Exporters and Shocks*

Doireann Fitzgerald[†] Stefanie Haller[‡]

January 2012

Version 2.1 PRELIMINARY, DO NOT CITE

Abstract

Aggregate exports are not very responsive to movements in real exchange rates, though they respond strongly to trade liberalizations. We use merged plant census and customs micro data for Ireland to explore the reasons for this. We estimate the elasticity of both export participation and export sales to tariff shocks and shocks to macro variables. We control for costs by focusing on the within-firm-year effects of shocks that vary across destination markets. We find that both participation and sales respond weakly to movements in real exchange rates, but more strongly to changes in tariffs. This is consistent with a story where real exchange rate movements are perceived by firms to be less persistent than trade liberalization shocks, and there are market-specific costs of adjustment for continuing exporters as well as sunk costs of export entry.

1 Introduction

Aggregate exports are not very responsive to movements in real exchange rates. Calibrated international business cycle models typically assume a very low elasticity of substitution between home goods and foreign goods (in the range 0.5 to 1.5) in order to match comovements

*This work makes use of data from the Central Statistics Office, Ireland, which is CSO copyright. The possibility for controlled access to confidential micro data sets on the premises of the CSO is provided for in the Statistics Act 1993. The use of CSO data in this work does not imply the endorsement of the CSO in relation to the interpretation or analysis of the data. This work uses research datasets which may not exactly reproduce statistical aggregates published by the CSO. We thank the staff of the CSO for making this project possible. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. We are grateful for comments and suggestions from Costas Arkolakis, Richard Rogerson, Kim Ruhl and Jim Tybout.

[†]Department of Economics, Stanford University, Stanford, CA 94305, dfitzger at stanford.edu

[‡]Economic and Social Research Institute, Whitaker Square, Sir John Rogerson's Quay, Dublin 2, Ireland, stefanie.haller at esri.ie

of relative prices and relative quantities at a business cycle frequency. However elasticities of substitution in this low range are at odds with evidence on the response of bilateral exports to trade liberalizations. One explanation for this puzzle is that there are costs of adjustment in exporting at the plant level, such as sunk costs of entry. If the business cycle shocks that drive exchange rates are less persistent than trade liberalization shocks, the extensive margin of exports may react more to trade liberalizations than to real exchange rate movements (see Ruhl (2008)). In addition, if there are costs of adjustment for continuing exporters, the exports of continuing exporters may respond more to trade liberalizations than to real exchange rate movements (see, e.g. Drozd and Nosal (2011)).

We use 10 years of merged plant census and customs micro data for Ireland to test key aspects of these proposed explanations. We document the response of export participation and export sales at the firm-market level to both macro shocks and trade liberalization shocks. Our empirical strategy builds on the extensive literature on estimating export participation equations in the presence of sunk costs of entry, and also on a more recent literature that finds evidence of post-entry export dynamics. Crucially, we make use of the structure of our data set to focus in so far as is possible on producer responses to the relative-demand-shifting component of shocks. Our identification strategy relies on the fact that we observe exports at the level of the plant and the destination market. By comparing responses across markets within the same plant, we can clean out the first-order effect of cost changes (which may be correlated with both macro shocks and trade liberalization) on export status.

Our results suggest that participation responds to the relative-demand-switching component of real exchange rate movements in the direction one would expect: depreciations of the home currency against that of the destination market tend to induce entry and reduce exit, and vice versa for appreciations. However the size of these effects is small. Meanwhile our estimated elasticities of entry and exit with respect to tariff liberalizations are considerably greater in magnitude than those with respect to real exchange rate movements, though for some classes of firms, the estimated effects go in the opposite direction to what one might expect, with increases in tariffs being associated with increased entry and reductions in tariffs with increased exit.

On the intensive margin, we find point estimates of the firm-level elasticity of export sales with respect to movements in real exchange rates that are always below one, and not significantly different from one, consistent with low elasticities of aggregate exports with respect to movements in real exchange rates. Our point estimates of the elasticity of export sales with respect to tariff liberalizations are higher, between four and five, close to long run estimates of the elasticity of aggregate exports with respect to tariff liberalizations. However these estimates are not significantly different from zero in all samples. These results are consistent with the hypothesis that there are market-specific costs of adjustment on the intensive margin.

Moreover, we provide independent evidence consistent with market-specific costs of adjustment for continuing exporters. We find that the probability of exit is negatively related to a firm's attachment to a particular market, as measured by lagged sales in that market or number of years in the market. We also find that the growth rate of market-specific sales is negatively related to tenure in the market.

Our work is related to several literatures. It is related to an older theoretical literature which argues that the expenditure-shifting effects of exchange rate movements may depend on sunk costs of exporting at the plant level (Baldwin (1988), Baldwin and Krugman (1989) and Dixit (1989)). It is also related to several more recent papers that propose that entry and exit can (or cannot) explain facts about international real business cycles. These include Ghironi and Melitz (2005), Alessandria and Choi (2007), Atkeson and Burstein (2008) and Ruhl (2008). Methodologically, the paper builds on work by authors who have estimated both reduced form and structural dynamic discrete choice models of export supply with sunk costs of exporting (see Roberts and Tybout (1997), Bernard and Wagner (2001), Bernard and Jensen (2004) and Das, Roberts and Tybout (2007)).

Relative to this latter literature, we innovate along three dimensions. Recent evidence documents that the hazard of exit is declining in the number of years a plant participates in a market. Moreover, conditional on survival, recent entrants grow faster than incumbents (see Ruhl and Willis (2008a), Eaton, Eslava, Kugler and Tybout (2008)). The first generation of sunk cost models cannot match these facts, and several authors have recently proposed alternatives based on learning (Ruhl and Willis (2008b), Eaton, Eslava, Krizan, Kugler and Tybout (2010)), search (Chaney (2009)) and innovations to productivity (Arkolakis (2009)) which can do better at matching these facts. Related work in the macro literature which focuses on accumulation of customer base includes Foster, Haltiwanger and Syverson (2010), Gourio and Rudanko (2010) and Drozd and Nosal (2011). Motivated by this literature, our empirical approach allows for the effect of market-specific costs of adjustment on participation and sales. Second, because our data allows us to observe export status by destination market, we can identify the effect of market-specific shocks on export status. Third, our empirical strategy exploits the fact that we observe exports in multiple markets to control for the first-order effect of heterogeneity in costs on export status using plant-year fixed effects. This approach has been succesfully used in the price literature (e.g. Knetter (1989), Fitzgerald and Haller (2010)) but not so far in the literature on export entry and exit.

The effect of exchange rate shocks on entry and exit was previously addressed by Campa (2004), who uses Spanish data and finds quantitatively small effects of exchange rate movements on entry and exit. He estimates for continuing exporters an elasticity of export sales with respect to the real exchange rate that is less than one. Campa does not observe the destination breakdown of exports every year for firms in his sample, potentially affecting precision. However our findings on the real exchange rate are quite similar to his. Berman, Martin and Mayer (2011) use French data to estimate the responses of participation and sales to real exchange rates. However their results are tricky to interpret, as their empirical strategy diverges from what is the standard in the literature. While there is considerable interest in the effect of trade liberalizations on entry, exit and sales at the intensive margin, empirical work on this has been hampered by the difficulty of constructing appropriate plant- and market-specific tariff measures. An exception to this is Lileeva and Trefler (2010). However their focus is not so much on the timing of effects of liberalization on participation and sales, but on the consequent effects on productivity upgrading.

The first section of the paper describes our data. The second section describes our empirical strategy. The third section describes our results. The fourth section describes the model of exporter behavior we use to motivate our empirical strategy. The final section concludes.

