NBER WORKING PAPER SERIES

EXPORTERS AND SHOCKS: DISSECTING THE INTERNATIONAL ELASTICITY PUZZLE

Doireann Fitzgerald
Stefanie Haller

Working Paper 19968
http://www.nber.org/papers/w19968

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2014

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. Expert research assistance was provided by Jackie Chan and Adrian Corcoran. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. Stefanie Haller is grateful for financial support from the IRC under the grant titled "Exporters and Shocks." We are grateful for comments and suggestions from Costas Arkolakis, Ina Jäkel, Richard Rogerson, Kim Ruhl and Jim Tybout. The views expressed here are those of the authors and not those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

Exporters and Shocks: Dissecting the International Elasticity Puzzle
Doireann Fitzgerald and Stefanie Haller
NBER Working Paper No. 19968
March 2014
JEL No. F14,F41

ABSTRACT

Aggregate exports are not very responsive to real exchange rates, though they respond strongly to trade liberalizations, a fact sometimes referred to as the International Elasticity Puzzle. We use micro data on firms and exports for Ireland to dissect the puzzle. Our identification strategy uses within-firm-year cross-market variation in real exchange rates and tariffs to identify the responses of export participation, export revenue and the product dimension of exporting to these variables. We show that (i) the weak response of export revenue of long-time market participants to real exchange rates is key to the behavior of aggregate exports, (ii) export participation also responds less to real exchange rates than to tariffs, but this alone cannot explain the puzzle; and (iii) the revenue response of long-time market participants cannot be accounted for by product entry responses. Hence any model that can successfully account for the puzzle needs to match the intensive margin responses of exporting firms.

Doireann Fitzgerald
Federal Reserve Bank of Minneapolis
90 Hennepin Ave
Minneapolis, MN 55401
and NBER
doireann.fitzgerald@gmail.com

Stefanie Haller
University College Dublin
School of Economics
Belfield
Dublin 4
Ireland
stefanie.haller@ucd.ie

An online appendix is available at:
http://www.nber.org/data-appendix/w19968

1 Introduction

Aggregate exports are not very responsive to movements in real exchange rates. As a result, standard calibrated models of international business cycles typically assume a low elasticity of substitution between home goods and foreign goods (in the range 0.5 to 1.5) in order to match comovements 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 exports to tariff liberalizations and changes in trade costs. It is typical to use elasticities on the order of 3 to 12 to calibrate the exact same models to match trade facts. This discrepancy is sometimes referred to as the International Elasticity Puzzle.[1] This puzzle is not purely a theoretical curiosity. The absence of a tractable framework which can match both macro facts and trade facts is a major obstacle to quantitative work on any question that simultaneously touches on intertemporal and intratemporal trade.[2] In this paper we provide direct evidence on the origins of the International Elasticity Puzzle at the firm level. We thereby provide guidance on how to modify standard models in order to resolve the puzzle.

We use 14 years of micro data on firms and exports for Ireland to provide this evidence. Using firm-year fixed effects to control for costs, we estimate the responses of export participation, export revenue and the product dimension of exporting at the firm-market level to real exchange rates and tariffs (the "shocks" of our title). We find that controlling for costs, both export participation and the export revenue of long-time market participants respond more strongly to variation in tariffs than to variation in real exchange rates. This revenue response is due not to product entry and exit within the firm and the market, but to the response of the continuing products which account for the bulk of the export revenue of long-time participants. Our estimates imply that revenue responses rather than participation responses are quantitatively key to matching the behavior of aggregate exports. We conclude that in order to resolve the International Elasticity puzzle, models must match the intensive margin responses of exporting firms.

We start by illustrating the International Elasticity Puzzle using a standard model of the exporter's problem. In the absence of adjustment frictions at the firm-market level, and conditional on costs and foreign real demand, the model implies that the elasticity of aggregate exports with respect to the real exchange rate is the negative of the elasticity with

[1] Ruhl (2008) coins the term, and also provides a useful survey of the evidence.

[2] In addition to Ruhl (2008), recent papers which address the puzzle are Arkolakis, Eaton and Kortum (2012) and Feenstra, Obstfeld and Russ (2012).

respect to gross ad valorem tariffs (in a CES world, this elasticity is equal to the elasticity of substitution between different goods). The intuition is that conditional on costs and foreign real demand, revenue and profits depend only on the *ratio* of the real exchange rate and the gross ad valorem tariff. Of course in the real world, since exchange rates and tariffs have different persistence and volatility, costs of adjustment may induce firms to respond quite differently to the two variables.[3] Our goal is to provide clean evidence on whether, and along which margins, firm responses do in fact differ.

As a preliminary to our analysis of the micro data, we confirm the International Elasticity Puzzle at the aggregate level for Ireland. For the period covered by our micro data, we find that relative exports across pairs of export markets respond more elastically to the relative tariff faced by Irish exporters in the two markets than they do to the bilateral real exchange rate between the pair of markets.

The main part of our analysis makes use of confidential firm and export micro data for Ireland for the period 1996-2009, constructed by merging customs records with the annual census of manufacturing, mining and utilities, and a comprehensive survey of products produced. In our data, on average more than 75% of export revenue is accounted for by firms which are both incumbent and continuing in the relevant export market.[4] For incumbent and continuing exporters defined at the market level, over 90% of export revenue is accounted for by incumbent and continuing products.[5] Even before we estimate the responses of different margins of exporting to shocks, these facts suggest that the responses of long-time market participants, and in particular the response of their revenue from continuing products, are key to understanding the behavior of aggregate exports.

We next estimate the responsiveness of different margins of exporting behavior at the firm-market level to both market-specific macro variables (real exchange rates and real demand) and firm-market-specific tariffs, constructed using tariff data at the product level from the WTO. Our empirical strategy builds on a substantial literature that estimates reduced form dynamic discrete choice models of export supply with sunk costs of exporting (e.g. Roberts and Tybout (1997), Bernard and Wagner (2001), Bernard and Jensen (2004)). We use firm-year fixed effects to control for heterogeneity in costs across firms and variation in

[3]Ruhl (2008) is the first to propose adjustment frictions as an explanation for the International Elasticity Puzzle. We discuss his contribution at greater length below.

[4]We refer to firms which do not enter in the current period as incumbents, and firms which do not exit in the next period as continuing firms. Eaton et al (2008) document a similar fact for Colombia, and Bernard, Jensen, Redding and Schott (2009) present related facts about the key role played by the intensive margin in export growth for the US.

[5]This is very similar to the findings of van Beveren, Bernard and Vandenbussche (2012) using comparable Belgian data.

costs within firms over time. This is important, because costs and real exchange rates are not fundamental shocks, but are jointly determined in general equilibrium, and it is only responses of exports *conditional* on costs that are predicted to be equal (in absolute value) in standard models.

We find that export participation responds to real exchange rates and tariffs broadly in the direction one would expect. Though the exact response of participation varies across the distribution of firm-market pairs, real depreciations of the home currency against that of the destination market tend to increase entry and reduce exit for all firms (vice versa for appreciations). For medium and large firms, tariff reductions tend to increase entry and reduce exit (vice versa for tariff increases). The quantitative impact on participation is larger for tariffs than for real exchange rates. For example, for medium-sized firms with 100-249 employees (the firms whose participation is most sensitive to tariffs) a 10% depreciation of the real exchange rate increases the entry rate for the average market from a baseline of 3% to 3.1% and reduces the exit rate from a baseline of 23% to 22.7%. This would imply an increase in steady state participation of 0.5 percentage points, from 11.5% to 12%. For the same firms, a comparable reduction in tariffs from 10% to zero increases the entry rate from the baseline of 3% to 3.3%, and reduces the exit rate from 23% to 20%. This would imply an increase in steady state participation of 2.6 percentage points, from 11.5% to 14.2%.

In examining the responses of export revenue conditional on participation, we restrict our attention to long-time market participants, as selection bias is a concern for participants who are close to their entry and exit thresholds. Our point estimate of the elasticity of firm-market-level export revenue with respect to real exchange rates is 0.64. Our point estimate of the elasticity of export revenue with respect to tariffs is -2.42. Both elasticities are significantly different from zero, and they are significantly different from each other in absolute value. For the same long-term market participants, looking only at their continuing products, the elasticity of firm-product-market-level export revenue with respect to real exchange rates is 0.80, while the elasticity with respect to tariffs is -2.65. Meanwhile, for these long-term participants, the share of export revenue accounted for by entering and exiting products is more responsive to tariffs than to real exchange rates, but contrary to what one would expect with sunk costs of product introduction, we find that real exchange rate depreciations and tariff reductions tend to *reduce* the share of revenue accounted for by entering products.

