuts to Globalization: A Structural Estimate of Trade and Growth." NBER Working Paper No. 12512.

[13] Broda, C. and D. Weinstein. 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121: 541-585.

[14] Cabral, L., and J. Mata. 2003. "On the Evolution of the Firm Size Distribution: Facts and Theory." *American Economic Review* 93: 1075-1090.

[15] Carson, R. and Y. Sun. 2007. "The Tobit model with a non-zero threshold." *Econometrics Journal*. 10: 488-502.

[16] Chaney, T. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98: 1707-1721.

[17] Clerides, S., S. Lach, and J. Tybout. 1998. "Is Learning by Exporting Important? Microdynamic Evidence from Columbia, Mexico, and Morocco." *Quarterly Journal of Economics* 113: 903-48.

[18] Cramer, J. S. 1987. "Mean and Variance of R^2 in small and moderate samples." *Journal of Econometrics* 35: 253-266.

[19] Damijan, J. P, C. Kostevc, and S. Polanec. 2007. "Exporters' Strategies: Geographic and Product Diversification Patterns." Mimeo, University of Ljubljana.

[20] Das, M., M. Roberts, and J. Tybout. 2007. "Market Entry Costs, Producer Heterogeneity and Export Dynamics." *Econometrica* 75: 837-873.

[21] Eaton, J., M. Eslava, M. Kugler, and J. Tybout. 2007. "Export Dynamics in Colombia: Firm-Level Evidence." NBER Working Paper No. 13531.

[22] Eaton, J., and S. Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70: 1741-1779.

[23] Eaton, J., S. Kortum, and F. Kramarz. 2004. "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review, Papers and Proceedings* 94: 150-154.

[24] Eaton, J., S. Kortum, and F. Kramarz. 2010. "An Anatomy of International Trade: Evidence from French Firms." Mimeo, University of Chicago.

[25] Heckman, J. 1979 "Sample Selection Bias as a Specification Error," *Econometrica* 47: 152-161.

[26] Helpman, E., M. Melitz, and Y. Rubinstein. 2008. "Estimating Trade Flows: Trading Partners and Trading Volumes." *Quarterly Journal of Economics* 123: 441-487.

[27] Helpman, E., M. Melitz, and S. Yeaple. 2004. "Export versus FDI with Heterogeneous Firms." *American Economic Review* 94: 300–316.

[28] Honoré, B. 1992. "Trimmed LAD and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects." *Econometrica*, 60: 533-565.

[29] Honoré, B. and E. Kyriazidou 2000. "Estimation of Tobit-Type Models with Individual Specific Effects." *Econometric Reviews*, 19(3): 341-366.

[30] Hummels, D., and P. Klenow, 2005, "The Variety and Quality of a Nation's Exports" *American Economic Review* 95(3): 704-723 (June).

[31] Kee, H. L., and K. Krishna. 2008. "Firm-Level Heterogeneous Productivity and Demand Shocks: Evidence from Bangladesh." *American Economic Review, Papers and Proceedings* 98: 457-462.

[32] Krugman, Paul R. 1979. "Increasing Returns, Monopolistic Competition, and International Trade." *Journal of International Economics* 9: 469-479.

[33] Kugler, M. and E. Verhoogen. 2010. "Prices, Pland Size, and Product Quality." Mimeo, Columbia University.

[34] Kuhn, E. and M. Lavielle. 2005. "Maximum likelihood estimation in nonlinear mixed effects models." *Computational Statistics and Data Analysis.* 49: 1020-1038.

[35] Johnson, R. 2010. "Trade and Prices with Heterogeneous Firms." Mimeo, Dartmouth College.

[36] Lawless, M., and K. Whelan. 2008. "Where Do Firms Export, How Much, and Why?" Mimeo, University College Dublin.

[37] Luttmer, E. 2007. "Selection, Growth, and the Size Distribution of Firms." *Quarterly Journal of Economics* 122: 1103–1144.

[38] Melitz, M. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71: 1695-1725.

[39] Nguyen, D. 2010. "Demand Uncertainty and Trade Failures." Mimeo, University of Copenhagen. Revise and resubmit at *Journal of International Economics.*

[40] Rauch, J. 1999. "Networks versus markets in international trade." *Journal of International Economics* 48: 7-35.

[41] Statistics Denmark. 2003. *Danish External Trade Statistics.* Copenhagen, available at http://www.dst.dk/Vejviser/Portal/Udenrigshandel/METODE/omudenrigshandelsstatistik.aspx

[42] Walker, S. 1996. "An EM Algorithm for Nonlinear Random Effects Models." *Biometrics*, 52: 934-944

[43] Wooldridge, J. 2002. *Econometric Analysis of Cross Section and Panel Data.* Cambridge: The MIT Press.

[44] Zuehlke, T. 2003. "Estimation of a Tobit model with unknown censoring threshold." *Applied Economics.* 35: 1163-1169

A A model of sales variation

This appendix presents a simple extension of Melitz (2003) from which the firm level sales equation in section 3 may be derived. Consider a country exporting N products to foreign destinations $j \in J$. For each product $n \in N$, there are W_n firms each producing a unique variety ω. A portion W_{nj} of these firms supply to destination j. For the rest of this section, we focus our attention on a single product and therefore drop the subscript n without loss of generality. The utility gained in destination j from consuming varieties of this product is represented by u_j :

$$u_j = \sum_{\omega=1}^{W_j} \exp\left(\frac{x_{\omega j}}{\sigma}\right) (q_{\omega j})^{\frac{\sigma-1}{\sigma}}, \tag{12}$$

where $q_{\omega j}$ is consumption of variety ω in j and $\sigma > 1$ is a measure of the substitutability among the different varieties.