2 Data

2.1 Micro data: two data sources

Our work makes use of two sources of micro data: the Irish Census of Industrial Production (CIP), and the Irish customs data. The CIP, which covers manufacturing, mining and utilities, takes place annually. Firms are required to fill in a return for all plants with 3 or more employees. In this paper, we make use of the data for the years 2000 to 2009 and for NACE Revision 1.1 sectors 10-36 (manufacturing and mining).[1] Of the variables collected in the CIP, those relevant for our purposes are the 4-digit NACE classification, country of ownership, value of sales, share of sales exported (with a rough breakdown of destinations)

[1] We are currently extending the plant census data to cover the years 2006-2009, and matching the plant data aggregated up to the firm level with customs data on exports by product and destination for the years 2000-2009.

and employment. We also construct a plant age variable based on information in the CIP and administrative records. Additionally, the CIP collects information on investment, the wage bill, expenditures on intermediates and share of intermediates imported, but so far we have not made use of this.

Our second source of data is customs records of Irish merchandise exports. These are collected monthly, but we make use of data aggregated to an annual frequency for the years 2000 to 2009. These are collected at the Combined Nomenclature (CN) 8-digit level, by destination country. These data are matched by the Central Statistics Office to the CIP data using tax id numbers, which are distinct from the plant identifier in the CIP. The match is done on the basis of firms rather than plants, as tax id numbers are associated with firms, not plants. As we document below (and as noted by the CSO), this match is imperfect.

A key feature of customs data in the EU is that data for intra-European and extra-European trade are collected separately, using two different systems called Intrastat and Extrastat. Most importantly, for Ireland, the reporting threshold for intra-European trade (635,000 Euro per year) is much higher than the reporting threshold for extra-European trade (254 Euro per transaction). We have reason to believe that a substantial fraction of smaller exporters do not report intra-European exports to Intrastat, and have intra-European exports imputed through VAT returns. For these flows, we do not know the destination within the EU.

We make use of the data at the firm level, but in our baseline analysis, we focus only on firms that have only 1 plant in Ireland.[2] We drop firms that have a zero value for total sales or the number of employees in more than half of their years in the sample. We also drop firms if more than half of their observations were estimated or imputed by the Central Statistics Office due to non-response or incomplete returns. This affects small firms more than big firms. Further details on the data and how we have cleaned it are provided in the data appendix.

2.2 Summary statistics on the micro data

Because the Census of Industrial Production asks questions on the export behavior of respondents, and because we have access to customs data, we have two independent sources of information on the export behavior of firms. Both sources are subject to error.

Since the CIP is a survey, misreporting and non-reporting of exports by firms is possible.

[2]This accounts for about 75% of plants. In future work, we plan to relax this restriction.

For nonreporters, the CSO imputes exports in some cases, and we have tried in so far as possible to exclude these firms from our sample. As regards the volume of exports, the CIP asks firms to report the value of sales, and the share of sales that is exported. Given the frequency with which round numbers are reported (e.g. 50%, 25%) we have reason to believe that this variable is often measured with considerable error (though this may average out across firms).

For customs data, the matching process introduces errors, in that a substantial fraction of exports reported to customs cannot be matched back to CIP firms. Some of this may be attributed to exports of industrial products by non-CIP firms (e.g. firms classified as being in the services sector). Some may be due to the use of intermediaries. Some may be due just to failure to find a match when a match should be found. An additional potential source of error is introduced by imputation of low-value intra-EU flows through VAT returns.

We believe that the CIP measurement of export behavior is likely to be more reliable on the participation margin than on the intensive margin. On the other hand, the measurement of export behavior based on the match with customs data is more likely to be reliable on the intensive margin (by destination) than on the extensive margin.

Table 1 provides the first set of summary statistics on the quality of the match. The first column reports the value of total industrial exports over the sample period based on publicly available customs data. The second column reports total exports by our cleaned dataset of single-plant CIP firms. This is calculated by multiplying the reported export share by reported sales. Clearly, there is a discrepancy between the two measures - despite the fact that we sum exports only over single-plant firms, in several years, the CIP-reported exports are greater than total industrial exports as measured by customs reports. The third column reports the total CIP-reported exports of firms that are matched to positive customs exports. This total is less than total CIP-reported exports for the single-plant firm sample. The final column reports total customs exports in plants that are matched to positive customs exports. The total here is generally (though not always) less than the CIP-reported exports by the same set of firms. It is not possible to conclude from these figures that one source of data on export behavior is necessarily better than the other.

Table 2 reports another dimension of match quality - the number of exporters in the single-plant firm sample according to the CIP definition and the customs match definition. The first column reports the total number of firms in our single-plant firm sample. The second column reports the number of these firms that report positive exports in the CIP. The third column reports the number of these firms that are matched to positive exports from customs data. The discrepancy between the total number of exporters according to the two different definitions is not huge. However, the fourth column reports the overlap between the two definitions (i.e. firms reported as being exporters under both CIP and customs definitions), and it is far from perfect. Tthe firms classified as exporters under one definition but not the other are in general very small.

Table 3 reports a key dimension of the two different measures of exporting - the ability to observe export particpation by market. In the CIP, firms are asked to report the share of export sales destined for the UK and the share destined for the US. This allows us to compare the measurement of participation by market across the two exporter definitions for one Intrastat destination and one Extrastat destination. There is a very substantial gap between the number of firms reporting participation in the UK market under the two definitions. This is most likely accounted for by the fact that the UK is an Intrastat destination, and as such, the threshold for reporting exports to customs is much higher than the threshold for reporting exports to the US. Although there may be error in the CIP measure of participation in the UK market, this discrepancy strongly suggests to us that the customs definition may under-measure participation in Intrastat destinations. In contrast, the number of firms reporting exports to the US is roughly similar across the two definitions, but the overlap between the two definitions is quite low.

Finally, Table 4 reports summary statistics on employment, sales, firm age and export shares for exporters as classified by the CIP definition and by the customs definition, and for non-exporters as classified by the CIP definition. By these definitions, exporters are bigger and older than non-exporters, exactly as the past literature finds. The main discrepancy between the two exporter definitions is that export shares (as measured by matched customs exports over CIP-reported sales) are much lower than CIP-reported export shares.

We also have information on entry and exit under both definitions of export participation, which will be reported in later versions.

2.3 Tariff data

In identifying the effect of changes in tariffs on exports, we exploit the fact that our sample covers the period 2000-2004, which saw the last years of the implementation of tariff reductions mandated under the Uruguay Round. This affects exports from Ireland to the US, Ireland's single biggest trading partner over the sample period. It also affects a series of other destinations which account for a smaller fraction of Irish exports. In contrast, exports to EFTA countries and to the EU were not subject to tariffs at any time during the sample period. Because we observe exports at the level of the product, firm, destination and year, we can exploit variation in average tariffs across firms, destinations and years to identify export responses to changes in tariffs.

Our source for tariff data is the WTO. We collect tariff data for four destinations, the US, Japan, Canada and Australia. Other destinations are more marginal for Irish exporters, and the tariff data provided by the WTO for these destinations is often incomplete. The US accounts for on average 20% of Irish merchandise exports over this period, and the remaining destinations account for between 3 and 6% of merchandise exports. These countries differ in the sectoral composition, magnitude and timing of tariff changes over the sample period.

Although we have exports at the 8-digit CN level which is the tariff line for the EU, this does not correspond to the tariff line for any of the target markets. However the 6-digit CN is exactly the 6-digit HS, so we collect average tariffs at the 6-digit HS level, and match these to the 6-digit CN.

We then construct two measures of tariffs at the firm-market level for each year. The first is a measure appropriate to the intensive margin. It is constructed as the sales-weighted average of $(1 + \tau)$ across HS-6 categories for which the firm reports positive exports to the relevant destination by the firm, where τ is the tariff reported as a fraction, and sales are sales by HS-6 category to that particular destination in that year. The second is a measure appropriate to the extensive margin. Here, we are hampered by the lack of production data at a level more disaggregated than the 4-digit NACE. As a first pass, we construct a weighted average of $(1 + \tau)$ across all HS-6 categories for which the firm reports positive exports to *any* destination. The weights are total sales by HS-6 category in that firm in that year. The disadvantage of this measure is that it can be constructed only for firms that export *somewhere.*

Summary statistics on variation in tariff data are reported in Table 5, which reports the results of regressing HS6 unweighted average tariffs on time fixed effects and HS6 fixed effects, for each country separately. This table illustrates the fact that tariffs vary across countries, that reductions are concentrated over the period 2001-2004, but that the size of reductions and the time-pattern of reductions differs across countries.