A simple back-of-the-envelope calculation based on our estimates suggests that revenue responses are quantitatively key to explaining the behavior of aggregate exports. Even

shutting down the participation margin altogether, aggregate exports would respond with a counterfactually high elasticity to real exchange rates if the export revenue of continuing exporters were to respond to real exchange rates as it does to tariffs. This implies that in order to resolve the International Elasticity puzzle, models must match the intensive margin responses of exporting firms. While sunk costs of export participation are an obvious reason why participation should respond less to real exchange rates than to tariffs, it is less obvious why revenue should respond differently. We point to some potential mechanisms, leaving quantitative and empirical exploration of these mechanisms for future research.

Our work is related to several literatures. As already mentioned, it is very closely related to Ruhl (2008). In the context of a calibrated two-country general equilibrium model, he finds that sunk costs of export participation can generate big differences between the responses of aggregate exports to real exchange rates (driven by mean-reverting productivity) and tariffs (assumed to be persistent). We provide direct empirical support for his story, though we also point to a key empirical role for differential responses of export revenue conditional on participation in explaining aggregate elasticities. A more recent contribution to the literature on the International Elasticity Puzzle is Arkolakis, Eaton and Kortum (2012). They propose a model where customers shift relative demand slowly in response to relative price changes. Their model can explain why short-run and long-run aggregate responses to the same variable differ. But it cannot explain why over the same time horizon, exports respond differently to real exchange rates and tariffs, something we observe strongly in the micro as well as the macro data. Another paper which attempts to rationalize different elasticities is Feenstra, Obstfeld and Russ (2012). They argue that the elasticity of substitution between home and foreign goods is empirically different from the elasticity of substitution between foreign goods from two different countries. However this cannot explain why (as we document) both at the aggregate and at the firm level, the elasticity of relative exports across pairs of destination markets responds differently to movements in relative prices depending on what is driving the price movements.

Our results also speak to an older theoretical literature which argues that the expenditure-shifting effects of exchange rate movements may depend on sunk costs of export entry (Baldwin (1988), Baldwin and Krugman (1989) and Dixit (1989)) and to more recent research that proposes that entry and exit responses can (or cannot) explain facts about international real business cycles. Contributions in this area include Ghironi and Melitz (2005), Alessandria and Choi (2007), Atkeson and Burstein (2008). Our findings imply that export entry and exit play little role in explaining real exchange rate dynamics at a business cycle frequency,

as export participation is relatively insensitive to real exchange rates.

While we are the first to provide an integrated empirical treatment of exporter responses to both real exchange rates and tariffs, responses to these two variables have been previously explored by two separate lines of research. Campa (2004) estimates the responses of export participation and the export revenue of continuing exporters to exchange rates. Our findings on responses to exchange rates are similar to his. Our findings on responses to exchange rates are also very similar to those of Berman, Martin and Mayer (2012) who use French data to estimate price, quantity and participation elasticities to exchange rates. In the trade literature, there is considerable interest in the effect of trade liberalizations on export participation and export revenue. But empirical work on estimating firm-market-level responses to tariffs has been hampered by the difficulty of constructing appropriate firm- and market-specific tariff measures. An exception to this is Lileeva and Trefler (2010), who focus not on the contemporaneous effects of liberalization on participation and sales, but on the consequent effects on productivity upgrading. While there are methodological similarities between our work and theirs, there are no results which are directly comparable.

The next section of the paper outlines a model of the exporter's problem. The third section replicates the International Elasticity Puzzle using aggregate data for Ireland. The fourth section describes the empirical strategy we use in our analysis of the micro data. The fifth section describes our micro data and presents relevant summary statistics. The sixth section describes our empirical results. The seventh section presents a calculation of the relative importance of the different margins of firm responses based on these results. The final section concludes.

2 Model

To illustrate the International Elasticity Puzzle, and to frame our analysis of the micro data, we present a partial equilibrium model of the potential exporter's problem with both fixed and sunk costs of export participation, as in e.g. Roberts and Tybout (1997). To fix ideas, we present a version of the model where demand is CES (in the Appendix, we relax this). For simplicity, we condition on participation in the domestic market, and do not model the firm existence decision.[6] We also make a number of assumptions under which we can treat participation in different export markets as independent decisions after conditioning on a common marginal cost. This is in contrast to, e.g. Morales, Sheu and Zahler (2014) who

[6]We thus ignore firms which sell only to the foreign market, entry of firms that are born global, and exit of firms from export markets that is accompanied by firm death.

explicitly allow for dependence of sunk costs of entering a new market on the set of markets the firm is already in.

2.1 Demand and marginal cost

The factory gate price that firm i charges to buyers from market k expressed in the producer's home currency is P_t^{ik}. Let E_t^k be the nominal exchange rate between the producer's currency and the currency of market k, and let $\left(1+\tau_t^{ik}\right)$ be the gross ad valorem tariff firm i faces in market k. The foreign currency price actually faced by buyers in market k is the factory gate price converted to foreign currency, inclusive of the ad valorem tariff, i.e. $\left(1+\tau_t^{ik}\right)\left(P_t^{ik}/E_t^k\right)$.[7] Assume that the demand faced by firm i in market k at time t takes the form:

$$Q_t^{ik} = \left(\frac{\left(1+\tau_t^{ik}\right)\left(P_t^{ik}/E_t^k\right)}{P_t^{k*}}\right)^{-\theta} Q_t^k \exp\left(\eta_t^{ik}\right) \tag{1}$$

Here, P_t^{k*} is the aggregate price level in market k, expressed in the currency of market k, Q_t^k is real aggregate demand in market k, and η_t^{ik} is a random shock to demand. We make the simplifying assumption that P_t^{k*} is a good measure of competitors' prices in market k.[8]

Firm i faces nominal marginal cost C_t^i of serving all markets, expressed in terms of domestic currency.[9] Real marginal cost is $\tilde{C}_t^i = C_t^i/P_t$, where P_t is the aggregate price level in the home market.

2.2 Intratemporal optimization, revenue and profits

We assume that prices are flexible. Then it does not matter which currency is used to set prices, and moreover, the optimal price conditional on participation is chosen to maximize current profits only. A firm which participates in market k chooses the home currency price P_t^{ik} to maximize real profits at time t, where this is given by:

$$\Pi_t^{ik} = \left(\frac{P_t^{ik}-C_t^i}{P_t}\right)\left(\frac{\left(1+\tau_t^{ik}\right)\left(P_t^{ik}/E_t^k\right)}{P_t^{k*}}\right)^{-\theta} Q_t^k \exp\left(\eta_t^{ik}\right) \tag{2}$$

[7] We ignore trade costs here, though we allow for them in the empirical part of the paper.

[8] We do not allow for the possibility that exchange rates may have different effects on competitors' prices than tariffs do.

[9] Our identification strategy relies on marginal cost being scale-independent.

Given the CES assumption, the optimal home currency price is a constant markup over home currency nominal marginal cost:

$$P_t^{ik} = \frac{\theta}{\theta - 1} C_t^i \tag{3}$$

Maximized real revenue is then given by:

$$R_t^{ik} = \left(\frac{\theta}{\theta - 1}\right)^{1-\theta} \left(\tilde{C}_t^i\right)^{1-\theta} \left(\frac{\left(1 + \tau_t^{ik}\right)}{RER_t^k}\right)^{-\theta} Q_t^k \exp\left(\eta_t^{ik}\right) \tag{4}$$

Here, we have rearranged to show explicitly how revenue depends on the real exchange rate, $RER_t^k = E_t^k P_t^{k*}/P_t$. Maximized real gross flow profits are given by this expression divided by θ. Note that, holding fixed $\tilde{C}_t^i$ and Q_t^k, revenue and hence profits depend only on the *ratio* of the tariff to the real exchange rate.

