Firm Export Dynamics and the Geography of Trade

Martina Lawless*
Central Bank and Financial Services Authority of Ireland
PO Box 559, Dame Street, Dublin 2, Ireland

Revised Draft: January 9, 2009

*E-mail: martina.lawless@centralbank.ie; tel.: +353 1 4344073; fax: +353 1 4344086. Thanks to Forfás for provision of the anonymous data and to Ali Ugur for help with queries on the data. I would also like to thank Karl Whelan for many useful comments and Jonathan Eaton and Samuel Kortum whose suggestions on previous work motivated much of this paper. The views expressed in this paper are the author's own, and do not necessarily reflect the views of the Central Bank and Financial Services Authority of Ireland or the ESCB.

Abstract

Two recent trends in international economics have been an increased focus on the geography of trade (e.g. what factors determine where a country exports) and the emergence of new theoretical and empirical work examining exporting activity at the firm-level. However, data limitations have prevented much progress in combining these two areas, because very few countries provide firm-level data breaking down firm exports by their destination. This paper uses a unique survey of Irish exporting firms with information on over fifty destination for a five year period to fill some of the gaps in this empirical literature. In particular we investigate how well the predications of a model of exporting with firm heterogeneity fits with the patterns of this detailed data source. Amongst our findings are that firm productivity differences are a factor in explaining the number of export markets a firm has but the prediction of a hierarchy of markets could only be weakly upheld by the data. Firm involvement in individual export markets is found to be much more dynamic than export status. Entry and exit to markets is shown to be a quantifiably important component of overall export flows, with this factor becoming more important for less popular markets. The paper also shows how the patterns of entry and exit into export markets combine to determine the overall firm-level distribution of number of markets entered.

Keywords: Firm exports; Market coverage; Market entry and exit
JEL Classification: F10; D21

1 Introduction

Recent years have seen an important development in empirical research in international economics toward the analysis of firm-level data to gain a better understanding of the processes underlying international trade, with classic contributions including the work of Roberts and Tybout (1997) and Bernard and Jensen (1995, 2004). This literature has provided substantial insights into the characteristics of exporting firms. For instance, exporting firms tend to more productive than non-exporting firms and the exporting process is very persistent, so that firms rarely change their status as either exporters or non-exporters.

This evidence has suggested a process in which individual firms face substantial barriers to engaging in trading activity, so that only the most productive can afford to do so. The new empirical findings in turn motivated the development of a new range of theoretical trade models to explain these facts such as Melitz (2003) and Eaton, Kortum and Kramarz (2008). Firm heterogeneity in productivity is the foundation of this new generation of trade model and, combined with different types of trade costs, results in productivity thresholds that allow only the best firms to become exporters.

An important limitation of the empirical evidence to date is that most of the existing datasets report information on whether a firm exports and, if so, how much it exports. However, information on the destination of exports has generally been missing from these studies. One important exception to this rule has been the work of Eaton, Kortum and Kramarz (2004) on the export destinations of French exporters. Using a cross-section of data from 1986, they find great heterogeneity in firms' involvement in exporting. Most firms sell only in the domestic market; amongst exporters, the modal firm exports to only a single foreign market; and a small fraction of firms export to a large number of markets.

This paper adds to this literature by making use of a unique survey data of Irish firms. Like the Eaton-Kortum-Kramarz data, this survey contains detailed information on firm exports and destinations. However, unlike their data, this survey is a panel dataset that follows firms over a period of five years. This dataset allows us to better understand the

dynamic processes underlying the patterns of firm exporting. We examine how the facts which emerge from the data fit with the predictions of existing models.

In particular, we begin by presenting a simple model based on Melitz (2003) with heterogeneous firms differing in their productivity endowment and having to face both fixed and variable trade costs to export. The main predictions of the model are discussed and the paper then proceeds to examine to what extent these predictions appear to hold in the data. Amongst our findings are that firm productivity differences are a factor in explaining a firm's range of export markets but the model's prediction that firms should enter export markets according to a common hierarchy could only be weakly upheld by the data. Firm involvement in individual export markets is found to be much more dynamic than their overall status as exporters. Entry and exit to markets is shown to be a quantifiably important component of overall export flows, with this factor becoming more important for less popular markets. The paper also shows how the patterns of entry and exit into export markets combine to determine the overall firm-level distribution of number of markets entered.

The paper proceeds as follows: Section 2 outlines the model and discusses five main predictions that it leads to. Section 3 explains the data source in more detail. Sections 4 to 8 then each discuss specific predictions of the model. Section 4 examines the link between exporting activity and firm productivity. Section 5 looks for evidence of a hierarchy in the markets firms export to. Section 6 examines at export growth at the firm level and how important changes in market coverage are for the firm. Section 7 then looks across markets, decomposing export growth into contributions by entrant, exitor and continuing firms. Section 8 examines the dynamics of changing market coverage, using a matrix of transition probabilities to generate predictions about the long-run distribution of export market coverage. The hypothetical distribution from this Markov-chain process turns out to be remarkably consistent with observed distribution of firms across markets. Section 9 concludes.

2 A Model with Firm Heterogeneity and Trade Costs

This section presents a simple model of firm export participation which incorporates the key features of firm heterogeneity and both fixed and variable trade costs. This is followed by a discussion of the main implications of the model for firm exporting behaviour. The remainder of the paper will examine how well these implications correspond to the observed patterns in the data.

2.1 The Model

The model is based on that first introduced by Melitz (2003) and adapted by Chaney (2008). We assume that each country produces a continuum of separate differentiated products, and that consumers in country j have utility function across a set of differentiated goods produced in all countries that takes the Dixit-Stiglitz form

$$U_j = \left[\int x_j(k)^{\frac{\epsilon-1}{\epsilon}} dk \right]^{\frac{\epsilon}{\epsilon-1}}. \tag{1}$$

Thus, the demand for good i in country j is

$$x_j(i) = \frac{p_j(i)^{-\epsilon} Y_j}{P_j^{1-\epsilon}} \tag{2}$$

where $p_j(i)$ is the price charged in country j for good i, Y_j is real income in country j and P_j is the price level defined by

$$P_j = \left[\int p_j(k)^{1-\epsilon} dk \right]^{\frac{1}{1-\epsilon}}. \tag{3}$$

Given our empirical focus on Irish exports, we will focus on the model's predictions for the exports from a specific country. Each firm in the country produces a differentiated product using a Ricardian technology with cost-minimizing unit cost $\frac{c}{a}$. The productivity parameter a is assumed to be randomly drawn from a distribution with cumulative density function $G(a)$. In additon to production costs, there are two types of trade costs associated with exporting to country j. First, there are fixed costs F_j. These can be viewed as related

to the bureaucratic paperwork costs associated with exporting, to marketing costs, and to the costs of running a wholesale and retail distribution chain. It is likely that each of these costs increase with the scale of exports; however, it is also likely that many of these costs need to be incurred independent of the scale of subsequent export sales. Second, there are variable costs, which are modeled with the iceberg specification so that τ_j units have to be shipped from our country of interest to country j for one unit to arrive. These can be viewed as transport costs, tariffs, and the variable costs associated with marketing and distribution.