The utility function resembles a Dixit-Stiglitz utility function with a demand shifter. The demand shifter $x_{\omega j}$ represents destination j taste[15] for variety ω. Higher $x_{\omega j}$ corresponds to greater demand for that variety relative to other varieties in the same destination. Destination j's demand for variety ω can be derived as:

$$q_{\omega j} = (p_{\omega j})^{-\sigma} \exp(x_{\omega j}) \frac{Y_j}{P_j} \tag{13}$$

$$P_j = \sum_{\omega=1}^{W_j} \exp(x_{\omega j}) (p_{\omega j})^{1-\sigma}, \tag{14}$$

where $p_{\omega j}$ is the price of ω and Y_j is j's total expenditure on varieties. P_j is the corresponding Chamberlainian price index, which is unaffected by the actions of any single firm.

Firms share similar increasing returns to scale production technologies. Firm ω's cost

[15] Nguyen (2010) defines this parameter as "perceived quality". We can also think of it as $\omega's$ popularity or appeal in j.

$c_{\omega j}$ of supplying $q_{\omega j}$ units of output to destination j is

$$c_{\omega j}\left(q_{\omega j}\right) = f + \exp\left(\frac{b_{\omega}}{1-\sigma}\right)\tau_j q_{\omega j}, \tag{15}$$

where f and τ are fixed and variable costs identical to all firms supplying to j. The firm specific $\exp\left(\frac{b_\omega}{1-\sigma}\right)$ is the firm's marginal cost of product that is constant across all destinations. The b_ω term is a normalized measure of ω's productivity: a higher b translates to a lower marginal cost for the firm across all destinations.

Each firm $\omega \in \{1, ..., m_j\}$ draws its firm specific productivity b_ω. In addition, each firm draws a firm-destination specific taste parameter $x_{\omega j}$. The two random variables b_ω and $x_{\omega j}$ determine firm ω's potential sales $r^*_{\omega j}$ in destination j, which is presented in log form:

$$\ln r^*_{\omega j} = a_j + b_\omega + x_{\omega j} \tag{16a}$$

$$a_j = \ln\left(\frac{Y_j \tau_j^{1-\sigma}}{P_j}\right). \tag{16b}$$

The firm productivity draws, b_ω, are drawn an exogenous distribution, as in Melitz (2003). The firm's costs $c_{\omega j}$ can be subtracted from it sales $r^*_{\omega j}$ to generate its potential profit $\pi_{\omega j}$ gained from supplying ω to j:

$$\pi_{\omega j} = \frac{1}{\sigma} r^*_{\omega j} - f. \tag{17}$$

Firms will only supply to profitable destinations. Therefore, observed sales $r^*_{\omega j}$ can be characterized by:

$$r_{\omega j} = \begin{cases} r^*_{\omega j} & \text{for } r^*_{\omega j} \geq \sigma f \\ 0 & \text{for } r^*_{\omega j} < \sigma f. \end{cases} \tag{18}$$

B Bias of firm-destination effects estimation

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

$$y_m = \beta x_m + \varepsilon_m \tag{19}$$

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

$$\hat{\varepsilon}_m = y_m - \hat{\beta} x_m \tag{20}$$

We decompose the sum of the squared residuals:

$$\hat{\varepsilon}'\hat{\varepsilon} = \sum_{m=1}^{M} \hat{\varepsilon}_m^2 \tag{21}$$

$$= \sum_{m=1}^{M} \left(\beta x_m + \varepsilon_m - \hat{\beta} x_m\right)^2 \tag{22}$$

$$= \sum_{m=1}^{M} \varepsilon_m^2 + \left(\beta - \hat{\beta}\right)^2 \sum_{m=1}^{M} x_m^2 + 2\left(\beta - \hat{\beta}\right) \sum_{m=1}^{M} \varepsilon_m x_m \tag{23}$$

By the construction of ordinary least squares,

$$\beta - \hat{\beta} = -\frac{\sum_{m=1}^{M} \varepsilon_m x_m}{\sum_{m=1}^{M} x_m^2} \tag{24}$$

We can insert the difference between β and $\hat{\beta}$ from equation (24) into equation (23):

$$\hat{\varepsilon}'\hat{\varepsilon} = \sum_{m=1}^{M} \varepsilon_m^2 + \left(\frac{\sum_{m=1}^{M} \varepsilon_m x_m}{\sum_{m=1}^{M} x_m^2}\right)^2 \sum_{m=1}^{M} x_m^2 - 2\left(\frac{\sum_{m=1}^{M} \varepsilon_m x_m}{\sum_{m=1}^{M} x_m^2}\right) \sum_{m=1}^{M} \varepsilon_m x_m$$

$$\hat{\varepsilon}'\hat{\varepsilon} = \varepsilon'\varepsilon - \frac{\left(\sum_{m=1}^{M} \varepsilon_m x_m\right)^2}{\sum_{m=1}^{M} x_m^2} \tag{25}$$

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 $\bar{R}^2$ for an M observations, K explanatory variables regression:

$$\bar{R}_n^2 = 1 - \frac{\hat{\varepsilon}'\hat{\varepsilon}}{y'y}\frac{M-1}{M-K} \tag{26}$$

is now an inconsistent estimate of Q^2 and is upwardly biased by endogeneity:

$$\lim_{M\to\infty} \bar{R}_n^2 = 1 - \frac{E[\varepsilon^2]}{E\left[y^2\right]} + \frac{E[\varepsilon x]}{E\left[y^2\right]E[x^2]} = Q_n^2 + \frac{E[\varepsilon x]}{E\left[y^2\right]E[x^2]} \tag{27}$$

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.

C Monte Carlo simulation

Our Monte Carlo simulation procedure consists of the following six steps:

1. Pick a $\tilde{Q}^2 \in \{0.1, 0.2, .., 0.8, 0.9\}$. Choose s_b^2 and s_x^2 such that $\frac{s_b^2}{s_b^2+s_x^2} = \tilde{Q}^2$ and $s_b^2 + s_x^2 = 4$. Set c equal to 7.5.[16]

2. Draw a_j from a lognormal distribution for each of the J destinations. Draw b_ω from a $n\left(0, s_b^2\right)$ normal distribution for each of the W firms. Draw $x_{\omega j}$ from a $n\left(0, s_x^2\right)$ normal distribution for each of the $J \times W$ observations.