2.4 Macro variables

The macro variables we include in our regressions are the real consumption exchange rate between Ireland and the target market, and a measure of real local currency demand in the target market. Real exchange rates are constructed using data on annual average nominal exchange rates and CPIs from International Financial Statistics. Real demand in the target market is calculated as GDP in current domestic currency less exports in current domestic currency plus imports in current domestic currency, all deflated by the relevant CPI. The National Accounts data are taken from the OECD, while the CPIs come from International Financial Statistics. We do not cover all countries in for which exports are reported, only the 20 most important countries.

3 Empirical strategy

We now describe the strategy we use to investigate the responsiveness of export participation and export revenue to different types of shocks. We take a reduced form approach, motivated by the model sketched in section 5 of the paper. There are three key differences between our strategy and that used by the previous literature. First, we allow for the possibility that post-entry, firms may slowly accumulate market-specific experience (either passively or through active investments) rather than jumping straight to steady state size. Second, we exploit the fact that we observe export status for multiple markets for each firm to control for the first order effect of costs using fixed effects. Third, our focus is on responses to shocks rather than the steady state.

3.1 Participation

The standard export participation equation regresses a participation indicator on plant or firm fixed effects, year fixed effects, an indicator of lagged participation, and a vector of (lagged) variables that are intended to capture time variation in costs. We instead estimate separate entry and exit equations. This allows for sunk entry costs while relaxing degrees-of-freedom constraints which would otherwise be an issue for us. The probability of entry or exit is then allowed to depend on market fixed effects, costs (assumed common across markets within a firm), firm-market-specific experience, and "shocks.". We also allow the sensitivity of entry and exit to shocks to vary across firms with their costs and market-specific experience.

Let i index firms, let k index markets and let t index time. Let $X_t^{ik} \in \{0,1\}$ be a participation indicator. We approximate the probability of entry as follows:

$$\Pr\left[X_t^{ik}=1|X_{t-1}^{ik}=0\right]=G\left(\alpha^k+c_t^i+\boldsymbol{\beta}\mathbf{z}_t^{ik}+\boldsymbol{\gamma}'\left(\mathbf{s}_{t-1}^i\otimes\boldsymbol{z}_t^{ik}\right)+\varepsilon_t^{ik}\right) \tag{1}$$

α^k is a time-invariant market-specific effect which captures time-invariant components of trade costs and all time-invariant factors which lead the probability of entry for all firms to be greater in some markets than others. It also accounts for scaling of the macro variables in the shock vector $\mathbf{z}_t^{ik}$. c_t^i is a firm-year effect which captures the first-order effect of firm-year-specific marginal cost, as well as any other variables that are common across markets for a given firm at a given point in time. A more parametric alternative would be to include $\mathbf{s}_{t-1}^i$, a vector of correlates of costs (lagged one year due to simultaneity concerns). $\mathbf{z}_t^{ik}$ is a vector of macro variables and trade policy variables. This variable is interacted with $\mathbf{s}_{t-1}^i$, a vector of correlates of costs, to allow the sensitivity of entry to shocks to vary across plants with different costs. In the baseline specification, we do not allow the probability of entry or the sensitivity of entry to shocks to vary with past market experience. ε_t^{ik} captures variables that are idiosyncratic to the firm, market and year (e.g. idiosyncratic demand shocks).

We approximate the probability of exit as follows:

$$\Pr\left[X_t^{ik} = 0 | X_{t-1}^{ik} = 1\right] = G\begin{pmatrix} \alpha^k + c_t^i + \boldsymbol{\lambda}'\mathbf{d}_{t-1}^{ik} + \boldsymbol{\beta}\mathbf{z}_t^{ik} + \\ \boldsymbol{\gamma}'\left(\mathbf{s}_{t-1}^i \otimes \boldsymbol{z}_t^{ik}\right) + \boldsymbol{\rho}'\left(\mathbf{d}_{t-1}^{ik} \otimes \mathbf{z}_t^{ik}\right) + \varepsilon_t^{ik} \end{pmatrix} \tag{2}$$

In addition to the variables already described, $\mathbf{d}_{t-1}^{ik}$ is a vector of variables intended to capture firm i's experience in market k at time $t-1$ (e.g. lagged number of years in the market, lagged sales in the market). The vector of shocks is interacted with the market experience vector to allow the sensitivity of exit to shocks to vary across firms with different market-specific experience as well as across firms with different costs.

Since we observe participation status for multiple export markets, there are two potential approaches to controlling for c_t^i: fixed effects or random effects. Using the fixed effects approach, we identify the coefficients on the shock variables $\mathbf{z}_t^{ik}$ solely from within-firm-year variation in the entry or exit decision. This approach is appealing, but it does restrict our choice of the $G(\cdot)$ function. As a baseline, we estimate a linear probability model (linear $G(\cdot)$), which allows us to make use of all firm-market-years on which a full set of independent variables is available. This has all the usual problems that using a linear probability model entails. Note that while in the standard case, estimating a fixed effects model with a lagged dependent variable or functions of lags of the dependent variable (the variables in $\mathbf{d}_{t-1}^{ik}$) is problematic, in our case, the structure of the fixed effects is such that the usual bias does not apply.

We also experiment with a conditional logit for the $G(\cdot)$ function, where the estimated coefficients do not depend on the fixed effects (which are not actually estimated). This has the disadvantage that only cases where entry or exit is observed in some markets but not others are used to identify the parameters of interest. This restricts the size of the sample, discarding information that can be used to identify the parameters of interest and potentially reducing precision.

We could also adopt a random effects approach, estimating c^i as a random effect, and include $\mathbf{s}_{t-1}^i$ in the equation in levels, to capture time variation in costs. This would allow us to use a probit for the $G(\cdot)$ function, while making use of all firm-market-years for which a full set of independent variables is available. The initial conditions problem would be is present in this case for the exit equation, where functions of lags of the dependent variable are included. We have not so far implemented this approach.

For the vector $\mathbf{s}_{t-1}^i$ we have experimented with including indicators for firm size (measured by number of employees), plant age, foreign ownership, and the capital-labor ratio. The indicator approach allows for nonlinear dependence of entry thresholds on the underlying continuous variables, though it is costly in terms of degrees of freedom. In the baseline results, we restrict the vector to a parsimonious set of indicators for firm size.

For the vector $\mathbf{d}_{t-1}^{ik}$ we have experimented with including indicators for the lagged number of years in the market and the log of lagged sales in the market. Number of years in the market must be calculated using the 10 years of data available, so using indicators for this variable can restrict the sample size. Lagged sales are included as the log of Euro revenues from the relevant market, deflated by the Irish CPI.

The vector $\mathbf{z}_t^{ik}$ includes the log of the real consumption exchange rate between Ireland and market k, a measure of real aggregate demand in market k, and firm-market-year-specific tariffs.

Because the reporting threshold for data to be collected through Intrastat is much higher than that for data collected through Extrastat, we separately estimate (1) and (2) on a sample that consists only of exports to Extrastat countries, and on a broader sample that includes exports to Intrastat as well as Extrastat countries. We exclude exports to countries (i.e. accession countries in Eastern Europe and elsewhere) whose status changed over the sample period.

3.2 Sales

We estimate the sales equation in differences rather than in levels. We approximate the change in log sales of firm i in market k at time t as follows:

$$\Delta r_t^{ik} = \alpha^k + c_t^i + \boldsymbol{\beta}' \Delta \mathbf{z}_t^{ik} + \boldsymbol{\gamma}' \mathbf{a}_{t-1}^{ik} + \eta_t^{ik} \tag{3}$$

r_t^{ik} is the log of sales measured in Euros.[3] As in the case of entry and exit, α^k is a time-invariant market-specific effect which captures time-invariant reasons why sales growth might be higher for all firms in some markets than others. c_t^i is a firm-year effect which captures changes in costs (assumed to be the same across markets within a firm) as well as demand factors where changes are common across markets for a given firm at a given point in time. The vector $\mathbf{z}_t^{ik}$ is defined as above, and is included in differences. $\mathbf{a}_{t-1}^{ik}$is a vector of indicators for the (lagged) number of years the plant has been in market k. Inclusion of this vector allows the rate of growth of sales to differ with the number of years the plant has been in the market (i.e. post-entry dynamics). η_t^{ik} captures changes in variables that are idiosyncratic to the firm, market and year (e.g. idiosyncratic demand shocks).