The assumptions about market structure and trade costs imply that the optimal selling price to country j for a good produced with technology level a is

$$p_j(a) = \frac{\epsilon}{\epsilon - 1} \frac{\tau_j c}{a}. \tag{4}$$

Exports to market j of a firm with technology level a are therefore

$$s_j(a) = \left(\frac{\epsilon - 1}{\epsilon} \frac{P_j a}{\tau_j c} \right)^{\epsilon - 1} Y_j. \tag{5}$$

The export sales of the firm are related positively to the firm's own productivity level and to the destination country's GDP and price level. The firm's exports are negatively related to the variable trade costs. Once the decision to export is made, subsequent sales are independent of any fixed cost to exporting and these costs do not feature in the expression above.

Profits generated by this firm in country j are given by

$$\pi_j(a) = \mu \left(\frac{P_j a}{\tau_j c} \right)^{\epsilon - 1} Y_j - F_j \tag{6}$$

where $\mu = (\epsilon - 1)^{\epsilon - 1} \epsilon^{-\epsilon}$. These profits are positive as long as

$$a > \left(\frac{F_j}{\mu Y_j} \right)^{\frac{1}{\epsilon - 1}} \frac{\tau_j c}{P_j}. \tag{7}$$

This defines a cut-off level of productivity necessary for exporting to country j as

$$\bar{a}_j = \left(\frac{F_j}{\mu Y_j} \right)^{\frac{1}{\epsilon - 1}} \frac{\tau_j c}{P_j}. \tag{8}$$

In order to export profitably to country j, the firm must have a productivity level above this minimum level $\bar{a}_j$. Thus, the fraction of firms that export to country j is $1 - G(\bar{a}_j)$. The cut-off level of productivity is increasing in both types of trade costs and in domestic cost levels, while it is negatively affected by the GDP and price level of the destination country j.

2.2 Implications

This model of heterogeneous firms gives rise to the following testable predications, which will be brought to the data in the following sections.

First, the key assumption of the model is that firms are heterogeneous in their productivity levels. In order to profitably export, firms must be able to cover the fixed costs of entering the destination market j. This gives rise to a minimum level of productivity that firms must achieve in order to export (defined as $\bar{a}_j$ above). There will exist a relationship between exporting and productivity whereby exporters will be more productive than non-exporters. Furthermore, as the productivity cut-off will vary across markets, the model predicts that the higher productivity firms will be present in more markets. That exporters have a productivity advantage over non-exporters has become a stylised fact in the literature on firm-level exporting. Models of this sort imply that a link also exists between productivity and the market coverage of exporters.

The different productivity cut-offs across markets give rise to a second prediction, which is that firms will enter markets in a specific order. Assuming that trade barriers are the same for all firms, if the firm is efficient enough to enter the k-th market, then by definition it is efficient enough to export to all markets with a lower productivity threshold than k's. There should exist a "hierarchy" of markets with firms entering export markets in the order of these cut-off points. If the model's assumption that firms differ only in their productivity endowments is correct, then firms in, for example, three markets, should all be exporting to the *same* three markets and if their productivity should increase to bring them above the next productivity threshold, we should see them enter the *same* fourth market. Knowing

the identity of the countries to which the firms export over a period of time allows us to examine the extent to which this prediction holds.

A third prediction of the model relates to how a firm's sales should grow as they enter more export markets. The expression for the cut-off to any market j shows that firms will enter low-GDP high-trade cost markets later. This implies that firms will tend to progressively sell less in each additional market as they move down the hierarchy and begin to enter these less popular markets.[1] In addition, as productivity increases, firms will tend to increase their sales more in markets that are higher up the hierarchy. Recall from equation (5) that a firm's productivity affects its sales as well as its participation in each market. Putting these implications together, we would expect to see export growth at the firm level come more from adding to sales in existing markets than from sales in new markets.

Turning from firms to looking at specific markets, a fourth implication of the model is that entering and exiting firms will be the more marginal firms, and will likely contribute less to export growth than continuing firms. This is because entering and exiting firms are the lowest productivity firms in the market, having just crossed or fallen below the productivity cut-off. Continuing firms on the other hand should start out with higher productivity, and hence sales, as they are further away from the minimum productivity threshold required to export in that market.

Finally, a fifth implication is that firm entry into new markets should come from changes in firm characteristics such as productivity as well as market-specific changes related to GDP growth and trade costs. As such, we should expect to see firms changing number of markets gradually. With data on firm export destinations over five years, we are able to construct a transition matrix of changes in market coverage. To the extent that market-specific shocks are important, we should expect to see higher rates of market entry and exit for firms that

[1] An exception here is a country whose fixed trade costs are out of line with their variable trade costs. For instance, a country with high GDP and low variable trade costs could have a high productivity threshold because of very large fixed costs. In this case, a firm that just moves above the threshold for this market could end up having higher sales in this market than in some of its pre-existing markets.

start out in more markets.

3 Data

3.1 Description of Data

The data used in this paper come from a survey of Irish firms undertaken by Enterprise Ireland and Forfás, which is the Irish national policy advisory board for enterprise, trade and technology and operates under the Government Department of Enterprise, Trade and Employment. The focus of survey is on Irish-owned and predominantly exporting firms. Of the 751 firms in the sample, 83% are exporters, compared to 44% of all Irish-owned firms (Central Statistics Office 2004). As such the dataset is better suited to understanding the dynamics of exporting firms rather than the determinants of export status per se, which has been the focus of most of the existing literature in this area.

The Enterprise Ireland survey records information on a number of firm characteristics such as employment, sales, inputs and exporting activity. More importantly for our analysis, the survey records detailed information on exports to over fifty individual markets. Also, in contrast to the French dataset used by Eaton, Kortum and Kramarz (2004), which is a single cross-section, the Enterprise Ireland survey provides firm-level data on five years of exporting activity (2000-2004). Together, these two features allow for a far more detailed analysis of firm export decisions than has been possible to date. For instance, one can use these data to track changes over time in individual firms' portfolios of export markets.

Comparing the total exports of the firms covered by this survey to the census totals from the Irish Central Statistics Office (2000-2004), our data cover approximately two-thirds of exports from Irish-owned firms. This was a period during which exports did not change much: The aggregate data show export growth of 3% in 2000-2001, followed by a significant decline over the next three years, falling by over 10% in 2001-2002 for example. The survey data used in this paper follow a similar but slightly less extreme pattern, the decline in 2001-2002 is 5% and a return to positive growth is observed by the end of the sample. This

difference is likely due to a slight under-representation of small firms in our sample. The export participation patterns of these firms tends to be more volatile.[2]

That the firms are Irish-owned is an aspect of the sample selection that must be emphasised, as foreign-owned firms dominate aggregate Irish exports; this is primarily due to a history of economic policy focused on encouraging export platform foreign direct investment (FDI) to the country. In 2004, foreign owned companies accounted for just over 90 per cent of the country's manufacturing exports (Central Statistics Office, 2004). Therefore, although the current sample can be considered representative of indigenous Irish exporting firms, this constitutes only a small proportion of overall Irish exports. Although having similar data on foreign-owned exporters would extend the scope of the analysis, the Irish experience of FDI-dominated exports is far from being a common occurrence. So, it is probably fair to conclude that understanding the export decisions and patterns of indigenous Irish firms is more likely to yield conclusions that apply more broadly across countries.