3. Generate $r^*_{\omega j}$ and $r_{\omega j}$ according to equations (1) and (4).

4. Obtain parameter estimates $\hat{a}_j, \hat{c}, \hat{s}_x^2$ and $\hat{s}_b^2$ via the MCEM procedure described in section 4. Calculate $\hat{Q}^2 = \frac{\hat{s}_b^2}{\hat{s}_b^2+\hat{s}_x^2}$.

5. Repeat steps 2, 3, and 4 ten times.

6. Repeat steps 1-5 for all $\tilde{Q}^2 \in \{0.1, 0.2, .., 0.8, 0.9\}$.

[16] These values were chosen to so that between 30 and 70 percent of the observations would be truncated.

Table 1: Descriptive statistics, 2003

	CN8 product level		HS6 product level	
	Full sample	Restricted sample	Full sample	Restricted sample
Number of				
Products	5339	491	3331	480
Firms	4304	3607	4304	3790
Destinations	223	79	223	84
Observations	155426	56297	145304	66488
Trade volume (billion DKK)	182	68	182	91
Destinations per firm-product				
Mean	3.4	2.8	3.5	3.0
Median	1.0	1.0	1.0	1.0
Min	1.0	1.0	1.0	1.0
Max	138.0	58.0	138.0	58.0
Firms per product-destination				
Mean	2.2	10.3	2.6	10.9
Median	1.0	7.0	1.0	8.0
Min	1.0	5.0	1.0	5.0
Max	373.0	373.0	412.0	412.0

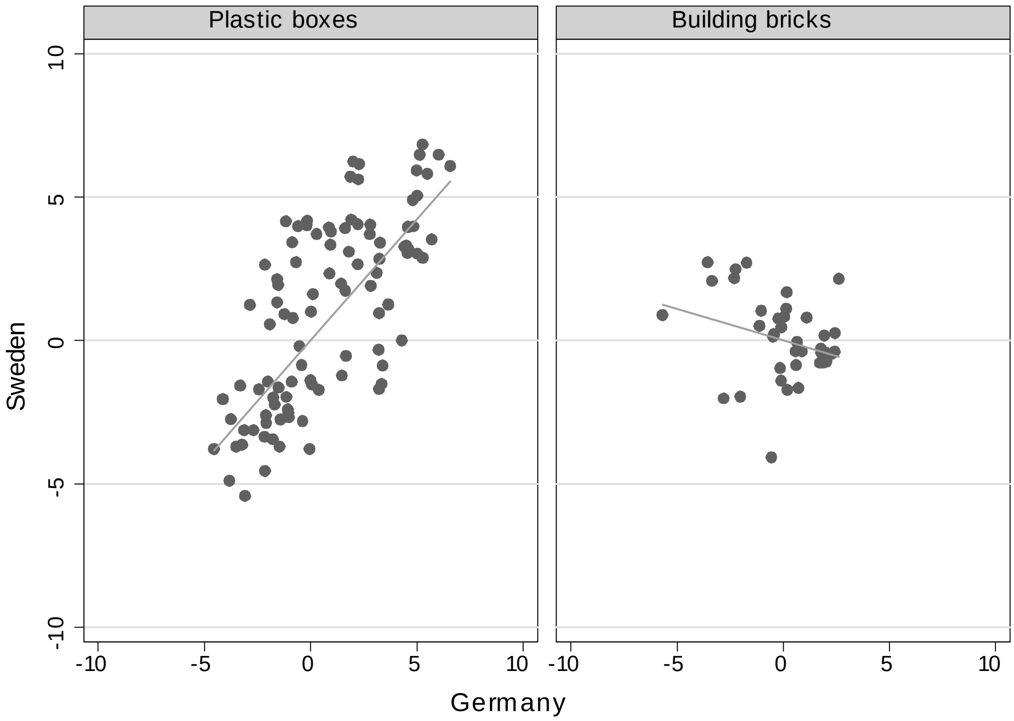

Figure 1: Sales relative to other Danish firms in Sweden and Germany for Danish Exporters of plastic boxes (left panel) and building bricks (right panel). Statistics for the lines with fitted values: Left panel: slope = 0.84, std.err. = 0.08, $R^2 = 0.54$. Right panel: slope $= -0.22$, std.err. = 0.12, $R^2 = 0.08$.