A major issue in estimating (3) is selection. We only observe sales for firm-market-years where both $X_t^{ik} = 1$ and $X_{t-1}^{ik} = 1$. But participation depends on unobserved idiosyncratic shocks, which also show up in the sales equation. Continued participation for a firm-market-year that experiences a negative idiosyncratic shock is more likely if the observed shocks $\Delta \mathbf{z}_t^{ik}$ are favorable. This implies that the expectation of η_t^{ik} conditional on the independent variables is not equal to zero for firm-market-years close to the participation thresholds. There are several aspects of the setup that make the standard approaches to controlling for selection (such as a Heckman correction) inappropriate or tricky to implement. We adopt the following approach as a baseline. We restrict attention to firm-market pairs where we observe continuous presence throughout the sample. The underlying assumption is that these pairs are sufficiently far from the exit threshold that we observe the full distribution of idiosyncratic shocks η_t^{ik}. For this sample, there is no measurable heterogeneity in $\mathbf{a}_{t-1}^{ik}$, so we drop this vector from the set of independent variables. Under all specifications, we calculate robust standard errors.

Given the requirement that firm-market pairs appear in all sample years, the distinction between Intrastat and Extrastat destinations is less likely to be problematic for the intensive margin than the extensive margin. But for comparability, we estimate (3) on the same two

[3] We deflate by the Irish CPI, but this is irrelevant due to the inclusion of firm-year fixed effects in the estimating equation.

samples as in the case of (1) and (2).

4 Results

We first present the results for entry, then the results for exit, and then the results on sales. We then discuss the economic significance of our findings.

4.1 Entry

The results from our baseline specification of the entry equation are reported in Table 6. The first column reports the results using only exports to a limited sample of Extrastat countries. The second column reports the results using exports to a broader sample of both Intrastat and Extrastat destinations. Along many dimensions, the results are similar across both samples. Entry responds positively to depreciations of the real exchange rate between Ireland and potential export destinations. It also responds positively to increases in real demand in a potential market. In the Extrastat sample, these effects are muted for larger firms (which are presumably closer to the entry threshold than smaller firms). In the full sample, the comparative static goes in the opposite direction, but is quantitatively much smaller. For the smallest firms, entry is increasing in tariffs, though this effect is entirely negated (or reversed in the larger sample) for bigger firms. Finally, the elasticity of the entry response to tariffs is an order of magnitude greater than the elasticity with respect to macro shocks.

We also examine robustness to estimating a conditional logit. Table 7 reports the results. The sample size is greatly reduced in this case. The result that point estimates of the elasticity of entry with respect to tariffs is greater than point estimates of elasticities with respect to macro shocks is unchanged.

4.2 Exit

The results from our baseline specification of the exit equation are reported in Table 8. The sample size is smaller here than in the case of entry, because the number of potential exiters is limited to those currently participating, which for most destinations, is a minority of firms. The first column reports the results using only exports to a limited sample of Extrastat countries. The second column reports the results using exports to a broader sample of both Intrastat and Extrastat destinations. Along many dimensions, the results are similar across both samples. Exit responds positively to appreciations of the real exchange rate between Ireland and potential export destinations. Exit responds positively to increases in real demand in a potential market (so churn in general is increasing in real demand in target markets). These effects do not vary significantly with firm size. For the smallest firms, exit is decreasing in tariffs, though this effect is reversed for bigger firms. Finally, the elasticity of the exit response to tariffs is greater than the elasticity with respect to macro shocks.

We also examine robustness to estimating a conditional logit. Table 9 reports the results. They are qualitatively similar to those in the linear probability baseline.

4.3 Sales

The results from estimating our baseline specification of the sales growth equation are reported in Table 10. Our point estimates of the elasticity of sales with respect to real exchange rates are systematically below 1. They are significantly different from zero (and not significantly different from 1) in the broader sample that includes exports to both Intrastat and Extrastat destinations. This is consistent with elasticities that are estimated on aggregate exports. Consistent with what is found in aggregate data, we find higher point estimates of elasticities with respect to foreign real demand than with respect to prices.

In contrast, our point estimates of the elasticity of sales with respect to tariffs are consistently higher in absolute value - in the neighborhood of 4 to 5. Again, these estimates are consistent with those estimated using more aggregate data (i.e. not at the firm level). However this elasticity is relatively imprecisely estimated, and is significantly different from zero only in the largest sample in which we estimate it.

5 Model

In line with recent evidence on exporter dynamics, we extend the standard model of sunk costs of exporting to allow for costs of adjustment on market-specific quantities for continuing exporters. We frame the friction as costly accumulation of market-specific demand, though there are alternative ways to model adjustment costs that would yield similar empirical implications. Several authors (Arkolakis (2009), Chaney (2010), Eaton et al (2010), Ruhl and Willis (2008)) have recently derived dynamic exporter behavior from first principles. For simplicity, we take a reduced-form approach.

The model has the following features. We assume that there are two distinct drivers of a plant's demand in a particular market. On the one hand, demand depends on the plant's own price, the price of its competitors, and the level of expenditure in the target economy. On the other hand, demand depends on the plant's customer base in the target economy. For simplicity, customer base is assumed to evolve independently of prices. Under the baseline scenario, plants invest today in future customer base that generates demand through a decreasing returns technology. Decreasing returns imply that there is a steady state level of customer base (conditional on market participation) that depends on plant characteristics and the aggregate state. We allow for convex adjustment costs that slow down convergence to steady state; it is cheaper to build up customer base gradually rather than doing it all at once. The model is related to those of Arkolakis (2008) in the trade literature, and Drozd and Nosal (2011), Foster, Haltiwanger and Syverson (2010) and Gourio and Rudanko (2010) in the macro literature.

For simplicity, we assume that the only link between markets for a given plant is through marginal costs (marginal cost is assumed identical across all markets served by the plant). We assume that there are no other spillovers across markets, in the sense that presence in one market does not affect latent initial demand in another market or the fixed cost of participating in another market.

5.1 Demand

Demand faced by plant i in market k at time t is as follows:

$$Q_t^{ik} = q\left(P_t^{ik*}, P_t^{k*}, Q_t^k\right)\left(D_t^{ik}\right)^{\alpha}\exp\left(\eta_t^{ik}\right) \tag{4}$$

where $\alpha \in (0,1)$. The first term of this expression is a function, $q(\cdot)$, of Q_t^k, aggregate expenditure in market k, P_t^{k*}, the market price in market k expressed in foreign currency, and P_t^{ik*}, plant i's price in market k, expressed in foreign currency. η_t^{ik} is an iid lognormally distributed random variable that captures idiosyncratic shocks to demand. D_t^{ik} is a persistent demand shifter that we will refer to as "customer capital." It captures the fact that there may be slow-moving determinants of the level of demand unrelated to the time-t price. At time t, D_t^{ik} is predetermined. D_t^{ik} accumulates according to the law of motion:

$$D_t^{ik} = \begin{cases} (1-\delta)\, D_{t-1}^{ik} + I_{t-1}^{ik} & \text{if } X_{t-1}^{ik} = 1 \\ \underline{D}^{ik} & \text{if } X_{t-1}^{ik} = 0 \end{cases} \tag{5}$$

where δ is the rate of depreciation of customer capital, X_t^{ik} is an indicator variable for i's participation in market k and $I_t^{ik} \geq 0$ is i's investment in customer capital. We assume that for all plants that produce and sell something to the home market, $\underline{D}^{ik} \leq \bar{D}_t^{ik}$, where $\bar{D}_t^{ik}$ is steady state D, which depends on plant characteristics and the aggregate state, but not on η_t^{ik}.[4] Notice that exit is assumed to imply full depreciation of customer capital in the sense that irrespective of what was accumulated prior to exit, on re-entry, customer capital will be reset to $\underline{D}^{ik}$.[5]

With this formulation, the choice of price today is a static decision, as it does not affect the future value of D. This simplifies the analysis considerably, though possibly at the expense of realism.[6]

This model nests an alternative where D_t^{ik} evolves independently of all actions by the plant other than its participation history. For example, demand may grow through an exogenous process conditional on participation as in Ruhl and Willis (2008a). In this case, $I_t^{ik} = 0$.