3.2 Descriptive Statistics

The unique feature of this dataset is that it allows us to follow firms' exporting activity in detail over time. Before exploiting this aspect of the dataset, we first report cross-sectional statistics similar to those that have been reported in other recent studies.

Previous work has found that international engagement by firms tends to be very concentrated. Bernard, Jensen and Schott (2005) find that the top 1% of US trading (i.e. both exporting and importing) firms accounted for 81% of US trade in 2000. In the case of our Irish data, exporting activity is also concentrated amongst a fairly small number of larger firms. Firms employing over 500 generated 30% of the total exports in 2004 even though they make up less than 3% of the firms in the sample. Firms with less than 25 employees, although the most numerous at almost 33% of the sample, account for a mere 3% of total exports.

[2]Gleeson and Ruane (2006) discuss the contribution by firm size in their decomposition of export participations and growth of Irish firms.

Table 1 presents summary information on the market coverage of the firms, both overall and by the different size groups. The distribution of the number of export markets is also shown in Figure 1. Consistent with the findings of Eaton, Kortum and Kramarz (2004) for France and of Bernard, Jensen and Schott (2006) for US firms, most firms export to only a small number of markets, with over one-third exporting to a single market.

The average number of markets exported to over the five-year period was 5.9, with a median of 2.8. The average number of destination markets per firm is higher than was found by Bernard, Jensen and Schott (2006). The firms in their analysis exported to 3.3 markets in 2000. The highly skewed nature of the distribution is common across the Irish, French and US firms. Only 17% of the firms in this paper export to more than 10 markets and just 3% to more than 25. Eaton, Kortum and Kramarz (2004) found approximately 20% of firms exporting to more than 10 markets and reported 1.5% exporting to over 50.

4 Exporting and Productivity

The key assumption of Melitz-style models is that firms differ in their productivity. In order to enter any export market, firms must have high enough productivity to be able to cover the additional fixed and variable costs associated with operating in that market. Previous research has shown exporters to be more productive than non-exporters; here we extend this finding by showing that productivity can also be loosely linked to extent of market coverage.

We begin by looking at some simple summary statistics. In the comparisons of exporting firms reported in Table 2, we find some consistent differences in the characteristics of firms selling in many markets relative to those in a small number of markets. Firms with greater market coverage tend to be larger in terms of employment (the first column) and there is some evidence suggesting they are more productive: the second column shows value added per employee tends to be higher in firms exporting to multiple markets (although the relationship is not strictly increasing for all levels of market coverage).

To examine in more detail this relationship between exporting and firm characteristics, we use a Bernard and Jensen (1999) style export premium test for a number of other firm features. The exporter premia test takes the form of a regression of the firm characteristics (in logs) on a dummy variable for exporting status and controlling for sector and year. Bernard and Jensen found positive and significant coefficients on export status in all regressions verifying that exporters are larger, more productive and pay higher wages than non-exporters. As the data in this paper consists mainly of firms that are already exporters, we do not replicate these regressions. Rather, we extend the premia test of Bernard and Jensen to establish if there are any consistent differences in firm characteristics for those firms exporting to higher numbers of markets relative to those in few markets.

To establish if these firm characteristics are higher for all exporters or if they can also be linked to the extent of exporting activity, a variation of the Bernard and Jensen regressions are run substituting number of export markets for export status:

$$lnX_i = a + b * lnMarkets_i + c * Sector + d * Year + e_i$$

where LnX_i denotes some of the characteristics of firm i namely employment, value-added per employee and average wage (intended here as a skill proxy). The first panel of Table 3 shows firms exporting to multiple markets have consistently higher levels of all three of the firm characteristics used. This fits with the model's implication that firms will export to more markets as their productivity increases and they are able to cross further productivity thresholds.

Going beyond cross-sectional relationships, can we find any link between *changes* in export coverage and changes in these same characteristics? The Melitz-Chaney model indicates that this should be the case: Growth in productivity in particular is necessary if the firm is to cross further productivity thresholds to expand its market coverage.

$$\%\Delta X_i = a + b * \%\Delta Markets_i + c * Sector + d * Year + e_i.$$

The second panel of Table 3 reports the coefficients for the relationships between growth in market coverage and growth in the firm characteristics. The results are considerably weaker

for growth rates than they were for levels; although all of the coefficients are positive, only average wage growth shows a statistically significant relationship with market coverage. The coefficient for value-added per employee growth is on the borderline of being significant at the 10% level.[3]

5 Destination Hierarchy?

In the Melitz model, firms differ only by their productivity levels, and thus have different levels of unit costs. If trade barriers are identical for all firms, then these cost differentials will determine which export markets can be profitably entered. Specifically, each market will have a cost threshold. If the firm is efficient enough to enter the k-th market, then by definition it is efficient enough to be exporting to all markets with a higher cost threshold than k's. Data on which markets firms export to can be used to examine the prediction of the existence of a hierarchy of markets.

5.1 Cross-Sectional Evidence

Eaton, Kortum and Kramarz (2008) point to "substantial deviations" from a predicted hierarchy in their data. Similarly, no rigid ordering of destinations is observed amongst the Irish firms. There is, however, some weaker evidence to support the idea of a sequence of export markets. This can be seen by dividing the firms into four groups according to the number of markets they export to: 1-3 markets, 4-6 markets, 7-10 markets and 11 or more. Markets are ranked by the number of firms exporting to them.[4] For each group of firms, the percentage exporting to each market is graphed in Figure 2. Almost all firms export to the top market (the UK). The second most highly ranked market is exported to by 15% of firms in the 1-3 markets group, 51% of firms in 4-6 markets and 75% of firms who export

[3]These results do not allow us to say anything about the direction of causation. If exporting results in learning by the firm, then increasing market coverage could increase productivity rather than vice versa. See Clerides, Lach and Tybout (1998) for more on this hypothesis.

[4]The set of markets and the number of firms exporting to each are reported on Table 6.

to more than 10. Moving left down the ranking of markets, there is no crossing of the lines representing the different firm groups - this can tentatively be interpreted as evidence that firms only export to less popular markets if they also have an export presence in the more popular destinations.

This suggests that a hierarchy of markets may exist, but is not necessarily the same for all firms. One potential explanation for deviations from the strict hierarchy is that preferences for products may not be identical in all markets. Another is that trade barriers across countries may not be the same across all sectors. Finally, it should be kept in mind that the popularity ranking of the markets in our dataset may not be precise enough. There are a large number of markets with very similar numbers of exporters, so their relative positions may be determined by the presence of just a few additional firms.

5.2 Entry and Exit in the Hierarchy

An alternative way to assess the idea of a hierarchy in destinations is to look at entry and exit by market popularity. To do this, the markets are ranked 1 to 53, with 1 being the most popular market (UK) and 53 being the least popular (Tunisia). This allows us to identify the lowest ranked market for each firm. If the theoretical prediction is at all accurate we would expect a firm entering an additional market to enter a less popular market than those it already serves. Likewise, exiting firms should be moving out of their least popular markets first. This is a fairly weak test in that it does not say firm's exporting to the k-th most popular market have to next move to the $k+1$-th as the strict hierarchy would suggest. It therefore allows for particular markets to be skipped over by firms (for whatever reason), so long as the general pattern is of movement from exporting to highly ranked markets first followed by movement into less popular destinations (or vice versa for exitors).