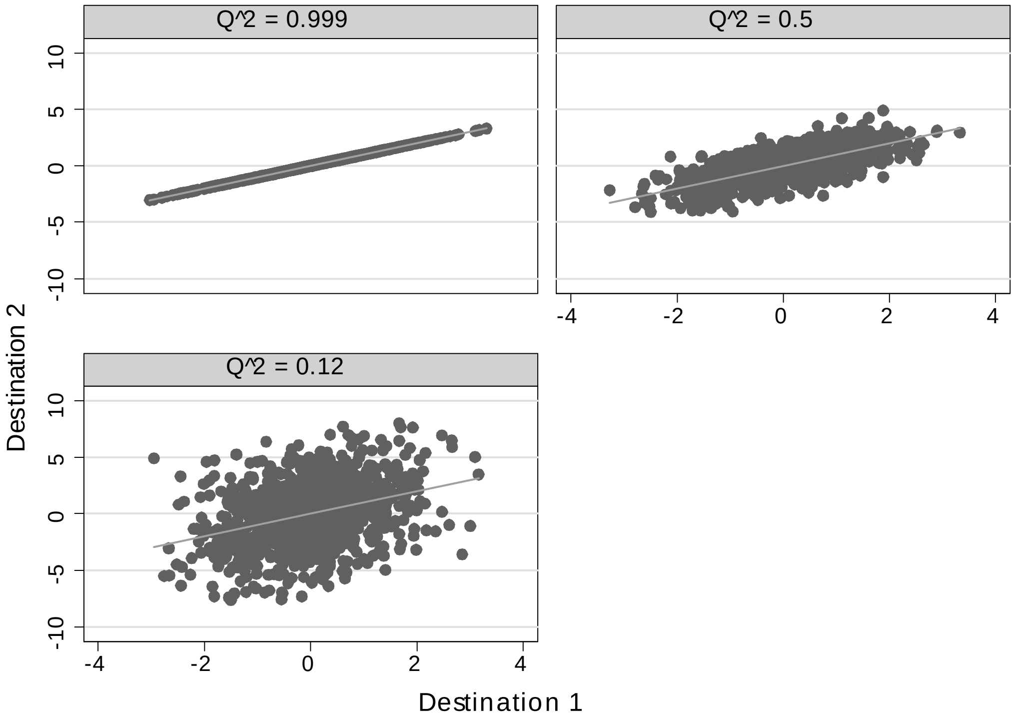

Figure 2: Simulated relative sales in two destinations for varying values of R^2.

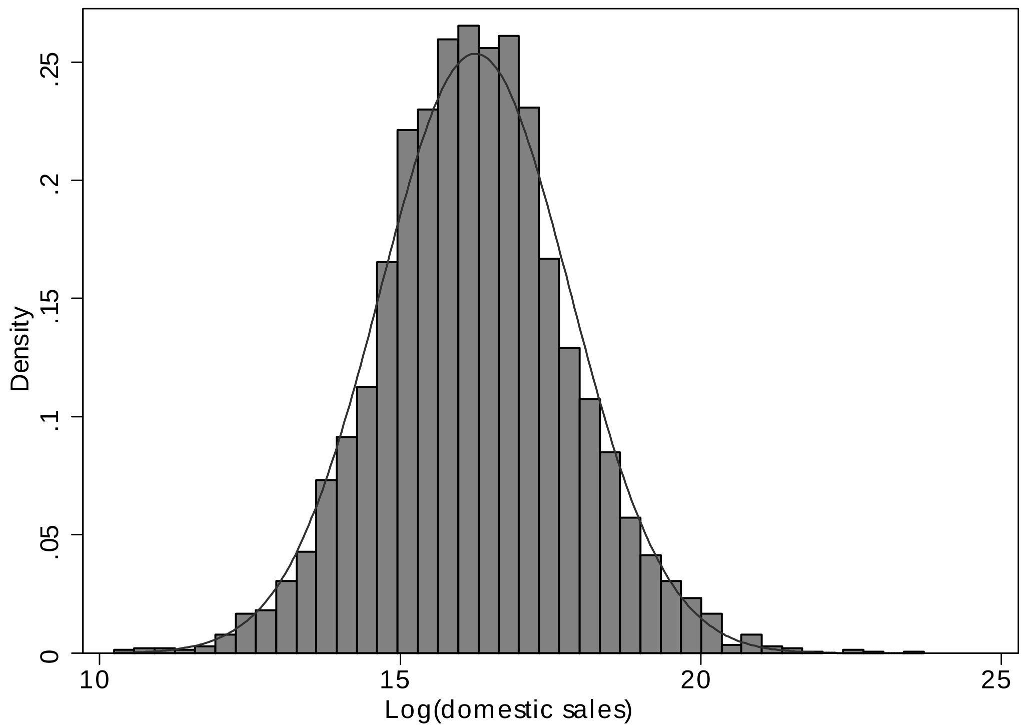

Figure 3: The distribution of log domestic sales for Danish manufacturing firms, 2003.

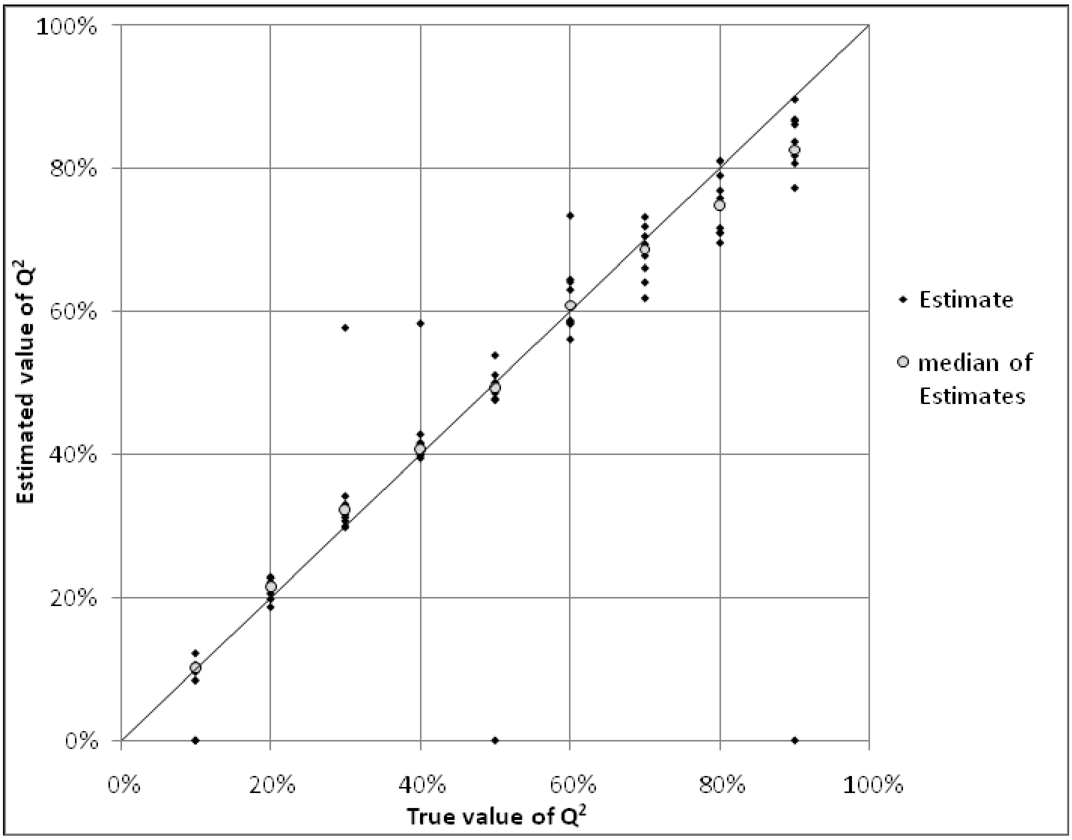

Figure 4: Monte Carlo Simulation results for nine values of Q^2 with ten repetitions each. The MCEM-MLE estimates are compared to known true values. The circles indicate the median values of the estimates. Estimates lying on the 45^o are exactly equal to the true value.