2.3 The International Elasticity Puzzle

To illustrate the International Elasticity Puzzle, suppose there are no fixed or sunk costs of market participation. Because the profits implied by (4) are always positive, all firms participate in market k. Further, assume that all firms in the home country have the same real marginal cost $\tilde{C}_t$. Let J_j denote the set of home firms that face tariff τ_t^{jk} in market k at time t. Summing across the export revenue from market k of all firms i such that $i \in J_j$, we get

$$R_t^{jk} = \left(\frac{\theta}{\theta - 1}\right)^{1-\theta} \left(\tilde{C}_t\right)^{1-\theta} \left(\frac{\left(1 + \tau_t^{jk}\right)}{RER_t^k}\right)^{-\theta} Q_t^k \sum_{i \in J_j} \exp\left(\eta_t^{ik}\right) \tag{5}$$

This implies that aggregate product-j export revenue from market k relative to market k' can be written:

$$\frac{R_t^{jk}}{R_t^{jk'}} = \left(\frac{RER_t^k / RER_t^{k'}}{1 + \tau_t^{jk} / 1 + \tau_t^{jk'}}\right)^{\theta} \frac{Q_t^k}{Q_t^{k'}} \frac{\sum_{i \in J_j} \exp\left(\eta_t^{ik}\right)}{\sum_{i \in J_j} \exp\left(\eta_t^{ik'}\right)} \tag{6}$$

where by definition, $RER_t^k / RER_t^{k'} = RER_t^{k'k}$. Clearly, the elasticity of relative export revenue with respect to the real exchange rate is equal to the negative of the elasticity of relative revenue with respect to the relative tariff faced by home firms in the two markets. We show in the Appendix that this prediction does not depend on the assumption of CES demand, though with non-constant markups, the elasticity of revenue with respect to the real exchange rate is not equal to the price elasticity of demand as it is in the CES case.

One manifestation of the International Elasticity Puzzle is the fact that when an equation like (6) is estimated using aggregate data, the elasticity of relative exports with respect to real exchange rates is lower than the elasticity with respect to relative tariffs (expressed in absolute value). We presently illustrate this fact using aggregate data for Ireland. Another manifestation is the fact that when a general equilibrium model which nests an exporter problem like the stripped-down version outlined here is calibrated to match comovements between prices and quantities, the appropriate choice of θ depends on the *source* of changes in relative prices.

2.4 Dynamic optimization and participation

Of course not all firms participate in all markets. Now suppose there is a fixed cost F^k in real terms of participating in market k in any period, and a sunk cost of entry S^k, also in real terms.[10] The firm's decision about whether or not to participate in market k is forward-looking decision in the presence of sunk costs.

Let $X_t^{ik} \in \{0,1\}$ be an indicator variable for participation in market k. Let $Z_t^{ik} = \left\{\tilde{C}_t^i, RER_t^k, Q_t^k, \tau_t^{ik}, \eta_t^{ik}\right\}$ be a vector of firm-and-market-specific exogenous state variables. The firm's dynamic problem can then be written:

$$V\left(X_{t-1}^{ik}, Z_t^{ik}\right) =$$

$$\max_{X_t^{ik} \in \{0,1\}} \left\{ \begin{array}{c} X_t^{ik}\Pi\left(Z_t^{ik}\right) \\ -X_t^{ik}F^{ik} - X_t^{ik}\left(1 - X_{t-1}^{ik}\right)S^{ik} \\ +\beta\mathbb{E}_t V\left(X_t^{ik}, Z_{t+1}^{ik}\right) \end{array} \right\} \tag{7}$$

The solution to this problem is a policy function for participation, where the optimal policy depends on the endogenous and exogenous state variables of the firm's problem, $\left\{X_{t-1}^{ik}, Z_t^{ik}\right\}$. A familiar way to write this optimal policy is in the form of a threshold for participation in terms of one exogenous state variable, where the threshold is a function of past participation,

[10] For illustrative purposes, we assume these are fixed, but more generally they could be stochastic as in e.g. Das, Roberts and Tybout (2007).

as well as the remaining exogenous state variables of the firm's problem. For example:

$$X_t^{ik} = \begin{cases} 0 & \text{if} \quad \eta_t^{ik} < \Psi\left(X_{t-1}^{ik}, \tilde{C}_t^i, RER_t^k, Q_t^k, \tau_t^{ik}\right) \\ \\ 1 & \text{if} \quad \eta_t^{ik} \geq \Psi\left(X_{t-1}^{ik}, \tilde{C}_t^i, RER_t^k, Q_t^k, \tau_t^{ik}\right) \end{cases} \tag{8}$$

Because of the sunk cost, $\Psi\left(0, \tilde{C}_t^i, RER_t^k, Q_t^k, \tau_t^{ik}\right) > \Psi\left(1, \tilde{C}_t^i, RER_t^k, Q_t^k, \tau_t^{ik}\right)$.

The existence of sunk costs per se need not imply that participation (and hence aggregate exports) respond differently to real exchange rates and tariffs. But the fact that tomorrow's participation threshold depends on today's participation decision implies that some part of the payoff from entering (or deciding to stay in) market k today comes through the expected *future* net benefit of participation, rather than just the net benefit today. This implies that the sensitivity of the participation thresholds to an innovation in RER_t^k or τ_t^{ik} will depend on the process from which the innovation is drawn. Although holding constant costs and foreign aggregate demand, profits depend only on the ratio of tariffs and real exchange rates, if tariffs are more persistent and less volatile than real exchange rates, participation today will respond more to a change in tariffs today than to the equivalent change in real exchange rates. This is because the impact of the tariff change on the expected future net benefit of participation is greater than the impact of the equivalent real exchange rate change. As a result, aggregate exports will respond more to tariffs than to real exchange rates. This explanation for the International Elasticity Puzzle is proposed by Ruhl (2008). We use micro data to test it directly by estimating reduced form approximations to (4) and (8).

3 Replicating the International Elasticity Puzzle

As a preliminary to our analysis of the micro data, we use aggregate data for Ireland to confirm that aggregate exports respond more to tariffs than to real exchange rates. We collect Comtrade data on bilateral exports from Ireland to a select set of destination markets at the HS6 digit level for the period 1996-2009.[11] We construct bilateral real exchange rates between Ireland and each destination market using data on annual average nominal exchange rates and CPIs from the IMF's International Financial Statistics. We use the CPI data and data on National Accounts from from the World Bank's World Development Indicators to construct a measure of real aggregate demand in each destination market. We collect data

[11]The destinations are listed in the Appendix. They are the same destinations we use in our analysis of the micro data.

on the tariffs faced by Irish exporters in these destinations at the HS6-digit level from the WTO. The macro and tariff data are described in detail in sections 5.3 and 5.4. Using this data, we take logs of (5), and estimate the following equation:

$$r_t^{jk} = \alpha^k + c_t^j + \boldsymbol{\beta}' \mathbf{z}_t^{jk} + \eta_t^{jk} \tag{9}$$

In this expression, r_t^{jk} is the log of revenue from exports of product j to country k. α^k is a destination-market fixed effect. c_t^j is a product-year fixed effect. The vector $\mathbf{z}_t^{jk}$ contains the log of the real exchange rate, the log of real demand in the destination market, and the log of one plus the ad valorem tariff.

The results from estimating (9) are reported in Table 1. The estimated coefficient on the log real exchange rate is 0.42 (standard error 0.15), the estimated coefficient on log real foreign demand is 0.54 (standard error 0.13), while the estimated coefficient on the tariff variable is -1.57 (standard error 0.24). Irish exports respond more elastically to tariffs than they do to real exchange rates, though the difference between the elasticity with respect to real exchange rates and the (absolute value of the) elasticity with respect to tariffs is lower than some estimates in the literature.

4 Empirical strategy

We now turn to our analysis of the micro data.

4.1 Export participation

Motivated by equation (8), the standard reduced form export participation equation (e.g. Roberts and Tybout (1997)) regresses a participation indicator on an indicator for lagged participation, a vector of (lagged) variables that are intended to capture time variation in costs at the firm or plant level, time fixed effects, and plant or firm fixed or random effects. In our baseline specification, we modify this approach along two dimensions.