Figure 3 plots changes in market coverage against changes in the rank of the least popular previous market. Although some exceptions do exist, the vast majority of observations fit with the conjecture that firms increasing market coverage are moving into lower ranked markets (upper-right quadrant) and those reducing market coverage are exiting their lowest

ranked markets (lower-left quadrant). As with the cross-sectional evidence on the hierarchy, the entry and exit patterns show some support for the model's predictions for an ordering of markets, albeit with a considerable amount of heterogeneity.

6 Firm Dynamics

This section looks at the implications of the model for how an individual firm's exports change, both within and across markets. The productivity cut-off derived earlier showed that firms will tend to enter low-GDP high-trade cost markets later. This implies that firms will tend to progressively sell less in each additional market as they move down the hierarchy and begin to enter these less popular markets. In addition, as productivity increases firms will tend to increase their sales more in markets that are higher up the hierarchy so we would expect to see export growth at the firm level come more from adding to sales in existing markets than from sales in new markets.

Referring back to the summary statistics presented in Table 2, the fourth column shows that firms selling in multiple markets have a fairly similar level of exports per market as those in only one or two markets. This result initially appears counterintuitive given that these firms with many markets export much more in total (third column). The apparent puzzle is resolved when exports to particular markets are compared. Taking the UK as an example: the fifth column of Table 2 shows that firms with more export markets sell much more in the UK than do firms with few export markets. The figure for average exports per market is reduced for firms with many markets because the additional markets they operate in tend to be smaller in scale. These figures denote a pattern of firm export growth in which firms *both* increase sales in their existing markets and expand their portfolio of markets into new destinations.

Table 4 examines the relative importance for firm growth of the addition of new markets relative to increasing exports in existing markets. This is done by decomposing the export growth of each firm into four components: the contribution of entering a new market, of

exiting a market and of increasing and decreasing sales in existing markets. Looking across all firms, entering and exiting markets had close to offsetting effects. The contributions of increasing and decreasing sales in existing markets were almost five times larger than those of entry and exit. For example, in 2000-2001 export sales in expanding existing markets contributed 19% to total export growth, this was offset by -14% in contracting existing markets. In the following year these rates were reversed, but remain substantially larger than the contributions of entry and exit of markets. The four components sum to give the total export growth for all firms in each year. Repeating the exercise separately for firms that are increasing and decreasing their total exports, the contributions of changing sales in existing markets always dominates the contribution of entering or exiting markets, as implied by the model. However, it is interesting to note that firms with expanding exports may still exit or have declining sales in some of their markets and firms that are contracting overall exports may still be growing in some markets.

Finally in this section, we look at how export growth rates vary over different groups of firms. In Table 5 we divide firms into quintiles according to their initial exports and look at their growth rates and the contribution to total export growth of each group of firms. We find that firms in the lowest quintile subsequently grew the fastest. However, mirroring the findings of Eaton, Eslava, Kugler and Tybout (2008) on Columbian firms, this fast-growing quintile contributed very little to total export growth. The top quintile in terms of initial exports contributed most to the overall growth over the period. Taking 2003-2004 as an example, aggregate export growth was 5%. The highest quintile of firms grew their exports at 4%, making up 73% of the total growth. The lowest quintile contributed just 2% of the total because of their small initial size, although their own exports grew the most rapidly of all the groups (by 21%). Linking this back to the implications of the Melitz-Chaney model, the largest exporters are likely to be the most productive and established in any market and therefore it is not particularly surprising that they make up the bulk of the contribution to total export growth.

7 Market Dynamics

We now turn from examining the number of markets that firms export to, and focus instead on the dynamics of entry and exit viewed from the viewpoint of specific markets. The Melitz-Chaney model implies that entering and exiting firms should be more marginal than existing exporters, as they will have productivity levels very close to the threshold for participating in the market. Section 7.1 looks at how prevalent entry and exit are across markets. Section 7.2 then presents a decomposition of export growth into contributions from entering, exiting and continuing firms.

7.1 Patterns of Entry and Exit by Market

Table 6 documents the distribution of firms across markets and the levels of entry and exit averaged over the time period. Unsurprisingly, given its proximity and historical links, the UK is the predominant export destination for Irish exporters. The 584 firms who sell at least some of their exports to this market represent 94% of the sample. The second largest market (the USA) has less than half of the number of firms exporting to it than the UK. With the exception of the US, the top ten markets for Irish firms are all located in Western Europe.

The most striking feature of the data presented in Table 6 is the extent to which entry and exit exists in all markets. Although the number of exporters in each destination changes only slightly over the period, the gross flows in to and out of markets strongly outweigh net changes in firm numbers. To take the UK as an example, there is an average increase in exporter numbers of 4 each year, which is a small change relative to the total. The underlying pattern, however, is perhaps more dynamic than this relatively small net change might imply. An average of 30 exporters began selling to the UK each year and an average of 26 exited. As a percentage of existing exporters, the UK has one of the lowest rates of entry and exit, presumably due to the fact that the majority of exporters are already present in that particular market. In general, the rates of entry and exit to and

from markets tend to increase as we move from more popular to less popular destinations.

In the context of the model, this information on entry and exit rates gives us an indication of how likely firms are to cross the productivity cut-offs and perhaps to what extent fixed costs of market entry inhibit turnover amongst exporters. Existing literature on firm exports (e.g. Roberts and Tybout, 1997) has emphasised the persistence of exporting activity, which is commonly interpreted as implying that substantial "sunk costs" are encountered in becoming an exporter. These findings suggest that entry and exit to individual markets is substantially more fluid than the existing results on export participation implied. They also point to the importance of firm-specific factors perhaps relating to demand and the competitive environment in individual export markets. One particularly interesting pattern documented in our data is that rates of entry and exit to and from markets are significantly higher in the less popular markets than in the more popular ones. This could be interpreted as weak evidence that firms find it more difficult to assess their prospects of success in advance in markets that are less familiar to the firm and on which information on market conditions may not be as widely available.

7.2 Decomposition of Gross Export Flows

One way to make use of the panel nature of the data to decompose net changes in exports is to separate out the positive contributions from firms who start exporting or increase their exports and the negative contributions due to firms who reduce their exports or stop exporting altogether.[5] This method follows a popular approach in the literature on job flows (for example, Davis, Haltiwanger and Schuh, 1996). This decomposition has been applied to exporting firms before by Wagner (2003) for German firms and Gleeson and Ruane (2006) using Irish data. Both of these papers evaluate the contribution of entry and exit to exporting, focusing in the German case on a sizable export boom, and in the Irish case on two exceptional episodes - an export boom with growth of 35% and in contrast

[5] A fifth possible category is firms with unchanged exports across two years, but there are no observations of this in the data. Additional breakdown by identifying re-starters and re-stoppers as Gleeson and Ruane (2006) do is also not feasible given the shorter time span of data available in this paper

a collapse of 26% in net exports. Wagner's results show the majority of export dynamics are accounted for by increases and decreases in exports by existing exporters. Gleeson and Ruane find a somewhat greater contribution from entry and exit, and demonstrate substantial volatility particularly for the smallest firms with exit and re-entry being a more common feature than had previously been identified in the literature.