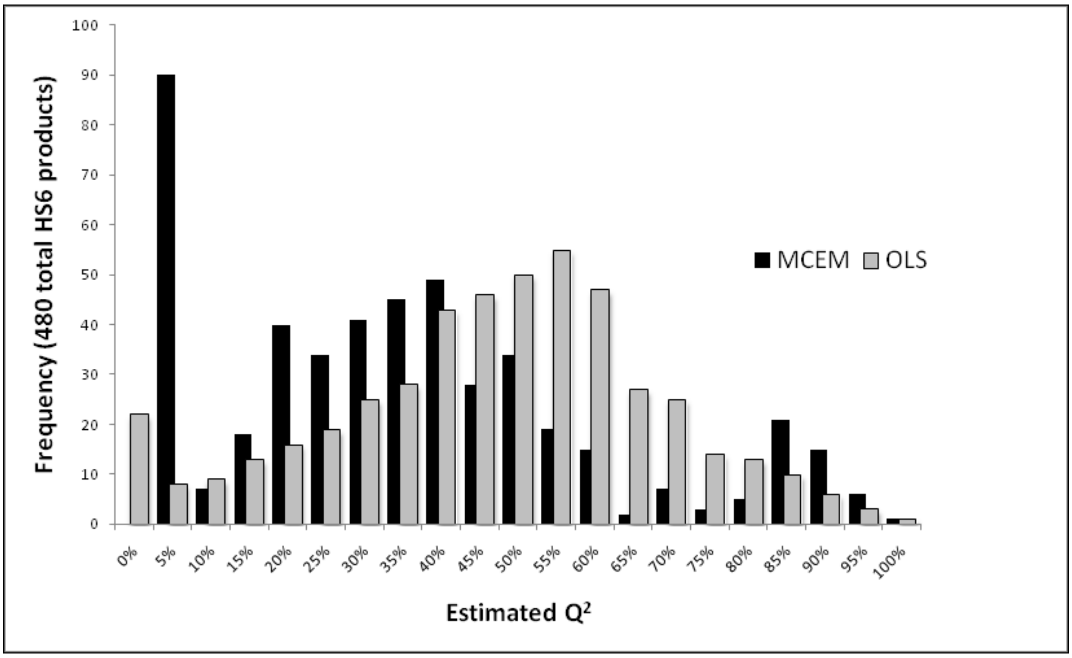

Figure 5: Estimated values for Q^2, the contribution of firm-specific effects, for 2003 Danish exports at the HS6 level.

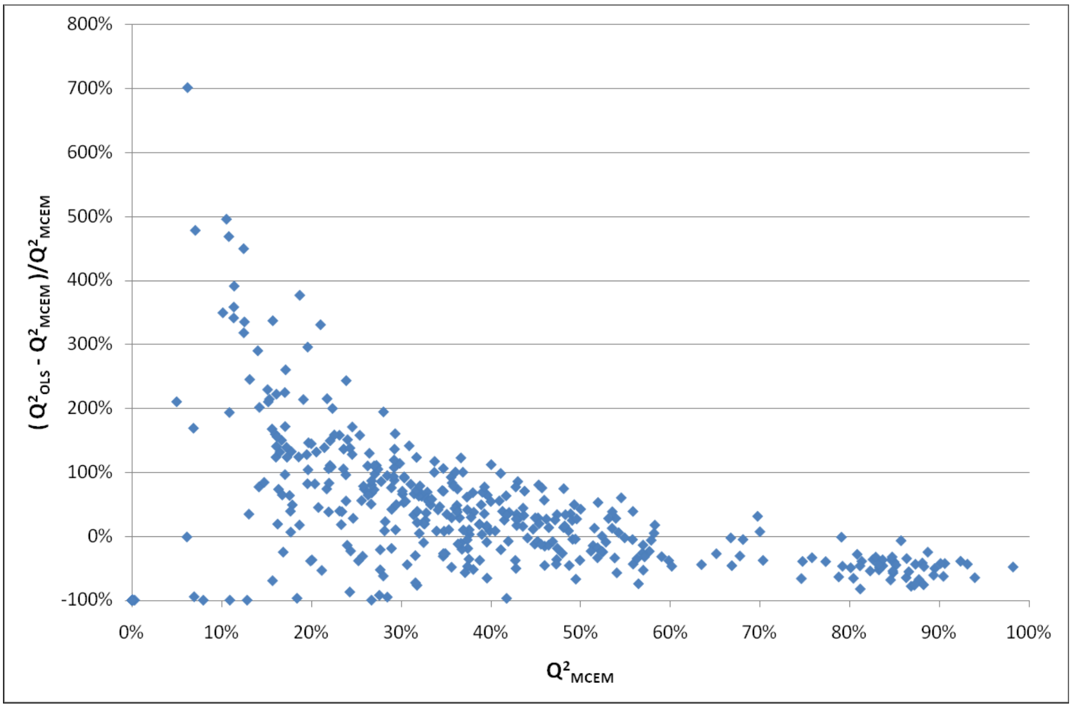

Figure 6: Comparison between MCEM and OLS estimates for the contribution of firm-specific effects.