5.2 Costs

We assume that plant i faces marginal cost $\tau_t^k \left(W_t / z_t^i\right)$ of serving market k. This cost is expressed in terms of domestic currency. The first term, τ_t^k, includes destination-specific and potentially time-varying tariffs and transportation costs (in our empirical analysis, we assume transportation costs are constant over the sample period). The second term, W_t/z_t^i, is the ratio of W_t, the cost of the input bundle, assumed the same for all plants, and z_t^i, plant i's idiosyncratic productivity. This term may vary across plants and over time but does not vary across markets for given i and t. There is also a fixed cost $W_t F^k$ of participating in market k in any period. Because of this cost, some plants will prefer not to participate in the export market. We do not allow for a sunk cost of entry, because as will become clear

[4]We can guarantee that this is the case if $\underline{D}^{ik}$ is sufficiently low and the fixed costs of selling in the dometic market are sufficiently high.

[5]Instead of assuming an initial draw, we could assume that plants must invest in D prior to entry. We have not yet derived the implications of varying this assumption. In addition we could potentially allow for a less stark assumption of a higher depreciation rate $\delta^H > \delta$ for plants that do not sell in a market. In the light of the evidence from the previous literature that spells of exporting previous to date $t-1$ do not greatly increase the probability of exporting at date t conditional on not exporting at $t-1$, we have not yet explored this possibility.

[6]Motivated by price evidence for relatively homogeneous goods, Foster, Haltiwanger and Syverson (2010) present a model of demand accumulation where the choice of price is a dynamic decision. Gourio and Rudanko (2010) instead assume that producers offer the same price to new customers and old customers, but make transfers to new customers to induce them to buy. Our investments I_t^{ik} could possibly be interpreted as transfers of this type, though we have not explored this interpretation to date.

later, the model is observationally equivalent to one with a sunk cost of entry in addition to the structure already described.

We assume that in order to increment consumer capital in market k by amount I_t^{ik}, the plant must spend an amount given by $W_t\left[I_t^{ik}+\phi\left(I_t^{ik}-\delta D_{t-1}^{ik}\right)\right]$. The adjustment cost function is assumed to have the following properties: $\phi(x)=0$ if $x\leq 0$, while if $x>0$, $\phi(x)>0$, $\phi'(x)>0$, $\phi''(x)>0$. The convex cost of adjustment implies that under constant market conditions, plants do not jump straight to their steady state customer capital.

Note that investment and the fixed participation cost depend on the home currency price of the domestic input bundle, but not the price of foreign inputs. This assumption could be relaxed.

5.3 Static optimization

Real flow profits from market k for a plant that sells a positive quantity are given by:

$$\Pi_t^{ik}=\frac{E_t^k P_t^{ik*}}{P_t^i}Q_t^{ik}-\tau_t^k\frac{W_t}{P_t^i z_t^i}Q_t^{ik}-\frac{W_t}{P_t^i}\left[F^k+I_t^{ik}+\phi\left(I_t^{ik}-\delta D_{t-1}^{ik}\right)\right] \tag{6}$$

The optimal price can be expressed as:

$$P_t^{ik*}=\frac{\theta_t^{ik}}{\theta_t^{ik}-1}\frac{\tau_t^k W_t}{E_t^k z_t^i} \tag{7}$$

where

$$\theta_t^{ik}=\frac{\partial\ln q\left(P_t^{ik*},P_t^{k*},Q_t^k\right)}{\partial\ln P_t^{ik*}} \tag{8}$$

For general specifications of the $q(\cdot)$ function, θ_t^{ik} is a function of P_t^{ik*}, P_t^{k*} and Q_t^k. We assume that the optimal price can be approximated as follows:

$$P_t^{ik*}=\mu_t^{ik}\frac{\tau_t^k W_t}{E_t^k z_t^i} \tag{9}$$

where

$$\mu_t^{ik}=\frac{\theta_t^{ik}}{\theta_t^{ik}-1}=\mu\left(\frac{\tau_t^k W_t}{E_t^k z_t^i},P_t^{k*},Q_t^k\right) \tag{10}$$

is the gross markup.

Plant i's real revenues from market k expressed in home currency can then be written

$$R_t^{ik}=\frac{1}{P_t^i}\mu_t^{ik}\frac{\tau_t^k W_t}{z_t^i}q\left(\mu_t^{ik}\frac{\tau_t^k W_t}{z_t^i E_t^k},P_t^{k*},Q_t^k\right)\left(D_t^{ik}\right)^{\alpha}\exp\left(\eta_t^{ik}\right) \tag{11}$$

and conditional on participation, real flow profits net of costs and investment in future customer capital can be written

$$\Pi_t^{ik} = R_t^{ik}\left(\frac{\tau_t^k W_t}{z_t^i}, E_t^k, P_t^{k*}, Q_t^k, D_t^{ik}, \eta_t^{ik}\right) - \frac{W_t}{P_t^i}\left[F^k + I_t^{ik} + \phi\left(I_t^{ik} - \delta D_{t-1}^{ik}\right)\right] \qquad (12)$$

Notice that D_t^{ik} enters the revenue function in the same way as productivity z_t^i. In this sense, our model of market-specific demand accumulation is indistinguishable from one where participation in a market allows plants to engage in market-specific technology upgrading. However market-specific technology seems less likely to depreciate on exit than customer base, which will have implications for the participation decision.

Note also that in the version of the model where customer base accumulates through a process that depends only on participation, the revenue function is unchanged, but the term $I_t^{ik} + \phi\left(I_t^{ik} - \delta D_{t-1}^{ik}\right)$ drops out of the flow profit function. Since I_t^{ik} is no longer a choice variable, this will affect the dependence of the growth rate of revenue on lagged shocks.

5.4 Dynamic optimization

As is standard in the empirical literature on export entry and exit, we ignore the plant existence decision, instead conditioning on positive lagged sales in the home market, assuming that this is the easiest market to enter.[7] We then focus on the decision to participate or not in a particular export market. We assume that the plant observes η_t^{ik}, z_t^i, E_t^k, W_t, τ_t^k, P_t^{k*}, P_t^i and Q_t^k before making its decision. Let Θ_t^k denote the aggregate state $\left\{E_t^k, W_t, Q_t^k, P_t^{k*}, P_t^i, \tau_t^k\right\}$. If plant i participated in market k at $t-1$, it inherits a predetermined D_t^{ik} from the previous period. Otherwise it reverts to its initial draw $\underline{D}^{ik}$. Let X_t^{ik} be an indicator variable, equal to 1 if the plant participates in market k at time t, equal to zero otherwise. The value of market k to firm i is:

$$V\left(D_t^{ik}, \eta_t^{ik}, z_t^i, \Theta_t^k\right) = \max_{\substack{X_t^{ik} \in \{0,1\} \\ I_t^{ik} \geq 0}} \left\{X_t^{ik}\Pi_t^{ik} + \beta \mathbb{E} V\left(D_{t+1}^{ik}, \eta_{t+1}^{ik}, z_{t+1}^i, \Theta_{t+1}^k\right)\right\}$$

with

$$D_{t+1}^{ik} = X_t^{ik}\left[(1-\delta) D_t^{ik} + I_t^{ik}\right] + \left(1 - X_t^{ik}\right)\underline{D}^{ik}$$

[7] We thus ignore entry of plants that are born global and entry and exit of plants that sell only to the foreign market.