The top panel of Table 7 reports the results of this method of decomposing exports. Confirming the findings of the existing research, the main contributions to net export change come from changes in exports by incumbent firms, with starters and stoppers featuring only marginally. For example, the 4% growth in net exports in 2000-2001 comes from increaser firms changing exports by 12% and decreasers by -8%. Starters and stoppers contribute less than 1% in all years. The survey nature of the data and its focus on exporters may explain why the entry and exit contributions are significantly smaller in these results than in Gleeson and Ruane (2006). It should also be noted that the period covered by the current data was one of relatively stagnant export performance, unlike the other papers which are examining significant boom or bust episodes.

Moving beyond the contribution of firms who start exporting or quit it altogether, the evidence of significant rates of entry and exit into individual markets raises the question of how important these transitions are for aggregate exports. This can be addressed by applying the Davis-Haltiwanger decomposition to each individual market. In other words, net export growth can also be separated into increases and decreases in exports across individual markets, and within these markets one can calculate the importance of entry, exit, increasers and decreasers. The four groups again combine to give the net change in exports, this time the contributions being from markets rather than from firms. While entry and exit to exporting was only a tiny proportion of export changes, the contribution of changes in market coverage are substantially larger. The middle panel of Table 7 shows firms entering a new market generate gross exports of between 2-5% of the total. That said, there is still confirmation for the Melitz-Chaney model prediction that existing firms are likely to be larger and more important for aggregate export flows: The bulk of gross

export flows still come from incumbent firms in increasing or decreasing exports.

The contributions of starters and stoppers to exporting in different destinations may however be under-estimated by the decomposition above. Comparing measurements of export of incumbent firms and firms switching markets has two potential drawbacks. The first is a matter of accountancy; incumbent firms are reporting exports for an entire year, whereas it is highly improbable that this is true of firms changing their markets. Unless all firms entering a new market do so on the first day of the accounting period, we are not comparing like with like. A second potential issue could come from assumptions made about the nature of sunk costs in exporting. If, as several papers have suggested (e.g. Roberts and Tybout, 1997), a sunk cost is encountered on becoming an exporter and this is borne completely within the first exporting period it is fairly reasonable to assume that firms take some time to establish their exporting activity.

An attempt to address this issue is to redo the decomposition of changes in exports by market over varying time periods. The final panel of Table 7 reports these results. Focusing on the final column, contributions to net export growth (or rather decline in this instance) between 2000 and 2004 are calculated. Firms who were exporting to a market in 2004 where they had not been present in 2000 added 7%, while firms who stopped exporting to a market they had exported to in 2000 contributed -10%. Increasers and decreasers essentially cancel one another out with gross changes of 24% each. Taking this approach, the contributions of firms entering and exiting markets is more evident, but the conclusion that incumbent firms are more important still holds.

8 Changes in Market Coverage

Finally, we examine how individual firms change their number of export markets over time. According to the model outlined earlier, changes in number of markets are the result of either changes in firm productivity or changes in export market characteristics such as GDP or trade costs. The model suggests that the productivity cut-offs represent significant

barriers to moving into new markets and therefore one would expect changes in market coverage to be gradual. Because the cut-offs are determined by market GDP and trade costs, market specific-shocks could also result in firms changing their market coverage. This can be further extended to a predication that firms in many markets would be more likely to change coverage as they are more likely to face a market-specific shock in one of their destinations.

Figure 1 showed a snapshot describing the average distribution for the number of markets prevailing over the five years sample available. We can now use the panel dimension of our dataset to examine the process by which firms move between various points in this distribution–for instance from n to $n+1$ or $n-1$ markets–as well as the forces that determine the shape of the cross-sectional distribution.

8.1 Entry and Exit Patterns

Table 8 is a summary of the matrix of transition probabilities, characterised as changes in market coverage by number of existing markets. The probabilities are calculated as averages of actual movements for the five years of sample data. Changing coverage by more than plus or minus four markets remains highly unusual so these are not reported. This results in the matrix containing a large number of zeros.

These calculations provide a detailed description of how firms enter exporting and how their export coverage tends to change over time. Firms that enter exporting activity for the first time tend to do so in a very gradual fashion, usually only entering a single market. After this initial entry, firms are more likely to add an additional market than they are to exit but the most common pattern is for the number of markets to remain unchanged: 81 percent of firms that exported to one market still exported to only a single market in the following year. When firms that export to a small number of markets do change their coverage, they tend to do so by adding or subtracting only one more market. As coverage of existing markets then becomes larger, however, they are far more likely to be observed adding or subtracting multiple markets, but rarely by more than four markets. Indeed, once

a firm exports to more than three markets, it is more likely to change its market coverage from one year to the next than it is to stay in the same cell. As well as the decrease in persistence as market coverage increases, the probability of exiting markets begins to overtake the probability of entry.

This general picture, of limited changes for small firms combined with a more dynamic pattern for firms in many markets, is consistent with with Melitz-Chaney model in which both firm-level productivity and market-specific factors determine export participation. Firms with low levels of productivity are only able to enter those markets with the lowest productivity thresholds and their observed transitions are largely dependent on firm-level productivity changes that either push them above or below one of these lower thresholds; changes in GDP or trade costs in less popular markets have no impact on these firms. Large firms in many export markets, however, are regularly affected by changes to trade costs and demand across a range of markets.

8.2 Transitions and Long-Run Distribution

The matrix of transition probabilities can be used to address two further issues: firstly, what do we know about how this distribution evolved, and secondly can anything be said about its long-run stability? To answer these questions, we follow that approach of Quah (1993) who applied a Markov transition probability methodology to the evolution of country income distributions. This method was also employed by Eaton and Eckstein (1997) for their analysis of the evolution of city size distributions.

In this context, we use a matrix of transition probabilities to generate a hypothetical distribution of firm market coverage. Let F_t denote the distribution of firms according to the number of export markets they serve at time t and assume that the evolution of F can be described by :

$$F_{t+1} = MF_t \tag{9}$$

where M is a 25x25 transition probability matrix, which maps where points in F_t end up

in F_{t+1}. Through a process of iteration, the s-period-ahead predictor for the distribution is given by:

$$F_{t+s} = M^s F_t. \tag{10}$$

Taking M^s to the limit as $s \to \infty$ we can represent the long-run (ergodic) distribution of F_t if it exists and is unique. Using the transition probabilities in Table 8, and calculating higher powers of M^n until it converged, we generated the asymptotic distribution of market coverage consistent with the transitions observed during our five years of data. Remarkably, the implied long-run distribution fits very closely with the cross-sectional distribution observed in our data: Figure 4 plots the fitted distribution against that observed in the data.

This result shows that, while we cannot observe transitions prior to 2001, the currently-observed cross-sectional distribution is consistent with the hypothesis that transitions of this form have been determining the distribution all along. This is important because it implies that the transition probabilities reported here—which show relatively high rates of adjustment of market coverage—are not anomalous. Figure 2 also suggests that the current cross-sectional distribution is likely to be stable. In other words, even though the market transitions reported here relate to a period of increased globalisation, they imply that the cross-sectional distribution is not likely to change much.

9 Conclusions

The empirical literature on firm exporting has been restricted by a lack of data, most notably in relation to the breakdown by destination markets. This paper uses a survey of Irish firms over a five-year period, which contains detailed information on exports to over fifty markets. This allows us to document some new facts on the distribution of firms' export markets and on changes in market coverage.