First, because we observe participation at the level of the firm-market pair, we can use firm-year fixed effects to control for time-varying marginal cost at the firm level. This is important given our focus on the sensitivity of participation to real exchange rates and tariffs. It allows us to clean out that part of both real exchange rates (and less importantly) tariffs that may have an effect on participation in *all* export markets through their effect on marginal cost. As a result the effects of these variables on participation are identified purely

from within-firm-year variation across markets.

Second, we allow the response of participation to real exchange rates and tariffs to differ across firm-market observations by interacting these variables with variables which vary at the level of the firm, the market and the year. We do this because the participation of firm-market pairs close to their entry or exit thresholds should be more sensitive to real exchange rates and tariffs than the participation of firm-market pairs that are far from these thresholds. And as equation (8) implies, these thresholds (expressed in terms of the idiosyncratic demand shock η_t^{ik}) differ systematically with firm-market characteristics. Imposing a uniform sensitivity across all firm-market-year observations could lead us to underestimate the impact of shocks on the participation decision.

The resulting estimating equation is:

$$\Pr\left[X_t^{ik}=1\right] = \alpha^k + c_t^i + \boldsymbol{\lambda}'\mathbf{s}_{t-1}^{ik} + \boldsymbol{\beta}'\mathbf{z}_t^{ik} + \boldsymbol{\gamma}'\left(\mathbf{s}_{t-1}^{ik}\otimes\mathbf{z}_t^{ik}\right) + \varepsilon_t^{ik} \tag{10}$$

Here, α^k is a market-specific effect which captures any time-invariant factors (e.g. trade costs) which lead the probability of participation for all firms to be greater in some markets than others. c_t^i is a firm-year fixed effect which captures the first-order effect of marginal cost on participation, as well as any effects of other variables that are common across all markets for a given firm at a point in time. Because we estimate a linear probability model, it is straightforward to condition out these fixed effects. $\mathbf{s}_{t-1}^{ik}$ is a vector of variables which in the baseline specification includes an indicator for lagged participation, a vector of indicators for lagged firm size,[12] the indicators for size interacted with lagged participation, and the log of lagged real revenue from the relevant market interacted with lagged participation. $\mathbf{z}_t^{ik}$ is a vector of macro and tariff variables. The baseline vector contains the real exchange rate between the home market and market k, a measure of real demand in market k, along with a firm-market-year-specific ad valorem tariff. These variables are constructed as described below in sections 5.3 and 5.4. Finally, ε_t^{ik} is the error term (implicitly, a transformation of the idiosyncratic demand shock). We cluster standard errors at the market-year level.

We also examine the possibility of dynamic responses to the variables of interest by estimating dynamic entry and exit equations based on differencing the entry and exit equations implied by (10), and including lags of $\Delta\mathbf{z}_t^{ik}$.

[12]The level effect of firm size is soaked up by the firm-year fixed effects.

4.2 Export revenue

Based on (4), we estimate the following empirical specification for export revenue conditional on participation:

$$r_t^{ik} = \alpha^k + c_t^i + \boldsymbol{\beta}' \mathbf{z}_t^{ik} + \eta_t^{ik} \tag{11}$$

Here, r_t^{ik} is log export revenue of firm i from market k at date t. As in the participation equation, α^k is a market-specific fixed effect which captures time-invariant reasons why revenue might be higher for all firms in some markets than others. c_t^i is a firm-year fixed effect which captures marginal cost, as well as any other factors common across markets for a given firm at a point in time. $\mathbf{z}_t^{ik}$ is as above a vector of macro and tariff variables, whose construction is described in sections 5.4 and 5.3. η_t^{ik} captures anything that is idiosyncratic to the firm, market and year (e.g. idiosyncratic demand shocks).

A potentially important issue in estimating (11) is selection. We only observe export revenue for firm-market-years where $X_t^{ik} = 1$. As we describe above, participation depends on the same unobserved idiosyncratic shocks to demand which appear in the revenue equation. This implies that for firm-market-years close to their entry and exit thresholds, the expectation of η_t^{ik} conditional on the independent variables need not equal zero.

There are several aspects of our setup which make the standard approaches to controlling for selection (such as a Heckman correction) inappropriate or tricky to implement. The most obvious of these is the absence of any variable that plausibly affects export participation but not export revenue. Since we use a linear probability model in the first stage (motivated by our desire to use fixed effects to control for marginal cost), we cannnot rely on identification through nonlinearities. So instead, what we do is to restrict our estimation sample to a set of firm-market pairs we think likely to be far from their participation thresholds. Our implicit assumption is that for these observations, the distribution of idiosyncratic shocks η_t^{ik} is not truncated. Our baseline estimation sample is firm-market observations for which we observe at least 10 years of continuous participation, and no in-sample exit. If entry is observed in-sample, we drop the first two years of participation. This approach is not ideal, but it has at least the merit of transparency. We check the robustness of our results to varying the estimation sample. Under all specifications, we cluster standard errors at the market-year level.

We also examine the possibility of dynamic responses to the variables of interest by

estimating a dynamic version of (11) in differences:

$$\Delta r_t^{ik} = \alpha^k + c_t^i + \sum_{s=0} \boldsymbol{\beta}_s' \Delta \mathbf{z}_{t-s}^{ik} + \xi_t^{ik}$$

As with the baseline revenue equation, we cluster standard errors at the level of the market and the year.

4.3 Product-level analysis

The firms in our data produce and sell multiple products.[13] The model we describe above can be extended to take account of this in a straightforward way by defining the problem at the level of the firm, the product and the market instead of at the level of the firm and the market. Based on this, we exploit the product-level data to investigate the responses of long-time market participants to real exchange rates and tariffs along two dimensions. First, we estimate the response of firm-product-market revenue from continuing products sold by long-time market participants to real exchange rates and tariffs. Second, we examine the extent to which long-time market participants respond to real exchange rates and tariffs by introducing new products or by retiring existing products.

4.3.1 Revenue from continuing products

We estimate an analog of the revenue equation (11) at the level of the firm, the product and the market. Our baseline estimating equation is:

$$r_t^{ijk} = \alpha^k + c_t^{ij} + \boldsymbol{\beta}' \mathbf{z}_t^{ijk} + \eta_t^{ijk} \tag{12}$$

We include all continuing-product observations for firm-market pairs included in the baseline estimation of (11). We include firm-product-year fixed effects c_t^{ij} to control for marginal cost. Implicitly we assume that marginal cost is the same across markets within a firm, a product and a year. The construction of the tariff variable used in this equation is described below in section 5.4. As with the baseline revenue equation, we cluster standard errors at the level of the market and the year.

[13] There is a growing literature on the product dimension of firms' exporting behavior, e.g. Bernard, Redding and Schott (2011).

We also estimate a version of (12) in differences:

$$\Delta r_t^{ijk} = \alpha^k + c_t^i + \sum_{s=0} \boldsymbol{\beta}_s' \Delta \mathbf{z}_{t-s}^{ijk} + \xi_t^{ijk}$$

Here, we include only firm-year fixed effects, assuming that the time-varying component of marginal cost is common across products within a firm, while constant firm-product differences in cost levels are differenced out. We again cluster standard errors at the level of the market and the year.

4.3.2 Entry and exit of products

While it is easy to define the set of possible markets a firm might export to, it is less obvious how to define the range of products it might export. Clearly the firm could export products that it already produces, but it could also start producing and exporting entirely new products in response to changes in exchange rates or tariffs. Given this, we do not treat entry and exit of products symmetrically with participation at the market level. Instead, we regress the share of export revenue at the firm-market level that is due to entering products and the share of export revenue that is due to exiting products on a similar set of independent variables to those included in (11), our revenue equation at the firm-market level:

$$shenter_t^{ik} = \alpha^k + c_t^i + \boldsymbol{\beta}' \mathbf{z}_t^{ik} + \eta_t^{ik} \tag{13}$$

$$shexit_t^{ik} = \alpha^k + c_t^i + \boldsymbol{\beta}' \mathbf{z}_t^{ik} + \eta_t^{ik} \tag{14}$$

Here, $shenter_t^{ik}$ and $shexit_t^{ik}$ are given by:

$$shenter_t^{ik} = \frac{\sum_{j \in J_t^{ik}, j \notin J_{t-1}^{ik}} R_t^{ijk}}{\sum_{j \in J_t^{ik}} R_t^{ijk}}$$

and

$$shexit_t^{ik} = \frac{\sum_{j \in J_{t-1}^{ik}, j \notin J_t^{ik}} R_t^{ijk}}{\sum_{j \in J_{t-1}^{ik}} R_t^{ijk}}$$

In these expressions, R_t^{ijk} is revenue of firm i from product j in market k at date t, and J_t^{ik} is the set of products that firm i sells in market k at time t. For exiting products, we also experiment with lagging the dependent variable one period relative to the independent variables. The interpretation of c_t^i in (13) and (14) is somewhat different from the interpretation

of the fixed effects in (11) and (12), as it captures variation in entry and exit shares that is common across markets within a firm at a given point in time. This may be due to costs, but it may also be due to the product cycle specific to that firm. The tariff variable in (13) and (14) is constructed as in the baseline participation equation, using production weights. The estimation sample is restricted to the firm-market pairs included in the estimation of (11). We cluster standard errors at the level of the market and the year.