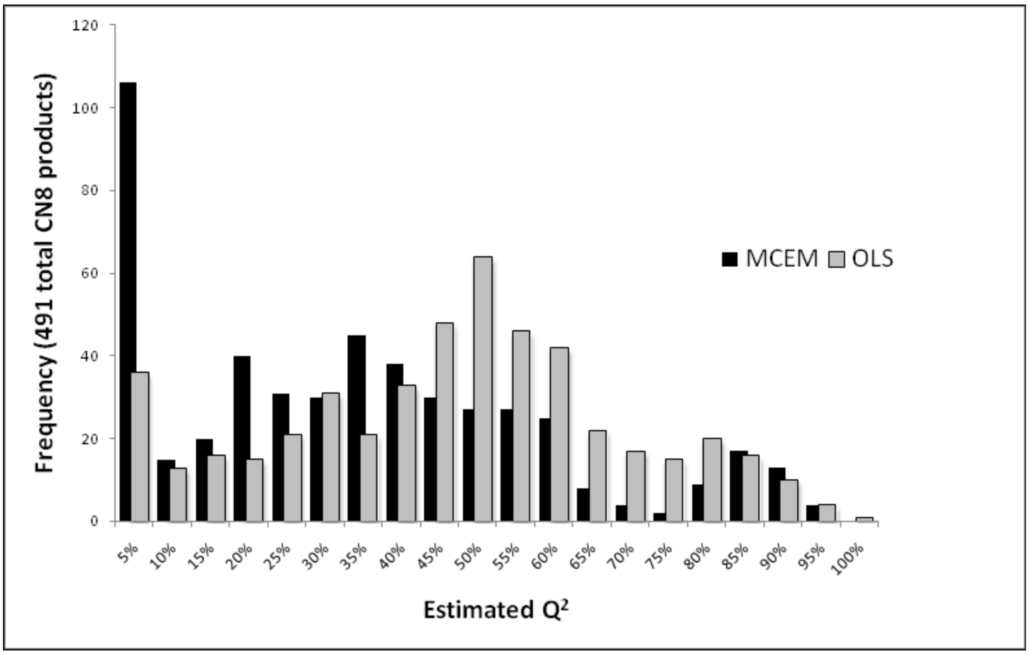

Figure 7: Estimated values for Q^2, the contribution of firm-specific effects, for 2003 Danish exports at the CN8 level.

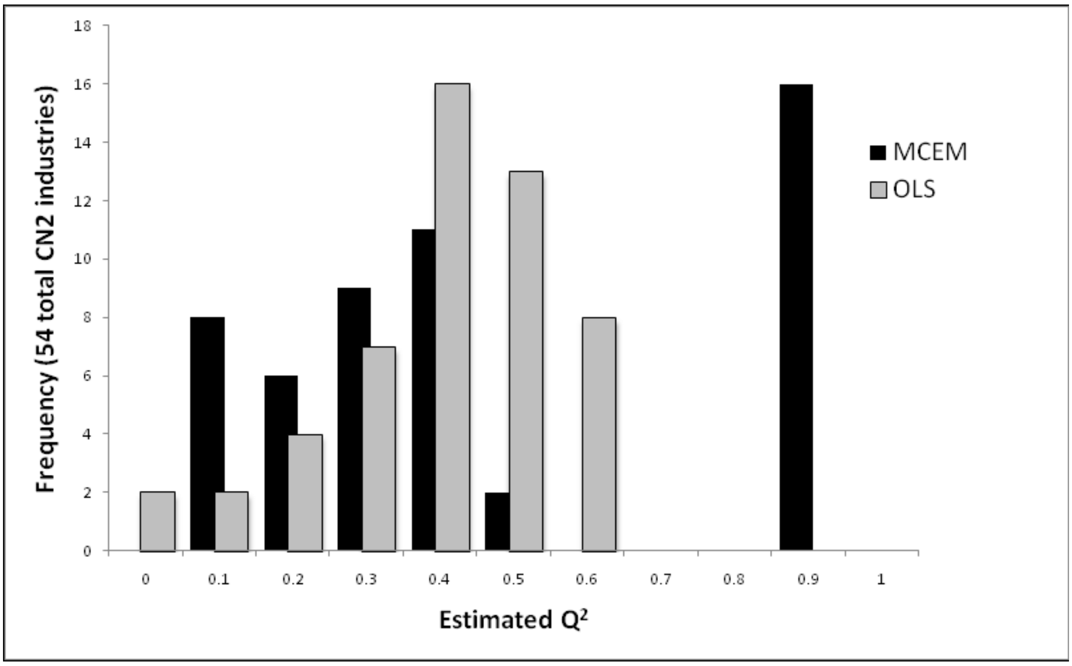

Figure 8: Estimated values for Q^2, the contribution of firm-specific effects, for 2003 Danish exports at the HS2 level.