The solution to this problem yields two policy functions:

$$X_t^{ik} = X\left(D_t^{ik}, \eta_t^{ik}, z_t^i, \Theta_t^k\right) \tag{13}$$

$$I_t^{ik} = I\left(D_t^{ik}, \eta_t^{ik}, z_t^i, \Theta_t^k\right) \tag{14}$$

Our interest is in the properties of $X_t^{ik}\left(D_t^{ik}, \eta_t^{ik}, z_t^i, \Theta_t^k\right)$, since we observe the participation decision. It is straightforward to show that X_t^{ik} is increasing in D_t^{ik} (see Appendix for proof). This implies that there is an underlying asymmetry in the participation decision that arises out of the accumulation of customer capital. This capital acts like a sunk cost, whose size varies across otherwise identical plants with the length of time they have been in market k.

5.5 Small exporters, exit hazard and sales growth conditional on survival

The model we have just laid out can explain why we observe exporters who export small amounts. They can do so because there is no up-front sunk cost of entry. At the same time, previous participation predicts future participation, because participation allows plants to accumulate a form of market-specific capital. The model can also generate a hazard of exit that is decreasing in the length of time a plant has been in a market. This follows from the fact a recent entrant will have a lower D than an otherwise identical plant that has been in the market for some time. Hence the recent entrant will be more vulnerable to idiosyncratic demand shocks. This model is also able to match decreasing growth rates conditional on survival, as the marginal product of customer capital and hence investment in customer capital and increases in demand decline as plants approach their steady state customer capital.

5.6 Costs of adjustment and comparative statics on entry and exit

In the standard model with a sunk cost of export entry, current participation in a market depends on lagged participation. In the model with slow accumulation of demand conditional on participation, current participation depends on D_t^{ik}. Where accumulation is an active process, D_t^{ik} depends on the history of costs and aggregate shocks as well as the length of time the plant has been in the market. Where accumulation is a passive process that is identical for all plants, as in Ruhl and Willis (2008a), D_t^{ik} depends only on the length of time the plant has been in the market.

5.7 Costs of adjustment and comparative statics on revenues

In the standard model with a sunk cost of export entry, the growth rate of revenues conditional on participation is affected by the sunk cost only through selection, as sufficiently low (or negative) growth rates trigger exit. In the model with slow active accumulation of demand, conditional on participation, the growth rate of revenues between period t and period $t-1$ depends on the difference between $D_{t-1}^{ik}$and expected steady state D at date t conditional on the state of the world at $t-1$. If accumulation is passive as in Ruhl and Willis (2008a), there is a deterministic component to revenue growth that depends on how long a plant has been in a market.

6 Conclusion

We document the response of export participation and export sales at the firm-market level to both macro shocks and trade liberalization shocks. We find that both participation and sales of continuing exporters are more responsive to tariff reductions than they are to macro shocks, in particular, movements in real exchange rates. Our results are consistent with a story where producers perceive macro shocks as being less persistent than trade liberalization shocks, and there are both sunk costs of entry, and costs of adjusting sales for continuing exporters. Moreover, we provide independent evidence consistent with market-specific costs of adjustment for continuing exporters. We find that the probability of exit is negatively related to a firm's attachment to a particular market, as measured by lagged sales in that market or number of years in the market. We also find that the growth rate of market-specific sales is negatively related to tenure in the market.

Our results provide support for recent papers by Ruhl (2008) and Drozd and Nosal (2011) which suggest that costs of adjustment for exporters may play a role in explaining sluggish responses of aggregate exports to real exchange movements. At the same time, they are consistent with the findings of the literature of substantial responses of trade aggregates to trade liberalizations. While further analysis is clearly merited - in particular with the goal of understanding why producers respond sluggishly to real exchange rates though univariat analysis finds them to be very persistent - we think that this is important progress towards resolving the international "elasticity puzzle."

References

[1] Alessandria, G. and H. Choi (2007), "Do Sunk Costs of Exporting Matter For Net Export Dynamics?" *Quarterly Journal of Economics*, 122 (1).

[2] Arkolakis, C. (2008), "Market Penetration Costs and the New Consumers Margin in International Trade," mimeo.

[3] Arkolakis, C. (2009), "A Unified Theory of Firm Selection and Growth," mimeo.

[4] Atkeson, A. and A. Burstein (2008), "Pricing to Market, Trade Costs and International Relative Prices," *American Economic Review*.

[5] Baldwin, R. (1988), "Hysteresis in Import Prices: The Beachhead Effect," *American Economic Review* 78 (4), 773-85.

[6] Baldwin, R. and P. Krugman (1989), "Persistent Trade Effects of Large Exchange Rate Shocks," *Quarterly Journal of Economics* 104 (4), 635-54.

[7] Berman, N., P. Martin and T. Mayer (2009), "How do Different Exporters React to Exchange Rate Changes? Theory, empirics and aggregate implications," mimeo.

[8] Bernard, A. and B. Jensen (2004), "Why Some Firms Export," *Review of Economics and Statistics* 86 (2), 561-569.

[9] Bernard, A., S. Redding, B. Jensen and P. Schott (2007), "Firms in International Trade," *Journal of Economic Perspectives*, Summer.

[10] Bernard, A. and J. Wagner (2001), "Exports and Success in German Manufacturing," *Weltwirtschaftliches Archiv* 133 (1), 134-157.

[11] Campa, J. (2004), "Exchange Rates and Trade: How Important is Hysteresis in Trade?" *European Economic Review* 48, 527-548.

[12] Chaney, T. (2010), "The Network Structure of International Trade," mimeo.

[13] Das, M., M. Roberts and J. Tybout (2007), "Market Entry Costs, Producer Heterogeneity and Export Dynamics," *Econometrica* 75 (3), 837-873.

[14] Dixit, A. (1989) "Entry and Exit Decisions Under Uncertainty," *Journal of Political Economy* 97 (3), 620-38.

[15] Drozd, L. and J. Nosal (2011), "Understanding International Prices: Customers as Capital," mimeo.

[16] Eaton, J., M. Eslava, C. Krizan, M. Kugler and J. Tybout (2010), "A Search and Learning Model of Export Dynamics," mimeo.

[17] Eaton, J., M. Eslava, M. Kugler and J. Tybout (2008), "The Margins of Entry into Export Markets: Evidence From Colombia," in E. Helpman, D. Marin, and T. Verdier, eds., *The Organization of Firms in a Global Economy,* Cambridge, MA: Harvard University Press.

[18] Fitzgerald, D. and S. Haller (2010), "Pricing-to-Market: Evidence From Producer Prices," mimeo.

[19] Foster, L., J. Haltiwanger and C. Syverson (2010), "The Slow Growth of New Plants: Learning About Demand," mimeo.

[20] Ghironi, F. and M. Melitz (2005), "International Trade and Macroeconomic Dynamics With Heterogeneous Firms," *Quarterly Journal of Economics* 120 (3), 865-915.

[21] Gourio, F. and L. Rudanko (2010), "Customer Capital," mimeo.

[22] Knetter, M. (1989), "Price Discrimination by U.S. and German Exporters," *American Economic Review* 79 (1), 198-210.

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

[24] Roberts, M. and J. Tybout (1997), "The Decision to Export in Colombia: An Empirical Model of Entry With Sunk Costs," *American Economic Review* 87 (4), 545-564.

[25] Ruhl, K. (2008), "The International Elasticity Puzzle," mimeo.

[26] Ruhl, K. and J. Willis (2008a), "New Exporter Dynamics," mimeo.

[27] Ruhl, K. and J. Willis (2008b), "Learning About New Exporters," mimeo.