These new facts largely correspond with a set of predictions drawn from a simple version of a Melitz model with heterogeneity in firm productivity and fixed and variable costs to

exporting. The evidence suggests that firm productivity differences are a factor in the number of markets a firm exports to. The model generates productivity cut-offs that a firm must cross in order to export to a given market; this results in a prediction of a hierarchy of markets that firms would be expected to enter in the same order. This prediction could only be weakly upheld by the data. Firms tended to export to more popular markets first but considerable variation across firms was observed, particularly as one looked at participation in less popular markets.

Predictions that firm-level export growth should be largely driven by their existing markets rather than by their entry into new markets was supported by the data. A similar prediction from the point of view of specific markets is that most of the growth would come from continuing firms and this also found support from a decomposition of export growth into entry, exit and continuing firms. However the rates at which firms enter and exit individual markets are much higher than models of sunk costs in exporting might have led one to suppose. Both of these results (firm growth and market growth) arise because firms that enter or exit a market should be close to the productivity threshold and are therefore less productive and sell less than continuing firms.

Finally, we analyzed how firms change their export market coverage. It appears that despite the well-documented persistence in exporting status, exporting is a very dynamic process and changes in market portfolios of exporters are a relatively common occurrence. Most firms start out by entering only a single market and then slowly add markets; in contrast, firms that are in a large number of markets exhibit frequent changes in their portfolio of destinations. Furthermore, the calculated transition rates can be used to generate predictions for a hypothetical long-run distribution, which in the Irish case turned out to fit remarkably well with the observed distribution.

References

Bernard, A.B., Jensen, J.B., 1995. Exporters, Jobs and Wages in US Manufacturing: 1976-1987. Brookings Papers on Economic Activity: Microeconomics 1995, 67-112.

Bernard, A.B., Jensen, J.B., 1999. Exceptional exporter performance: cause, effect, or both? Journal of International Economics 47(1), 1-25.

Bernard, A.B., Jensen, J.B, 2004. Entry, Expansion and Intensity in the US Export Boom, 1987-1992. Review of International Economics 12(4), 662-675.

Bernard, A.B., Jensen, J.B., Schott P.K., 2005. Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods. NBER Working Paper No.11404.

Central Statistics Office (2000-2004). Census of Industrial Production. Dublin.

Chaney, T., 2008. Distorted Gravity: The Intensive and Extensive Margins of International Trade. American Economic Review 98(4), 1707-1721.

Clerides, S.K., Lach, S., Tybout, J.R., 1998. Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico and Morocco. Quarterly Journal of Economics 113(3), 903-947.

Davis, S., Haltiwanger, J., Schuh, S., 1996. Job Creation and Destruction. MIT Press, Cambridge.

Eaton, J., Eckstein Z., 1997. Cities and Growth: Theory and Evidence from France and Japan. Regional Science and Urban Economics 27, 443-474.

Eaton, J., Eslava, M., Kugler, M., Tybout J., 2008. Export Dynamics in Colombia: Firm-Level Evidence, in: Helpman, E., Marin, D., Verdier, T. (Eds.), The Organization of Firms in a Global Economy. Harvard University Press, Cambridge, MA.

Eaton, J., Kortum, S., Kramarz, F., 2004. Dissecting Trade: Firms, Industries and Export Destinations. American Economic Review 94(2), pp. 150-152.

Eaton, J., Kortum, S., Kramarz, F., 2008. An Anatomy of International Trade: Evidence from French Firms. NBER Working Paper No.14610.

Gleeson, A.M., Ruane, F., 2006. Export Dynamics in Small Open Economies: Indigenous Irish Manufacturing Exports, 1985-2003. IIIS Discussion Paper No.140.

Melitz, M.J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. Econometrica 71(6), 1695-1725.

Quah, D., 1993. Empirical Cross-Section Dynamics in Economic Growth. European Economic Review 37, 426-434.

Roberts, M.J., Tybout, J., 1997. The Decision to Export in Columbia: An Empirical Model of Entry with Sunk Costs. American Economic Review 87(4), 545-564.

Wagner, J., 2003. On the Micro-Structure of the German Export Boom: Evidence from Establishment Panel Data 1995-2002. HWWA Discussion Paper No.249.

Figure 1

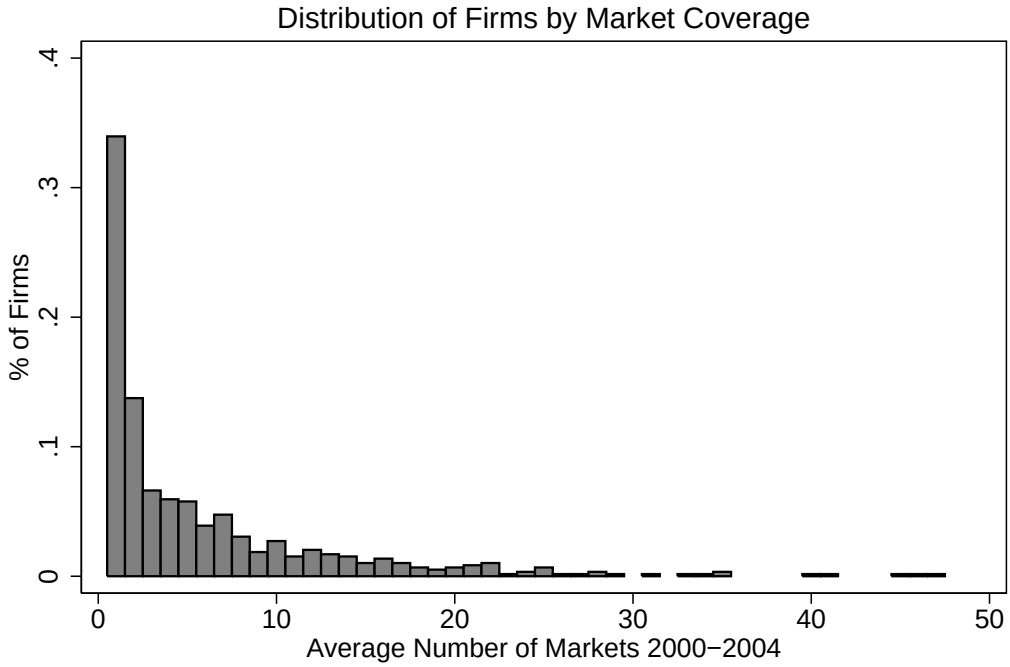
Distribution of Firms by Market Coverage
% of Firms
0
.1
.2
.3
.4
0
10
20
30
40
50
Average Number of Markets 2000−2004