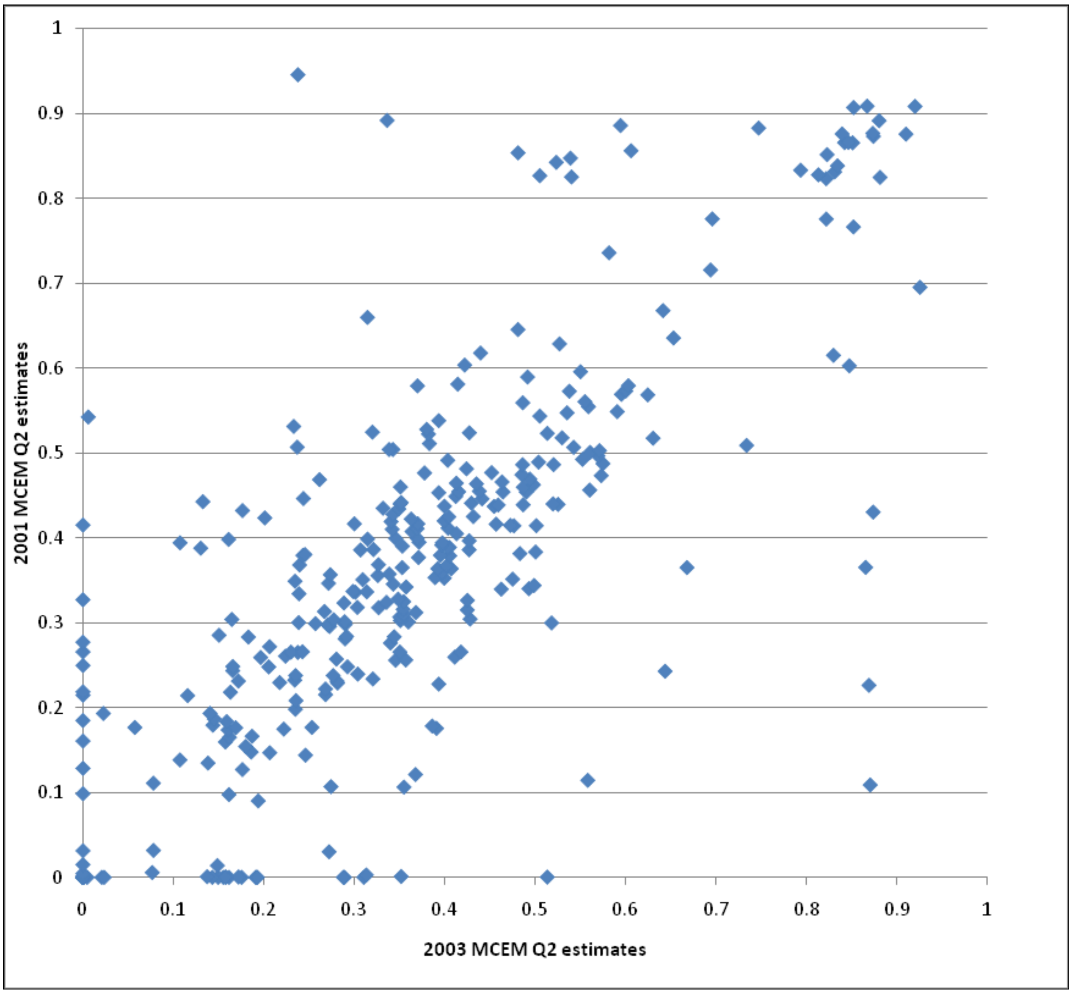

Figure 9: Point estimates for Q^2_{MCEM} for the years 2001 and 2003 for HS6 Danish Exports.

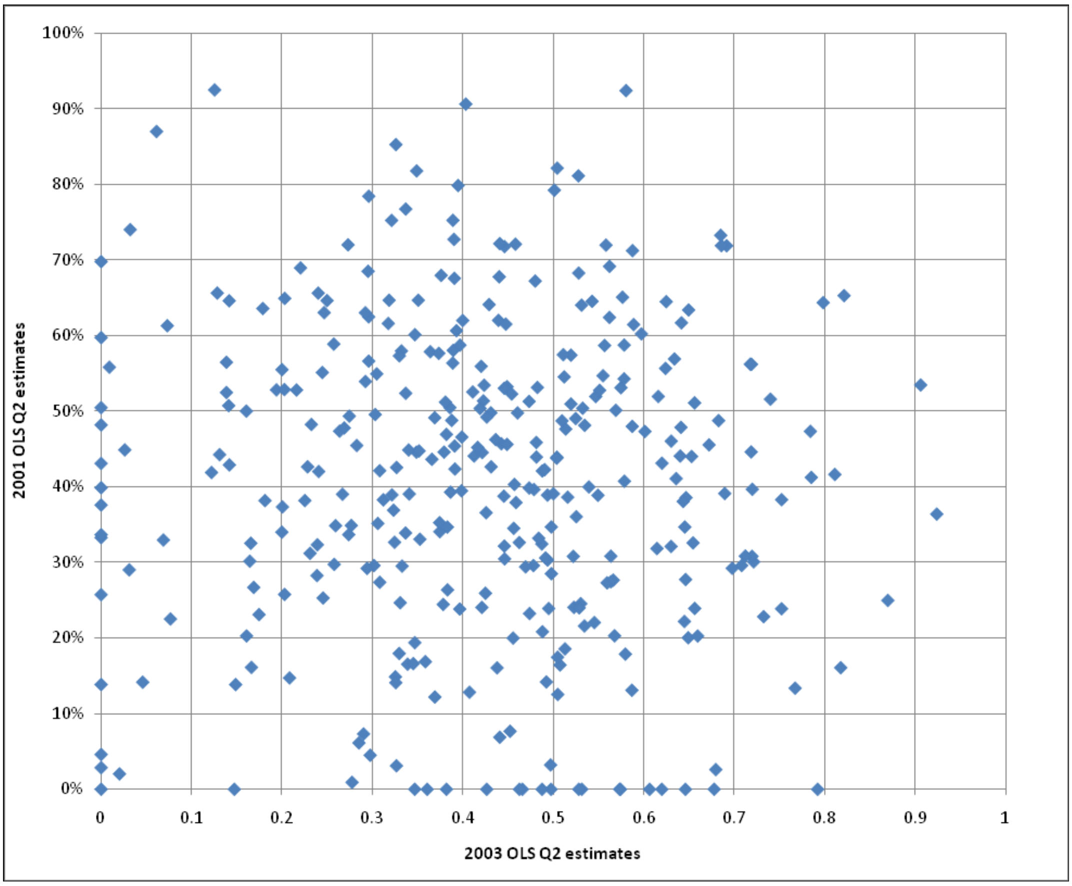

Figure 10: Point estimates for Q^2_{OLS} for the years 2001 and 2003 for HS6 Danish Exports.