Table 1: Match quality: Total exports (million Euros)

year	Customs	CIP	CIP match	Customs match
2000	77,649	73,353	59,493	62,063
2001	86,628	77,094	62,655	68,283
2002	87,996	80,946	72,729	69,077
2003	76,875	81,421	72,422	62,606
2004	78,836	84,641	75,262	64,778
2005	80,876	90,894	81,121	66,099
2006	80,524	93,160	81,604	64,498
2007	82,456	98,832	84,728	67,557
2008	80,042	81,138	76,438	63,146
2009	78,489	70,110	67,315	62,217

Notes: Customs refers to total industrial exports according to Customs data. CIP refers to total exports by our cleaned dataset of single-plant CIP firms, calculated based on CIP reported sales and CIP reported export shares. CIP match refers to total exports of matched firms, calculated based on CIP reported sales and CIP reported export shares. Customs match refers to total customs exports of single-plant CIP firms matched to customs data. Source: CSO and authors' calculations.

Table 2: Match quality: Number of exporters

year	Total CIP firms	Exporters (CIP)	Exporters (customs)	Exporters (both)
2000	3999	2085	2007	1651
2001	3981	2070	2047	1675
2002	4034	2067	2028	1659
2003	3968	2011	2089	1654
2004	3741	1882	2026	1592
2005	3539	1785	1895	1525
2006	3691	1821	1914	1524
2007	4205	1847	1895	1493
2008	4097	1754	1849	1426
2009	3873	1662	1817	1362

Notes: Total CIP firms refers to our cleaned dataset of single-plant CIP firms. Exporters (CIP) refers to the subset of Total CIP firms who report some positive exports in the CIP. Exporters (customs) refers to the subset of Total CIP firms who are matched with exports from customs data. Exporters (both) refers to firms in the sample classified as expoters by both definitions. Source: CSO and authors' calculations.

Table 3: Match quality: Number of firms exporting to different destinations

year	UK (CIP)	UK (customs)	UK (both)	US (CIP)	US (customs)	US (both)
2000	1842	818	738	678	747	472
2001	1800	823	733	730	769	493
2002	1785	811	721	730	756	491
2003	1733	813	728	697	714	475
2004	1626	797	711	633	677	440
2005	1521	771	684	616	651	427
2006	1563	744	667	606	608	399
2007	1586	728	650	603	574	385
2008	1502	719	641	568	503	342
2009	1386	718	601	565	524	353

Notes: Statistics are for our cleaned dataset of single-plant CIP firms. UK (CIP) refers to firms defined as exporters to the UK by CIP data. UK (customs) refers to firms defined as exporters to the UK by the match with customs data. Definitions for the US are analogous. Source: CSO and authors' calculations.

Table 4: Match quality: Summary statistics on exporters and non-exporters

	Mean employees			Mean sales ('000 EUR)			Mean age			Mean export %	
	Nonex.	Exporters		Nonex.	Exporters		Nonex.	Exporters		Exporters	
year	CIP	CIP	customs	CIP	CIP	customs	CIP	CIP	customs	CIP	customs
2000	23	102	109	4,425	43,266	44,331	16	18	18	46	37
2001	23	100	107	5,320	45,300	45,772	17	19	19	46	36
2002	21	96	103	4,181	47,362	52,091	17	20	20	45	34
2003	20	95	101	3,204	50,388	53,849	17	21	21	45	35
2004	21	98	104	3,624	55,247	57,766	17	21	22	44	34
2005	23	101	107	4,459	61,535	63,886	18	22	23	44	35
2006	22	100	107	4,096	63,371	66,281	17	22	23	43	34
2007	19	97	106	3,618	65,917	69,239	16	22	23	42	32
2008	17	92	101	3,746	58,890	66,840	15	22	23	42	32
2009	19	85	93	8,668	52,648	60,103	16	23	24	44	32

Notes: Statistics are for our cleaned dataset of single-plant CIP firms. CIP and customs refer to the definition of an exporter. Statistics for exporters under the customs definition are only for exporters reporting more than 500 Euro per year in exports. Export share under the customs definition is calculated as total exports from the customs match divided by sales reported in the CIP. Values greater than 100 are replaced by 100. Source: CSO and authors' calculations.

Table 5: Tariff variation over time and across countries

	Australia		Canada		Japan		US	
const	4.41	(0.02)**	4.29	(0.01)**	3.37	(0.01)**	3.98	(0.01)**
2001	-0.28	(0.02)**	-0.11	(0.02)**	-0.09	(0.01)**	-0.02	(0.01)**
2002	-0.28	(0.02)**	-0.27	(0.02)**	-0.16	(0.01)**	-0.12	(0.01)**
2003	-0.28	(0.02)**	-0.39	(0.02)**	-0.24	(0.01)**	-0.21	(0.01)**
2004	-0.28	(0.02)**	-0.45	(0.02)**	-0.32	(0.01)**	-0.28	(0.01)**
2005	-0.98	(0.02)**	-0.65	(0.02)**	-0.30	(0.01)**	-0.38	(0.01)**
2006	-0.98	(0.02)**	-0.65	(0.02)**	-0.34	(0.01)**	-0.38	(0.01)**
2007	-0.97	(0.02)**	-0.67	(0.02)**	-0.36	(0.01)**	-0.37	(0.01)**
2008	-0.97	(0.02)**	-0.67	(0.02)**	-0.36	(0.01)**	-0.37	(0.01)**
2009	-0.97	(0.02)**	-0.79	(0.02)**	-0.36	(0.01)**	-0.38	(0.01)**
hs6 f.e.	yes		yes		yes		yes	
R^2	0.96		0.98		0.99		1.00	
N	51444		511124		51036		48324	

Notes: Dependent variable is unweighted average ad valorem tariff at the HS6 level. All HS6 codes with at least one ad valorem tariff are included (i.e. the set of HS6 codes is not restricted to those for which positive exports from Ireland to the relevant destination are observed). No attempt is made to impute ad valorem equivalents for specific or mixed tariffs. Source: WTO.

Table 6: Entry: Linear probability model

	(1)		(2)		(3)		(4)	
Ownership	All		All		Irish-owned		Irish-owned	
Destinations	Extrastat only		All countries		Extrastat only		All countries	
	coeff	s.e.	coeff	s.e.	coeff	s.e.	coeff	s.e.
rer_t^k	0.04	(0.02)**	0.02	(0.01)**	0.07	(0.02)**	0.02	(0.01)**
$emp2_{t-1}^i * rer_t^k$	-0.01	(0.00)**	0.00	(0.00)**	-0.01	(0.00)**	0.00	(0.00)**
$emp3_{t-1}^i * rer_t^k$	-0.01	(0.01)**	0.00	(0.00)**	-0.02	(0.01)**	0.00	(0.00)**
dem_t^k	0.03	(0.01)**	0.01	(0.00)*	0.02	(0.01)	0.01	(0.00)*
$emp2_{t-1}^i * dem_t^k$	-0.01	(0.00)**	0.00	(0.00)**	-0.01	(0.00)**	0.00	(0.00)**
$emp3_{t-1}^i * dem_t^k$	-0.01	(0.00)**	0.00	(0.00)**	-0.02	(0.00)**	0.00	(0.00)**
tau_t^{ik}	0.35	(0.09)**	0.56	(0.07)**	0.30	(0.10)**	0.56	(0.07)**
$emp2_{t-1}^i * tau_t^{ik}$	-0.20	(0.10)**	-0.59	(0.07)**	-0.17	(0.10)**	-0.59	(0.07)**
$emp3_{t-1}^i * tau_t^{ik}$	-0.35	(0.10)**	-0.91	(0.08)**	-0.27	(0.11)**	-0.91	(0.08)**
Market f.e.	yes		yes		yes		yes	
Firm-year f.e.	yes		yes		yes		yes	
# firm-mkt-years	93748		494817		83926		494817	
# firm-years	35372		35501		32670		35501	
# firms	7538		7548		7059		7061	
R^2	0.36		0.25		0.33		0.25	
R^2-adj	-0.02		0.19		-0.11		0.19	

Notes: Estimation method is OLS. Dependent variable is an indicator for entry. Sample consists of all firm-mkt-years at risk for entry, and where there is positive lagged and current sales in the Irish market. Robust standard errors are calculated. ** indicates significance at the 5% level. * indicates significance at the 10% level. The omitted category is firms with 1-14 employees in the previous period. $emp2_{t-1}^i$ indicates that the firm had 15-99 employees and $emp3_{t-1}^i$ that the firm had 100+ employees in the previous period. The Extrastat only sample includes the Australia, Canada, Japan, Norway, Switzerland and the US. The "All countries" sample includes those countries and additionally Austria, Belgium, Denmark, Finland, France, Germany, UK, Italy, Netherlands, Portugal, Spain and Sweden.