5 Data

5.1 Micro data

Our work makes use of three sources of confidential micro data made available to us by the Central Statistics Office in Ireland: the Irish Census of Industrial Production (CIP), Irish customs records, and the Irish Prodcom survey.

5.1.1 Census of Industrial Production

The CIP, which covers manufacturing, mining and utilities, takes place annually. Firms with 3 or more persons engaged are required to file returns.[14] We make use of data for the years 1996 to 2009 and for NACE Revision 1.1 sectors 10-40 (manufacturing, mining and utilities). Of the variables collected in the CIP, those relevant for our purposes are the country of ownership, total revenue, share of revenue that is exported, employment, wage bill, expenditures on intermediates and imported share of intermediates.

In constructing our sample for analysis, we use the CIP as a basis. We drop firms that have a zero value for total revenue or zero employees in more than half of their years in the sample. We perform some recoding of firm identifiers to maintain the panel dimension of the data, e.g. in cases where ownership changes. As an export platform, Ireland is home to some large foreign-owned firms which may have very small shares of domestic value added relative to sales, and which do not participate at all in the domestic market. In order to focus on firms whose behavior we think likely to be more generally applicable, we drop firms which report in the CIP that less than 0.5% of their revenue comes from sales to the Irish market. Further details on the data and how we have cleaned it are provided in the Appendix.

[14]Multi-plant firms also fill in returns at the level of individual plants, but we work with the firm-level data since this is the level at which a match with the customs data can be performed.

5.1.2 Customs records

Our second source of data is customs records of Irish merchandise exports for the years 1996 to 2009. We also have import data, but we do not make use of it here. The value (Euros) and quantity (kg) of exports are aggregated to an annual frequency at the level of the tax ID, the Combined Nomenclature (CN) 8-digit product, and the export market (country). These data are matched by the Central Statistics Office to the CIP data using tax ID numbers, which are distinct from the firm and plant identifiers in the CIP, along with other confidential information. We have access only to the customs records that are matched with a firm ID number. This includes firms that are present in the CIP and a few firms that are not. We restrict attention to the records that match with CIP firms.

The CSO informs us that their ability to match the customs records to firm IDs is limited before 1997, and is best for the period 2000-2009. In the Appendix, we provide summary statistics which are consistent with this. Due to the possibility that a substantial fraction of exporters are not identified as such in earlier years, we restrict our baseline estimation of the export participation equation to the period 2000-2009.

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. For Ireland, the reporting threshold for intra-European exports (635,000 Euro per year) is much higher than the reporting threshold for extra-European exports (254 Euro per transaction).[15] Because our measure of participation is truncated for Intrastat markets, we do not make use of these markets in our baseline estimation of the participation equation. However we do make use of these markets to examine the behavior of long-time market participants.

Overall, we restrict attention to Ireland's largest export markets, and to Eastern European countries which joined the EU during the period of our sample. In the Appendix, we list the export destinations used in our empirical analysis. The full sample of countries (both Intrastat and Extrastat) accounts for on average 92% of the exports of our baseline sample of CIP firms over the period 1996-2009 while the Extrastat countries account for on average 22% of exports.

A final important feature of the customs data is that the 8-digit CN classification system changes every year. We concord the product-level data over time at the most disaggregated level possible following the approach of Pierce and Schott (2012) and van Beveren, Bernard

[15]Intra-European exports below the threshold are imputed based on VAT returns. We observe these imputed flows, but we do not make use of them.

and Vandenbussche (2012).[16]

5.1.3 Prodcom survey

Our third source of data is the Prodcom survey for the years 1996 to 2009. The Prodcom survey is an annual survey of the value and volume of all products manufactured by the firm and sold in the relevant year. The survey basis is all firms in the CIP, excluding some mining sectors. Products are classified at the 8-digit level according to the Prodcom classification. While in principle this survey covers almost all CIP firms, in practice, coverage is imperfect (on average 94% of total CIP turnover). We make use of the value data, but not the volume data from this survey. We use these data solely for the purpose of creating measures of tariffs at the firm-market level. In order to use the Prodcom data for this purpose, we concord the Prodcom classification with the HS6 classification for tariff data.

5.2 Summary statistics on firms and exports

All of the summary statistics we report here are for our baseline sample of CIP firms, excluding those reporting less than 0.5% of revenue originating from the Irish market.[17] Figure 1 plots real revenue and real export revenue (deflated by the PPI) for all the firms in the sample over the period 1996-2009. This period was a turbulent one for the Irish economy. A prolonged boom (1996-2007) was followed by very deep bust (2008-2009). However our empirical strategy controls for any cost implications of this turbulence by focusing on relative responses across different export markets.

Table 2 reports the share of total exports by market in our sample for a select set of export markets. These shares differ from those for all CIP firms (which are very similar to those based on published data on all merchandise exports), but the relative ranking of different markets is not very different. The main export markets for our firms are the US, the UK and other EU markets, principally the Euro Zone.

We replicate using our data many of the stylized facts about non-exporters and exporters which are familiar from micro data for other countries. Table 3 reports summary statistics on non-exporters and exporters separately for all years in our sample. Firms are classified as exporters if they are matched with positive exports from the customs data.[18] As is

[16] Van Beveren, Bernard and Vandenbussche show that after concording the product data appropriately, there is less product churn than naive calculations on un-concorded data would suggest.

[17] In the Appendix we report basic statistics without imposing this restriction.

[18] Throughout the period, this may exclude small exporters who export exclusively to Intrastat destinations, and exporters to all destinations who may be misclassified as non-exporters because of difficulty in matching

standard in both developed and developing countries, we find that exporters are bigger than firms which sell exclusively to the domestic market, in terms of both employees and total revenue. Export intensity is relatively high, a feature of the data which is typical of small open economies in Europe.

Table 4 reports transition rates into and out of exporting. As has been documented elsewhere, there is a good deal of persistence in export status. For individual markets, rates of entry are lower and rates of exit are higher, but exporting to a particular market is still persistent.[19] It is also worth noting that almost half of export entrants are "born global." Unfortunately, our empirical strategy does not capture this mode of export entry, as it relies on lagged variables at the firm level.

In Table 5, we report the share of total matched export revenue accounted for by export entrants and exiters on a year-on-year basis. We also report for incumbent firms the share of revenue accounted for by entering products, and for continuing firms the share of revenue accounted for by exiting products.[20] On average, on a year-on-year basis, 87% of export revenue is accounted for by incumbent firms selling incumbent products, while 84% of export revenue is accounted for by continuing firms selling continuing products. This is important, because it suggests that the elasticity of this category of revenue with respect to macro variables and tariffs is likely to be quantitatively very important in explaining the elasticity of aggregate exports with respect to these variables on a year-on-year basis.

In Table 6, we report average export revenue per entrant as a share of revenue per incumbent, and average export revenue per exiter as a share of revenue per continuer. We also report for incumbent exporters the average revenue per entering product as a share of the average revenue per incumbent product, and for continuing exporters, the average revenue per exiting product as a share of the average revenue per continuing product. Entrants and exiters export less on average than incumbents and continuers (this size difference would be even greater ignoring firms that are born global, and those that exit from the domestic market simultaneously with exiting from exporting). The size difference is even more marked at the product level.[21] These findings are similar to those of Eaton et al (2008) for Colombia, Bernard, Jensen, Redding and Schott (2009) for the US, and van Beveren, Bernard and

customs data to firm IDs.