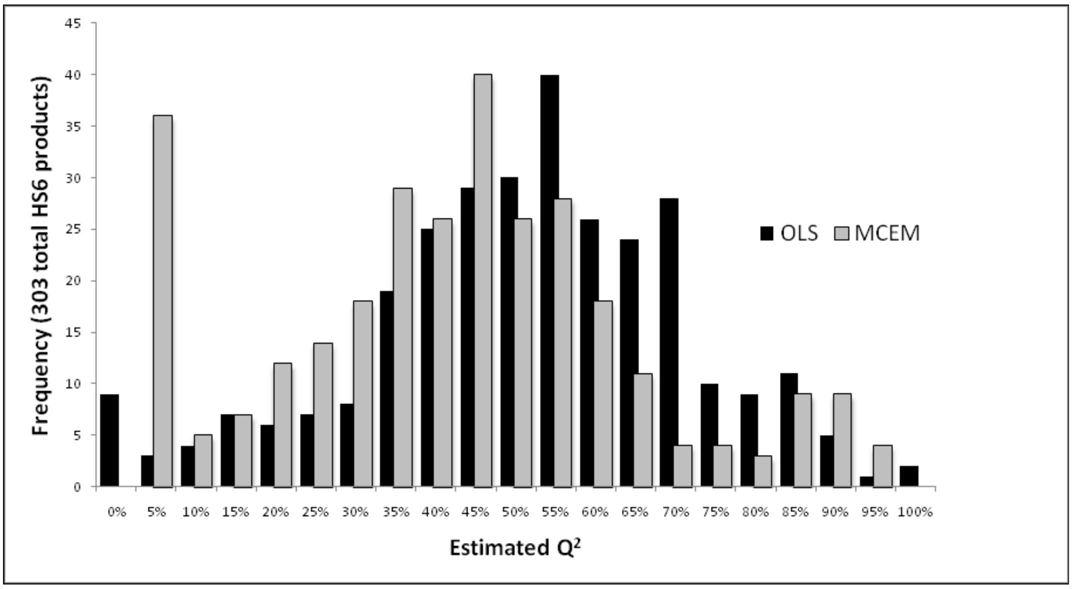

Figure 11: Estimated values for Q^2, the contribution of firm-specific effects, for 2003 Danish exports at the HS6 level. The sample includes only firms that also exported to the same destination in 2002.

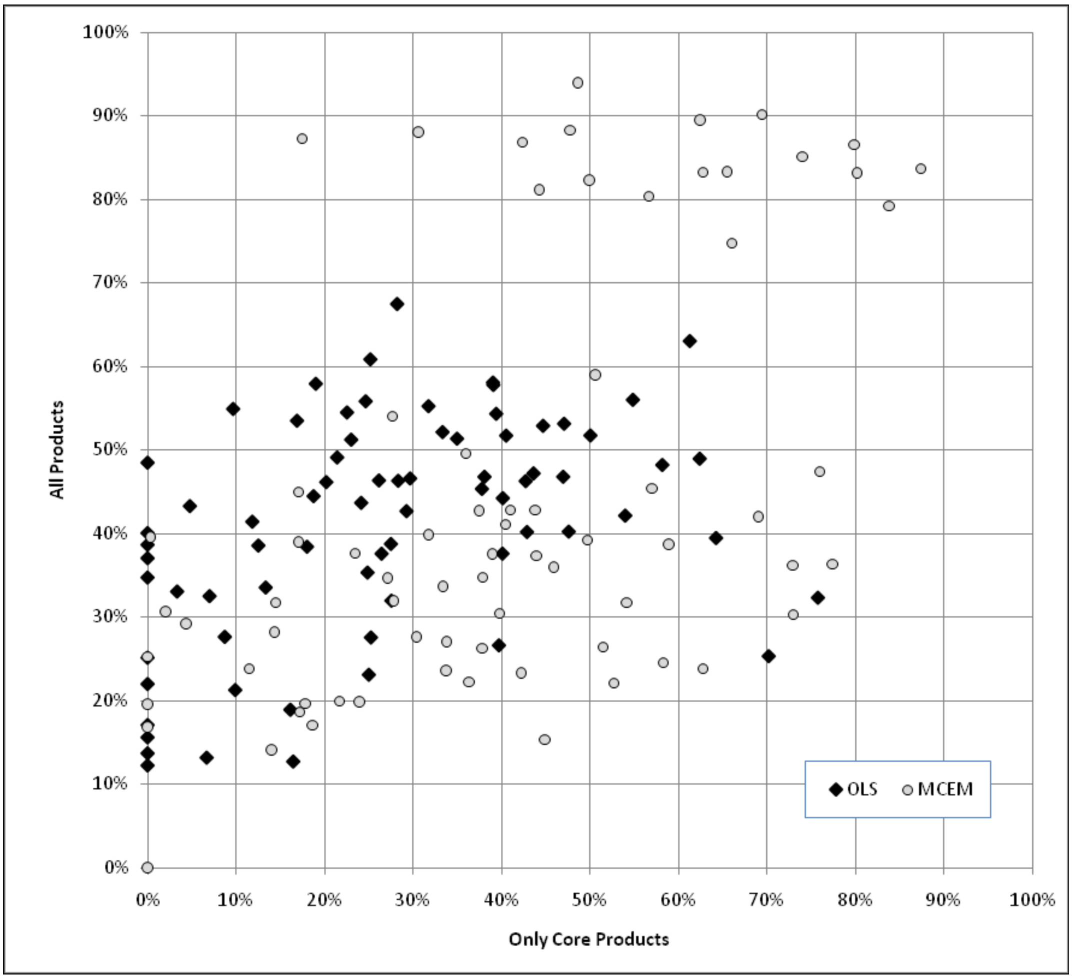

Figure 12: Estimates of Q^2 using All and only Core products, using the MCEM and OLS techniques, for 2003 Danish HS6 Exports.