Table 7: Entry: Conditional logit

	(1)		(2)	
Destinations	Extrastat only		All countries	
	coeff	s.e.	coeff	s.e.
rer_t^k	-0.02	(0.43)	0.73	(0.29)**
$emp2_{t-1}^i * rer_t^k$	0.03	(0.06)	-0.05	(0.02)**
$emp3_{t-1}^i * rer_t^k$	-0.05	(0.08)	-0.08	(0.02)**
dem_t^k	1.74	(0.68)**	0.49	(0.43)
$emp2_{t-1}^i * dem_t^k$	0.02	(0.05)	-0.07	(0.03)**
$emp3_{t-1}^i * dem_t^k$	-0.05	(0.06)	-0.08	(0.03)**
tau_t^{ik}	8.55	(1.58)**	12.15	(1.59)**
$emp2_{t-1}^i * tau_t^{ik}$	-4.04	(1.86)**	-13.54	(1.93)**
$emp3_{t-1}^i * tau_t^{ik}$	-9.19	(2.24)**	-21.76	(2.48)**
Market f.e.	yes		yes	
Firm-year f.e.	yes		yes	
# firm-mkt-years	9954		47337	
# firms	1120		1466	
pseudo-R^2	0.20		0.08	

Table 8: Exit

	(1)		(2)		(3)		(4)	
Ownership	All		All		Irish-owned		Irish-owned	
Destinations	Extrastat only		All countries		Extrastat only		All	
	coeff	s.e.	coeff	s.e.	coeff	s.e.	coeff	s.e.
rev_{t-1}^{ik}	-0.25	(0.05)**	-0.08	(0.01)**	-0.34	(0.04)**	-0.07	(0.01)**
rer_t^k	-0.25	(0.07)**	-0.08	(0.03)**	-0.30	(0.11)**	-0.08	(0.05)*
$emp2_{t-1}^i * rer_t^k$	-0.00	(0.01)	-0.00	(0.00)	0.00	(0.02)	-0.00	(0.00)
$emp3_{t-1}^i * rer_t^k$	-0.02	(0.01)	0.00	(0.00)	-0.01	(0.02)	0.01	(0.00)*
$rev_{t-1}^{ik} * rer_t^k$	0.02	(0.00)**	0.00	(0.00)**	0.02	(0.00)**	0.00	(0.00)**
dem_t^k	0.32	(0.13)**	0.09	(0.05)*	0.29	(0.20)	0.11	(0.08)
$emp2_{t-1}^i * dem_t^k$	-0.01	(0.01)	-0.00	(0.01)	-0.00	(0.01)	0.00	(0.01)
$emp3_{t-1}^i * dem_t^k$	-0.01	(0.01)	0.00	(0.01)	-0.00	(0.01)	0.01	(0.01)
$rev_{t-1}^{ik} * dem_t^k$	0.01	(0.00)**	0.00	(0.00)**	0.01	(0.00)**	0.00	(0.00)
tau_t^{ik}	-1.78	(0.69)**	-0.87	(0.52)*	-2.19	(0.80)**	-1.66	(0.60)**
$emp2_{t-1}^i * tau_t^{ik}$	2.42	(0.73)**	1.86	(0.52)**	2.05	(0.86)**	2.03	(0.60)**
$emp3_{t-1}^i * tau_t^{ik}$	2.38	(0.75)**	1.98	(0.56)**	2.57	(0.92)**	2.62	(0.56)**
$rev_{t-1}^{ik} * tau_t^{ik}$	-0.18	(0.07)**	-0.19	(0.05)**	-0.02	(0.09)	-0.19	(0.65)**
Market f.e.	yes		yes		yes		yes	
Firm-year f.e.	yes		yes		yes		yes	
# firm-mkt-years	11327		37470		6361		20968	
# firm-years	4897		7396		3207		5225	
# firms	1210		1571		878		1181	
R^2	0.62		0.43		0.67		0.47	
R^2-adj	0.32		0.29		0.33		0.29	

Notes: Estimation method is OLS. Dependent variable is an indicator for exit. Sample consists of all firm-mkt-years at risk for exit, and where there are positive lagged and current sales in the Irish market. Robust standard errors are calculated. ** indicates significance at the 5% level. * indicates significance at the 10% level. The omitted category is firms with 1-14 employees in the previous period. $emp2_{t-1}^i$ indicates that the firm had 15-99 employees and $emp3_{t-1}^i$ that the firm had 100+ employees in the previous period. The Extrastat only sample includes the Australia, Canada, Japan, Norway, Switzerland and the US. The "All countries" sample includes those countries and additionally Austria, Belgium, Denmark, Finland, France, Germany, UK, Italy, Netherlands, Portugal, Spain and Sweden.

Table 9: Exit: Conditional logit

	(1)		(2)	
Destinations	Extrastat only		All countries	
	coeff	s.e.	coeff	s.e.
rev_{t-1}^{ik}	-0.81	(0.30)**	-0.72	(0.10)**
rer_t^k	-2.78	(0.75)**	-1.30	(0.42)**
$emp2_{t-1}^i * rer_t^k$	0.12	(0.11)	-0.03	(0.04)
$emp3_{t-1}^i * rer_t^k$	-0.05	(0.13)	0.01	(0.04)
$rev_{t-1}^{ik} * rer_t^k$	0.02	(0.02)	-0.01	(0.01)
dem_t^k	4.22	(1.21)**	1.18	(0.66)*
$emp2_{t-1}^i * dem_t^k$	0.06	(0.08)	-0.03	(0.05)
$emp3_{t-1}^i * dem_t^k$	-0.05	(0.09)	-0.04	(0.05)
$rev_{t-1}^{ik} * dem_t^k$	0.01	(0.02)	0.00	(0.01)
tau_t^{ik}	-14.64	(4.86)**	-15.14	(3.91)**
$emp2_{t-1}^i * tau_t^{ik}$	13.84	(5.34)**	13.11	(4.20)**
$emp3_{t-1}^i * tau_t^{ik}$	18.17	(5.54)**	15.47	(4.35)**
$rev_{t-1}^{ik} * tau_t^{ik}$	-1.08	(0.77)	1.02	(0.43)**
Market f.e.	yes		yes	
Firm-year f.e.	yes		yes	
# firm-mkt-years	4041		21031	
# firms	578		1029	
pseudo-R^2	0.39		0.38	

Table 10: Sales

	(1)		(2)		(3)		(4)	
Ownership	All		All		Irish-owned		Irish-owned	
Destinations	Extrastat only		All countries		Extrastat only		All	
	coeff	s.e.	coeff	s.e.	coeff	s.e.	coeff	s.e.
Δrer_t^k	0.70	(0.76)	0.84	(0.32)**	0.75	(1.15)	0.94	(0.45)**
Δdem_t^k	4.36	(2.64)*	2.34	(0.69)**	4.02	(3.91)	2.47	(0.94)**
Δtau_t^{ik}	-3.97	(2.79)	-5.45	(2.64)**	-4.15	(3.48)	-5.54	(3.40)
Market f.e.	yes		yes		yes		yes	
Firm-year f.e.	yes		yes		yes		yes	
# firm-mkt-years	3100		13057		1650		7231	
# firm-years	1639		3468		953		2338	
R^2	0.59		0.35		0.66		0.42	
R^2-adj	0.12		0.11		0.18		0.13	

Notes: Estimation method is OLS. Dependent variable is an the change in log Euro revenue deflated by the Irish CPI. Sample consists of all firm-mkt pairs for which continuous participation is observed throughout the sample. Robust standard errors are calculated. ** indicates significance at the 5% level. * indicates significance at the 10% level.