Figure 2

Market Participation and Firm Market Coverage

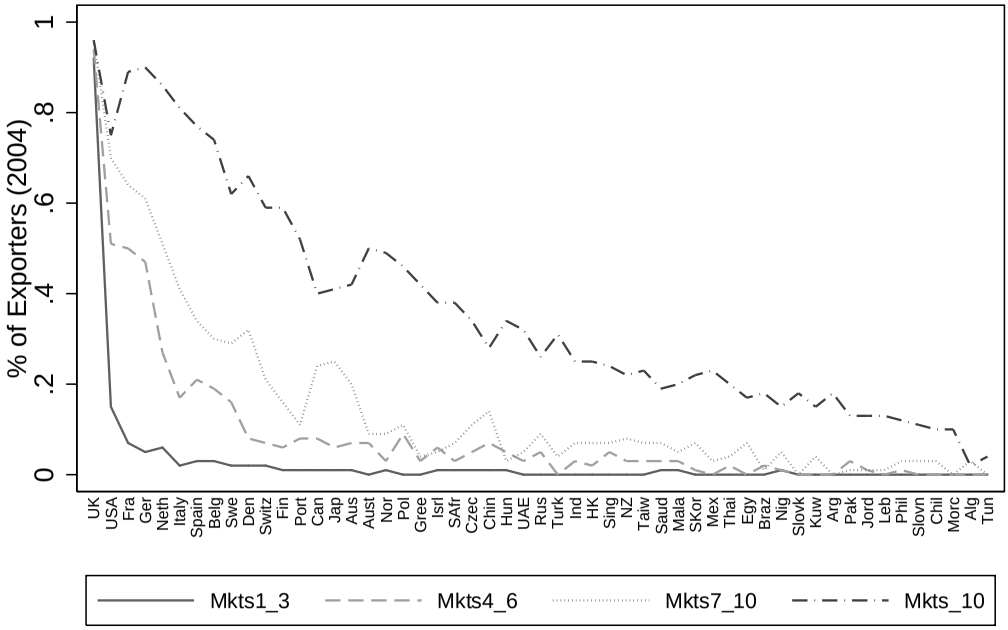

Figure 3

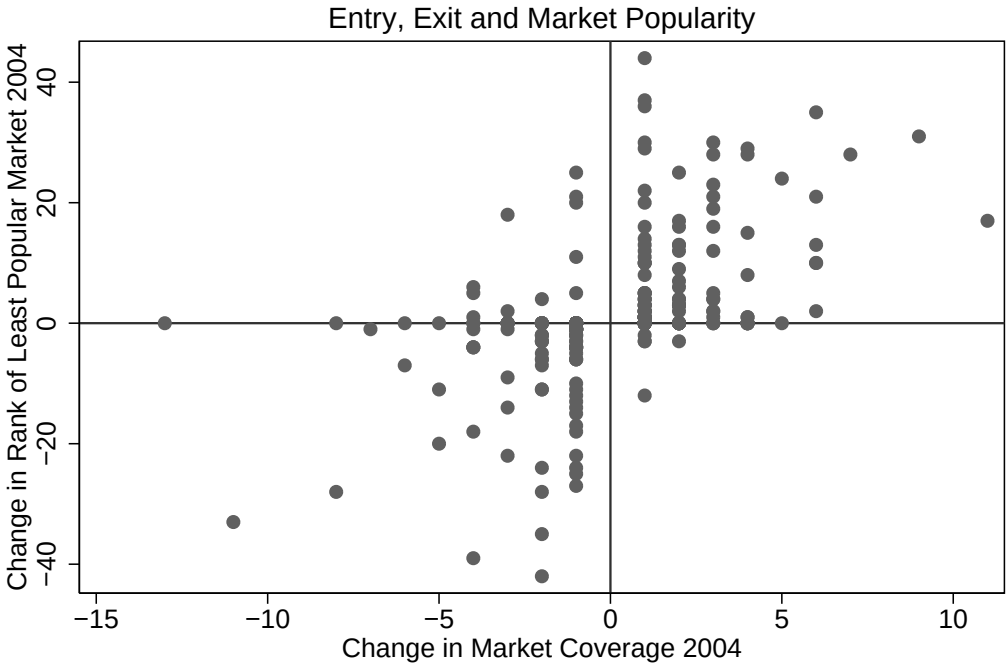
Entry, Exit and Market Popularity
Change in Rank of Least Popular Market 2004
40
20
0
−20
−40
−15
−10
−5
0
5
10
Change in Market Coverage 2004

Figure 4

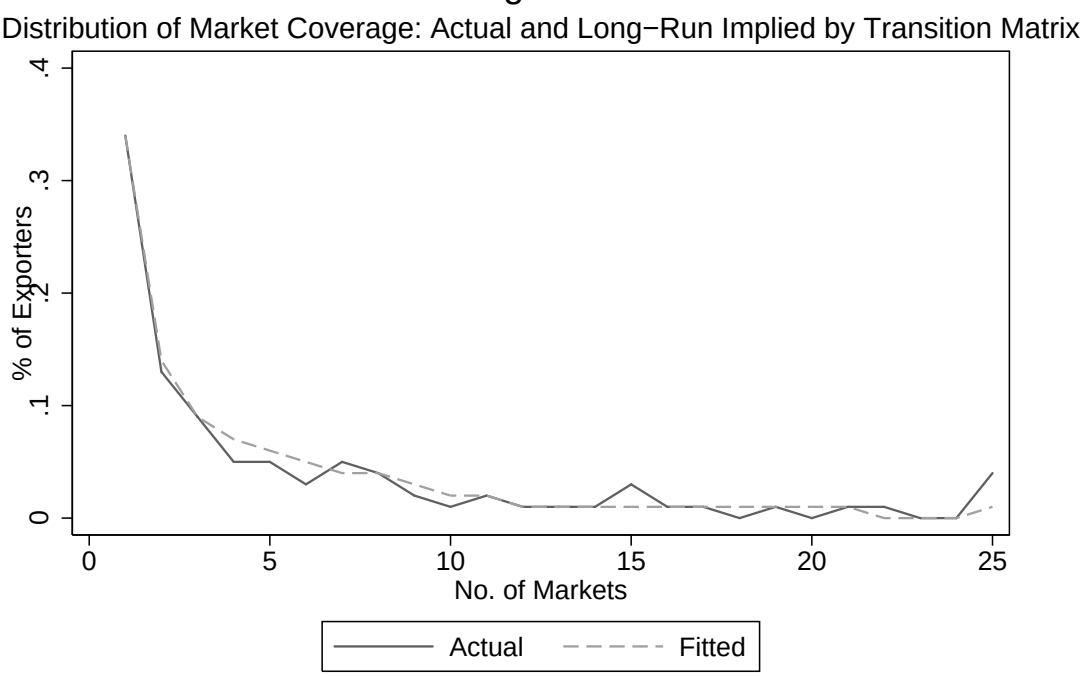
Distribution of Market Coverage: Actual and Long−Run Implied by Transition Matrix
% of Exporters
0
.1
.2
.3
.4
0
5
10
15
20
25
No. of Markets
Actual
Fitted

Table 1: Market Coverage and Firm Size (Average 2000-2004)

		Firm Employment					
	All Firms	1-24	25-49	50-99	100-249	250-499	500+
Average Markets	5.93	4.70	4.87	5.93	8.05	12.29	9.88
Median Markets	2.80	2.00	2.00	3.20	5.40	9.20	7.10
% Exporting to 1 Market	0.34	0.43	0.40	0.28	0.23	0.16	0.13
% Exporting to 2-5 Markets	0.33	0.30	0.35	0.41	0.27	0.14	0.31
% Exporting to 6-10 Markets	0.15	0.14	0.11	0.15	0.23	0.26	0.20
% Exporting to 11-25 Markets	0.14	0.11	0.12	0.13	0.22	0.32	0.30
% Exporting to $>$ 25 Markets	0.03	0.02	0.02	0.03	0.05	0.12	0.07