[19] In the Appendix we report transition rates for a number of Extrastat markets, including the US.

[20] In the Appendix we report similar statistics, where instead of defining entry as entry into exporting, we define entry as entry into a particular export market (we define exit analogously). The relative shares are broadly similar under this alternative definition, and these are the numbers we refer to in the Introduction.

[21] As with Table 5, in the Appendix we report these statistics defining entry as entry into a particular export market rather than entry into exporting, and defining exit analogously. Defined in this way, entrants and exiters are even smaller relative to incumbents and continuers.

Vandenbussche (2012) for Belgium.[22] They imply that in order to generate large *short run* movements in aggregates due to entry and exit, there must be very large responses of entry and exit *rates* to shocks.

In the last row of Tables 5 and 6 we define an entrant as a firm which did not export in 2000, but which did export in 2009 (this includes firms which did not exist in 2000 but which exported in 2009), and an exiter as a firm which exported in 2000, but did not export in 2009 (including firms which did not exist in 2009). This allows us to get at the long-term impact of export entry and exit. We find (as is the case for other countries) that over longer horizons, entrants and new products account for a much bigger share of total export sales, consistent with entrants and entering products growing faster than incumbents and incumbent products. This suggests that the impact on aggregate exports of any participation responses to macro variables and tariffs will be bigger in the long run than in the short run.

5.3 Macro data

We include two macro variables in our regressions. The first is the real consumption exchange rate between Ireland and the relevant export market. The second is a measure of real local currency demand in the relevant export market. Real exchange rates are constructed using data on annual average nominal exchange rates and CPIs from the IMF's International Financial Statistics. The bulk of the variation in real exchange rates is driven by variation in nominal exchange rates (this is illustrated in the Appendix). Real demand in the target market is calculated as GDP in current local currency, less exports in current local currency, plus imports in current domestic currency, with this aggregate deflated by the relevant country's CPI. The National Accounts data are taken from the OECD where available, and otherwise from the World Bank's World Development Indicators. The CPIs are taken from International Financial Statistics. More details are in the Appendix.

In our baseline empirical specifications, the coefficients on real exchange rates and foreign demand are identified by variation across markets within a firm and a year in the deviation of the real exchange rate (or demand) from its in-sample average for that market across all sample years. In the Appendix, we use the raw real exchange rate and demand data to provide histograms of the identifying variation in these variables.

[22] Bernard et al (2014) point out that at least some of this is due to partial year effects for entrants who do not export for the full calendar year in their year of entry.

5.4 Tariff data

To construct measures of tariffs at the firm-market or firm-product-market level, we require raw data on tariffs at the level of the individual product and market. We take this tariff data from the WTO. Irish exporters do not face tariffs in EU, Eastern European or EFTA export markets in our sample period. For WTO member countries with which the EU does not have any special agreement, Irish exporters face the MFN tariff. In some non-EU non-EFTA markets (e.g. Turkey), Irish exporters face a preferential tariff offered to exporters from the EU. In all of the cases where Irish exporters face tariffs, conditional on Irish membership of the EU, the actual level of tariffs is determined by multilateral bargaining under the auspices of the WTO (MFN tariffs) or bilateral bargaining between the EU the relevant partner country (preferential tariffs). Given that Ireland is one small member of the EU, this makes us fairly confident that we can assume that tariffs are exogenous to our firms. We collect MFN tariff data for the following countries: Australia, Brazil, Canada, China, Hong Kong, India, Japan, Malaysia, Mexico, New Zealand, Russia, Saudi Arabia, South Africa, Thailand, United Arab Emirates and the US. We also collect information on preferential tariffs offered to EU countries by Turkey. For several of the developing countries in the sample, tariff data are not available for all of the sample years. Full details of the tariff data coverage are in the Appendix.

To make use of the tariff data, we concord it with our export and production data. The most disaggregated level at which a concordance is possible is the 6-digit level. The CN 6-digit and the Prodcom 6-digit correspond to the Harmonized System (HS) 6-digit classification, which is used by all countries as the basis for their tariff lines. We restrict attention to HS6 product-market-years for which there are no non-ad-valorem tariffs,[23] and for which there is no within-HS6 tariff variation.[24] For some purposes, HS6 categories must be concorded over time, as the HS6 classification changes in 2002 and 2007. For simplicity, when a concordance is required, we focus on HS6 categories for which there is a one-to-one match.

In our baseline empirical specifications, coefficients on tariffs are identified by variation across markets within a firm and a year in the deviation of tariffs from their in-sample average for that market across all sample years, and across all firms. In the Appendix, we use the raw tariff data at the product level to illustrate the fact that for tariffs, the key

[23] Non-ad-valorem tariffs affect incentives to export differently depending on the firm's export price, something we do not want to have to deal with.

[24] As a robustness check, we relax the requirement that there be no variation in tariffs within the HS6 category, and use the unweighted average of tariffs within HS6s where there is tariff variation.

source of variation is cross-sectional variation across markets, while there is very little time series variation. The main source of time-series variation in our tariff data is the gradual phasing-in over the period 1996-2005 of MFN tariff reductions agreed in the context of the Uruguay Round of the WTO, which was signed in 1995.

5.4.1 Firm-market- and firm-product-market-specific tariffs

In creating a tariff variable for the right hand side of participation equation, we assume that the tariffs on the products a firm *produces* provide a good guide to the effective degree of protection faced by the firm in deciding whether or not to participate in a market. Our baseline measure of the tariff variable for the participation equation is:

$$ptariff_t^{ik} = \sum_j \left(\frac{psh_t^{ij} + psh_{t-1}^{ij}}{2} \right) \ln\left(1 + \tau_t^{jk}\right) \tag{15}$$

where psh_t^{ij} is the share of product j in firm i's total production at date t, and τ_t^{jk} is the ad valorem tariff on good j in market k at time t. In robustness tests, we also make use of measures which weight using contemporaneous or lagged production shares. Note that we weight the log of the gross tariff rather than taking the log of the weighted gross tariff.

In constructing the tariff variable for the revenue equation at the firm-market level we use a similar approach, but instead of using production shares across different products to weight different tariffs, we use export revenue shares from the relevant market:

$$rtariff_t^{ik} = \sum_j \left(\frac{rsh_t^{ijk} + rsh_{t-1}^{ijk}}{2} \right) \ln\left(1 + \tau_t^{jk}\right) \tag{16}$$

where rsh_t^{ijk} is the share of product j in firm i's total export revenue from market k at date t. In robustness tests, we also make use of alternative tariff measures which weight using contemporaneous or lagged export revenue shares, and the average of current and lagged production shares. For robustness tests where we estimate in differences, analogous weighting schemes are used, but on differenced tariffs, $\Delta \ln\left(1 + \tau_t^{jk}\right)$.

In principle, no weighting should be needed for the tariffs in our firm-product-level revenue equation, (12), because each CN 8-digit category maps into at most one HS 6-digit category. In practice, the concordance of CN 8-digit categories over time (necessary in order to include the appropriate fixed effects in the regression) leads to cases where a concorded product category maps into one or more HS6 categories. In these cases, we use an analogous weighting

procedure to that used at the firm-market level.

6 Results

6.1 Participation

Table 7 reports the results from estimating equation (10). As already mentioned, we restrict our baseline estimation sample to Extrastat countries and to the period 2000-2009, as this is the sample for which participation is most accurately measured. We first control only for lagged participation, then introduce macro variables and tariffs, and finally allow responses to these variables to vary across the distribution of firms by allowing for increasingly richer sets of interactions. We note that especially for tariffs, but also for the macro variables, imposing the same response across all firm-market pairs (column (2)) would lead to misleading inference, as responses vary significantly with firm-market characteristics. As a result, we make the specification which interacts the macro variables and tariffs with lagged participation, lagged firm size and lagged sales in the market our baseline (column (5)).

To make it easier to interpret the estimates in column (5), Table 8 reports the marginal effects of the variables of interest on participation, by firm size and lagged participation status.[25] For incumbents, these marginal effects are evaluated at the median of lagged log revenue for the relevant size class. We find that participation responds to real exchange rates in the direction predicted by theory: a depreciation of the home real exchange rate vis-a-vis the destination market induces entry, while an appreciation induces exit. Similarly, higher real demand in a destination market induces entry, while lower demand induces exit. For potential entrants, the response of participation to these variables varies with size - larger firms are more likely to be induced to enter by a given shock than smaller firms. For incumbents, responses are more uniform across size classes. For tariffs, the responses are less precisely estimated than for the macro variables. The larger two classes of firms respond as theory predicts: higher tariffs reduce entry, and increase exit. For these medium and large firms, the point estimates of the exit responses are quantitatively much greater than exit responses to real exchange rates. For smaller firms, the point estimate of the response to tariffs goes in the "wrong" direction, though the coefficients are not significantly different from zero.

In order to put the magnitude of these marginal effects in context, it is useful to compare the impact of a 10% depreciation of the real exchange rate with the impact of an equivalent

[25] Marginal effects for specification (4) are reported in the Appendix.

reduction in the ad valorem tariff from 10% to 0. For firms with 100-249 employees a 10% depreciation of the real exchange rate increases the entry rate from a baseline of 3% (the average entry rate across markets in the sample) to 3.1% and reduces the exit rate from a baseline of 23% to 22.7%. This implies an increase in steady state participation of 0.5 percentage points, from 11.5% to 12%. For the same firms, a reduction in tariffs from 10% to zero increases the entry rate from the baseline of 3% to 3.3%, and reduces the exit rate from 23% to 20%. This implies an increase in steady state participation of 2.6 percentage points, from 11.5% to 14.2%. Interpreting this, in the long run, participation responds significantly more to tariffs than to real exchange rates, though in the short run (i.e. year-on-year) the impact is more modest.

We examine the robustness of these results along a number of dimensions. All of the robustness results described here are reported in the Appendix. We estimate our baseline specification both on broader samples (Intrastat as well as Extrastat) and on subsamples of the baseline (Irish-owned firms). We examine robustness to using different tariff measures (different weighting schemes and a liberal tariff measure that includes HS6 categories with intra-HS6 tariff variation). We allow for different configurations of fixed effects. The main thrust of the results (participation responds more to tariffs than to real exchange rates) is unchanged under these variations. Using the broader sample of both Intrastat and Extrastat markets results in smaller estimates of the response of participation to real exchange rates, and bigger estimates of the response of participation to tariffs, though as we have already noted, participation is less-well-measured for the Intrastat markets than it is for Extrastat markets. We also split the real exchange rate into its nominal exchange rate and price parts, and find some evidence that responses to prices are bigger than responses to nominal exchange rates.

Finally, we estimate separate entry and exit equations. The marginal effects from the entry equation are very similar to those from the baseline participation equation. The sample size in the exit equation is small (there are fewer potential exiters than potential entrants), and the results are noisy. We also estimate these entry and exit equations in differences, including lags of the independent variables to allow for dynamic responses. The results are noisy, and sensitive to the number of lags included and the exact sample period. This is not surprising given the short time dimension of our sample, and the lack of time-series variation in tariffs.

6.2 Revenue

Table 9 reports the results from estimating our baseline specification of the revenue equation (11). The baseline sample includes both Intrastat and Extrastat destinations (with the exception of Eastern European Accession countries), over the period 1996-2009. As mentioned above, to deal with the possibility of selection bias in the simplest way possible, we restrict the estimation sample to firm-market observations for which we observe at least 10 years of continuous participation, no in-sample exit, and if entry is observed in-sample, we drop the first two years of participation.

The estimated coefficient on the real exchange rate is 0.64 (standard error 0.15), the estimated coefficient on real foreign demand is 0.32 (standard error 0.16), while the estimated coefficient on the tariff variable is 2.42 (standard error 0.46). The signs of the estimated coefficients are all as predicted by theory. Depreciations of the home currency against that of the destination market, increases in foreign real demand, and reductions of tariffs in the destination market are all associated with higher export revenue. But perhaps surprisingly, the magnitude of the coefficients is squarely in the ballpark of estimates based on aggregate instead of firm-level data. The International Elasticity Puzzle persists at the level of firm-market revenue for the long-time market participants which account for the bulk of export revenue.

We examine the robustness of these results along a number of dimensions. All of the robustness results described here are reported in the Appendix. We estimate our baseline specification both on different samples (all firm-market observations with positive exports rather than just long-time market participants, and all firm-market observations *not* included in the baseline) as well as on a range of different subsamples of the baseline (Irish-owned firms, small and medium firms, large firms, Extrastat-only destinations). The results are not very sensitive to whether or not we include firm-market observations that we know are close to entry and exit thresholds. For all subsamples of the baseline, we estimate an elasticity of revenue with respect to real exchange rates that is less than one. However our finding that firm-market level export revenue responds more elastically to tariffs than to real exchange rates is most stark in the case of large firms (250+ employees), and less striking in the case of small and medium firms, and domestically-owned firms. For the Extrastat-only sample, the elasticity of revenue with respect to tariffs is not significantly different from zero. This may be due to the fact that there are relatively few firms which satisfy the continuous participation requirement in Extrastat markets, combined with the fact that this sample has less variation in the tariff variable than the baseline sample.

The results are not very sensitive to changes in the way we construct the tariff measure. When we split the real exchange rate into the nominal exchange rate and price parts, we find very similar coefficients on both nominal and real components of the real exchange rate, and the other results are not affected.

We also allow for different configurations of fixed effects. Perhaps surprisingly, the results when we include only firm, year and market effects (rather than firm-year and market effects) are very similar to the baseline. When we include firm-market and year effects, the coefficient on tariffs is still negative, but smaller than the coefficient on the real exchange rate in absolute value, and not significantly different from zero. This is consistent with the coefficient on tariffs in the baseline specification being identified mainly by cross-sectional rather than time series variation in tariffs.

To investigate this further, we estimate in differences, including lags of the independent variables to allow for dynamic responses. As in the case of participation, the results are noisy and sensitive to the number of lags included and the exact sample period. The point estimates of accumulated responses are generally consistent with what we find in the cross section, but the standard errors are very large.

Our sample period also allows us to estimate the response of export revenue to a non-tariff trade barrier. We exploit the fact that during the period of our sample, a number of Eastern European countries acceded to the EU, the first wave in May 2004, the second wave in January 2007. While accession was not associated with tariff changes (Irish exporters faced zero tariffs in these markets throughout the sample period), it was associated with the elimination of customs controls between Ireland and these destination markets. Because the reclassification of the relevant destinations from Extrastat to Intrastat (and hence how precisely participation is measured) coincided with EU accession, we do not examine the effect of this change on participation. However we can examine how eliminating customs controls affects the export revenue of firms who export continuously to these markets.

To do so, we restrict the sample to destinations where tariffs are always zero (EU, Eastern Europe and EFTA). In addition to the standard macro variables, we include a dummy for EU membership. The results are reported in Table 10. The point estimate of the coefficient on this variable is 0.13 (standard error 0.08). Though imprecisely estimated, the coefficient is consistent with a substantial (i.e. 13%) increase in the export revenue of continuing participants from these markets. In the Appendix, we report that this effect is bigger and statistically significant (coefficient of 0.21, standard error 0.09) for small and medium firms, and smaller and much less precisely estimated for bigger firms. These results are consistent

with a strong response of revenue to non-tariff as well as tariff barriers for some firms, in contrast to anemic responses to real exchange rates.

6.3 Product-level analysis

In Table 11 we report the results from estimating equation (12). The baseline sample includes continuing products (i.e. products sold in the current year and in the previous year) for the firm-market observations used in the baseline estimation of (11). The estimated coefficient on the real exchange rate is 0.80 (standard error 0.16), the estimated coefficient on real demand in the destination market is 0.15 (standard error 0.17), while the estimated coefficient on the tariff is -2.65 (standard error 0.74). For continuing products within a firm and a market, revenue responds to real exchange rates and tariffs much as we estimate firm-market revenue to respond.

We examine robustness of these results along several dimensions. These results are reported in the Appendix. We apply a similar continuous presence requirement at the firm-product-market level as the requirement we impose at the firm-market level. We include firm-product and year fixed effects, rather than firm-product-year fixed effects, and firm-year and product fixed effects rather than firm-product-year fixed effects. The results are qualitatively similar for all of these variations. As in the case of firm-market revenue, we estimate in differences, including lags of the independent variables to allow for dynamic responses. The results from this exercise are very noisy.