Table 2: Firm Characteristics and Market Coverage (Average 2000-2004)

Markets	Employment	Value-Added per Emp.	Exports	Sales per Market	UK Sales
1	55	49	1978	1978	1878
2	55	50	2681	1341	2191
3	106	45	5995	1998	4482
4	71	42	4771	1193	2627
5	85	48	6375	1275	3986
6-10	121	61	10979	1391	5073
11+	166	106	29095	1509	8611

Table 3: Market Coverage Premia

	Advantage in Levels		
	Ln(Employment)	Ln(VA per Emp.)	Ln(Average Wage)
Ln No. Markets	0.27***	0.10***	0.06***
	(0.02)	(0.01)	(0.01)
Observations	3138	2946	3137
R^2	0.17	0.12	0.26
	Advantage in Growth Rates		
	Employment Growth	VA per Emp. Growth	Wage Growth
Growth Market Coverage	0.01	0.05	0.03**
	(0.01)	(0.03)	(0.01)
Observations	2516	2298	2515
R^2	0.02	0.03	0.01

Notes: All regressions include sector and year dummies. Standard errors in parentheses. *** indicates significance at 1% level, ** at 5% and * at 10%.

Table 4: Firm Export Growth Across Markets, Average 2000-2004

	2001	2002	2003	2004
All Firms				
Market Entry	0.03	0.05	0.03	0.02
Market Exit	-0.04	-0.05	-0.03	-0.01
Increase Existing Markets	0.19	0.14	0.11	0.14
Decrease Existing Markets	-0.14	-0.19	-0.17	-0.09
Total Growth	0.04	-0.05	-0.06	0.05
Firms with Growing Exports				
Market Entry	0.04	0.08	0.05	0.02
Market Exit	-0.04	-0.05	-0.01	0.00
Increase Existing Markets	0.25	0.27	0.23	0.20
Decrease Existing Markets	-0.05	-0.06	-0.09	-0.05
Total Growth	0.20	0.07	0.19	0.16
Firms with Declining Exports				
Market Entry	0.03	0.01	0.01	0.02
Market Exit	-0.05	-0.05	-0.03	-0.03
Increase Existing Markets	0.09	0.05	0.04	0.05
Decrease Existing Markets	-0.27	-0.28	-0.22	-0.15
Total Growth	-0.20	-0.27	-0.20	-0.11

Table 5: Quintile Export Growth

	Growth Rates				Contribution to Total Growth			
	2001	2002	2003	2004	2001	2002	2003	2004
Quintile 0-20	0.72	0.16	0.09	0.21	0.05	-0.01	-0.01	0.02
Quintile 21-40	0.35	0.17	0.11	0.08	0.09	-0.04	-0.02	0.03
Quintile 41-60	0.14	0.03	0.04	0.09	0.09	-0.02	-0.02	0.06
Quintile 61-80	0.13	0.16	0.00	0.07	0.28	-0.27	0.00	0.16
Quintile 81-100	0.02	-0.08	-0.08	0.04	0.49	1.34	1.05	0.73
Total	0.04	-0.05	-0.06	0.05	1.00	1.00	1.00	1.00

Table 6: Average Number of Exporters, Entry and Exit by Destination

	Exporters	Entry	Exit		Exporters	Entry	Exit
UK	584	30	26	Saudi Arabia	40	9	10
USA	228	30	25	Hong Kong	36	10	10
Germany	213	26	27	Hungary	38	12	9
France	210	26	22	China	39	11	7
Netherlands	183	26	22	S. Korea	31	8	8
Italy	144	21	19	Taiwan	32	7	6
Spain	136	24	20	India	35	11	9
Belgium	139	25	24	Brazil	23	5	6
Sweden	122	19	21	New Zealand	33	10	8
Denmark	110	20	17	Malaysia	31	7	6
Portugal	76	18	18	Egypt	26	7	7
Switzerland	87	19	15	Philippines	21	5	7
Japan	75	17	17	Argentina	19	4	4
Norway	74	15	16	Kuwait	23	6	6
Canada	71	15	14	Mexico	24	8	6
Austria	69	15	14	Lebanon	17	6	7
Finland	78	16	11	Nigeria	22	6	4
Poland	61	14	11	Slovak R.	14	6	6
Australia	65	16	13	Slovenia	19	6	5
South Africa	56	15	14	Jordan	17	6	6
Greece	59	12	11	Thailand	20	6	3
Russia	43	8	10	Pakistan	17	4	3
Israel	53	11	10	Chile	15	3	4
Turkey	41	11	14	Algeria	7	2	4
Czech R.	46	13	12	Morocco	8	3	3
UAE	44	11	12	Tunisia	5	3	2
Singapore	40	11	12				

Table 7: Contributions to Net Export Growth

	2000-2001	2001-2002	2002-2003	2003-2004
Summed over Firms				
Starters	0.0007	0.0034	0.0002	0.0002
Stoppers	-0.0024	-0.0033	-0.0031	-0.0006
Increasers	0.12	0.10	0.07	0.09
Decreasers	-0.08	-0.16	-0.13	-0.04
= Net Change	0.04	-0.05	-0.06	0.05
Summed over Markets				
Starters	0.03	0.05	0.03	0.02
Stoppers	-0.04	-0.05	-0.03	-0.01
Increasers	0.19	0.14	0.11	0.14
Decreasers	-0.14	-0.19	-0.17	-0.09
= Net Change	0.04	-0.05	-0.06	0.05
	2000-2001	2000-2002	2000-2003	2000-2004
Summed over Markets				
Starters	0.03	0.07	0.06	0.07
Stoppers	-0.04	-0.09	-0.10	-0.10
Increasers	0.19	0.20	0.21	0.24
Decreasers	-0.14	-0.19	-0.25	-0.24
= Net Change	0.04	-0.01	-0.08	-0.03

Table 8: Market Transition Rates

	Market Coverage in t-1								
	0	1	2	3	4	5	6-10	11-25	>25
+4 Markets	0.00	0.00	0.01	0.01	0.02	0.01	0.03	0.05	0.04
+3 Markets	0.01	0.00	0.02	0.03	0.03	0.04	0.03	0.05	0.09
+2 Markets	0.01	0.02	0.04	0.09	0.04	0.10	0.08	0.10	0.13
+1 Market	0.15	0.08	0.14	0.13	0.16	0.18	0.16	0.11	0.09
No Change	0.83	0.81	0.56	0.46	0.38	0.33	0.28	0.17	0.27
-1 Market	0.00	0.08	0.17	0.21	0.22	0.20	0.14	0.14	0.12
-2 Markets	0.00	0.00	0.04	0.06	0.08	0.05	0.12	0.07	0.09
-3 Markets	0.00	0.00	0.00	0.00	0.02	0.02	0.05	0.08	0.02
-4 Markets	0.00	0.00	0.00	0.00	0.01	0.02	0.02	0.06	0.06
Memo items									
Total Entry	0.17	0.11	0.22	0.27	0.29	0.37	0.36	0.38	0.35
Total Exit	0.00	0.08	0.22	0.27	0.33	0.30	0.36	0.45	0.38