In Table 12 we report the results from estimating (13) and (14). The sample is restricted to firm-markets included in the baseline estimation of (11). We report results both for the case where we use the contemporaneous value of the exit share in revenue as the dependent variable, as well as for the case where we use the lagged share. The estimates imply that the share of entering products in revenue decreases in response to both real exchange rate depreciations, and tariff reductions. The same is true for the share of exiting products in revenue, irrespective of whether the current or the lagged value is used as the dependent variable. The (absolute value of the) response to tariffs is greater for the entry and current exit shares. But the direction of these results is somewhat ambiguous. They suggest that new products account for a *smaller* share of revenue in markets and years where real exchange rates and tariffs are favorable than in markets and years where they are unfavorable. This is contrary to what the sunk cost model would predict, and contrary to our finding that exit shares are lower when real exchange rates and tariffs are favorable.

Taken together with the results on the response of product-level revenue, the results on

shares of entering and exiting products suggest that product entry and exit do not play a big role in accounting for the within-firm-market elasticities we document in Table 9.

7 Discussion

Having estimated the responses of different margins of exporting to real exchange rates and tariffs, we would like to know what is the relative importance of these different margins in explaining the International Elasticity Puzzle. We address this question with a simple back-of-the-envelope calculation. This approach is motivated by the fact that our empirical strategy prioritizes clean identification over estimating responses along all dimensions for all firms. As a result, we cannot perform an exact decomposition of the responses of aggregate exports to real exchange rates and tariffs into the responses of different margins.

Suppose we have an economy with a constant number of firms (i.e. ignore firm birth and death). We assume there is one potential export market. In steady state, the number of entrants into exporting is equal to the number of exporters, so the number of firms participating in the export market is constant. We assume that all exporters have identical export revenue conditional on exporting. We ignore the within-firm product margin. We then use our estimates of participation and revenue responses to examine the impact on aggregate exports of a permanent real depreciation of 10%, where the real exchange rate changes from 1 to 1.1, and the impact of a permanent reduction in an across-the-board ad valorem tariff from 10% to 0. In the absence of an elasticity puzzle, the response of aggregate exports to these two shocks should be the same. Implicitly, throughout this exercise, we hold everything else constant (including costs).

To perform this calculation, we need to choose baseline export entry and export exit rates (which imply a baseline steady state participation rate). We use the averages over the sample period of the entry and exit rates for medium-sized firms (100-249 employees) in the Extrastat destinations used in our baseline estimation of the participation equation. The average entry rate for these firms is 3% (this ignores entry by "born globals") and the exit rate is 23%, implying a steady state participation rate of 11.5%.

To calculate the impact of the changes in real exchange rates and tariffs on participation in this exercise, we first multiply our shocks by the estimated marginal effects for medium-sized firms (100-249 employees) reported in Table 8, and add the implied changes to the baseline entry and exit rates. The 10% real exchange rate depreciation increases the entry rate by 0.1 percentage points (to 3.1%) and reduces the exit rate by 0.3 percentage points (to

22.7%), while the tariff change increases the entry rate by 0.3 percentage points (to 3.3%) and reduces the exit rate by 3 percentage points (to 20%). We first calculate the short-run impact of the shocks year-by-year, by iteratively applying the changed entry and exit rates to lagged participation and non-participation shares. We then calculate the long-run impact by calculating the steady state share of exporting firms implied by the new entry and exit rates.

To calculate the impact of the shocks on export revenue conditional on participation, we apply the estimated coefficients (β) from Table 11 to the expression $R_{after} = (1.1)^{\beta} R_{baseline}$, where we normalize $R_{baseline}$ to 1. For real exchange rates, $\beta = 0.64$, while for tariffs, $\beta = 2.42$ (we compare an increase in the real exchange rate with a reduction in tariffs). In this simple exercise, we assume no dynamics in revenue responses.[26]

We combine both participation and revenue responses (by multiplying the new participation rate by R_{after}) to calculate the impact on aggregate exports, and use this to back out an implied elasticity of aggregate exports with respect to the relevant shock, calculated as $[\ln(EX_{after}) - \ln(EX_{before})] / [\ln(1.1) - \ln(1)]$. The results of this exercise are reported in the first two columns of Table 13. The responses estimated from the micro data imply a difference between the response of aggregate export revenue to real exchange rates and the response to tariffs even greater than what we estimate using aggregate data for the sample period (see Table 1).

To understand the importance of the participation margin in generating this result, we shut down this margin of adjustment altogether. When participation does not change in response to a shock, the response of aggregate exports is exactly equal to the response of the export revenue of individual firms. Then our estimates imply an elasticity of aggregate exports with respect to the real exchange rate of 0.64, and an elasticity of aggregate exports with respect to tariffs of 2.42. Hence, the participation margin is not crucial in order to generate differences in elasticities on the order of magnitude of those observed in the aggregate data. Moreover, since our estimates imply that the export participation margin makes aggregate exports respond *more* to real exchange rates rather than *less*, the weak response of export revenue conditional on participation to the real exchange rate is crucial in order to generate a low elasticity of aggregate exports with respect to this variable. In this sense, we think the key to the International Elasticity Puzzle lies not in export participation, but in how export revenue conditional on participation responds to the different shocks.

This obviously begs the question as to why (controlling for costs and foreign demand)

[26] As we show in Table 6, entrants are smaller than incumbents, and exiters are smaller than continuing firms. Incorporating this fact into this simple exercise does not much change the results.

export revenue at the firm-market level should respond differently to real exchange rates and tariffs. An obvious candidate, building on the work of Ruhl (2008), is some form of forward-looking behavior and lumpy adjustment which affects export revenue conditional on participation. This, combined with the different processes for macro variables and tariffs in the data could potentially rationalize different revenue responses to the two shocks. Because an explanation of this type would give the right comparative static at all time horizons, and would apply to pairs of export markets as well as to the home relative to the foreign market, it would be more likely to explain the facts we document than the explantions proposed by Arkolakis, Eaton and Kortum (2012) and Feenstra, Russ and Obstfeld (2012).

Within this class of explanations, one obvious possibility is menu costs of changing prices. Costs of this type create an inaction region with respect to shocks, where the size of the inaction region varies with the persistence and volatility of the process from which the shock is drawn. It could be that firms react to tariff changes but not to real exchange rate movements by changing prices, and that this plays some role in explaining the revenue responses we document. There is considerable evidence out there that prices are indeed sticky, and hence this is an explanation deserving of future research.

An alternative (and possibly complimentary) possibility is that there are costs of adjustment on the quantity side rather than the price side. For example, suppose that firms can engage in costly and irreversible investment in a capital-like customer base that shifts demand conditional on prices not just today but in the future. Under irreversibility, if real exchange rates are less persistent or more volatile than tariffs, an innovation in the real exchange rate today will optimally lead to a smaller investment response than the equivalent innovation in tariffs. As a result, revenue will respond less to the real exchange rate than to tariffs. In the Appendix, we describe a model with these features. This type of explanation is potentially appealing, both because recent research in other areas of macro and international macro has proposed slow-moving market shares as explanations for a variety of phenomena (e.g. Drozd and Nosal (2012), Foster, Haltiwanger and Syverson (2013) and Gourio and Rudanko (2013)), and because Arkolakis (2010) shows that a static model with some of these features can help explain trade facts. We leave quantitative exploration of the performance of such a model in explaining the International Elasticity Puzzle for future research.

8 Conclusion

The International Elasticity Puzzle refers to the fact that standard models incorporating international trade need to be parameterized differently in order to match the comovement of exports with real exchange rates and the response of exports to trade liberalizations. Our goal in this paper is to provide guidance as to how best to modify standard models in order to be able to match both sets of facts with a single set of parameters. In order to do this, we estimate the responses of export participation, export revenue and the product dimension of exporting at the firm-market level to both real exchange rates and tariffs. We exploit the structure of our data to control for marginal cost, identifying responses from within-firm-year cross-market variation in real exchange rates and tariffs.

We find that the weak response of export revenue of long-time market participants to real exchange rates is key to the behavior of aggregate exports. At the same time, export participation also responds less to real exchange rates than to tariffs, but this alone cannot explain the International Elasticity Puzzle. Finally, responses of product entry and exit of long-time market participants do not play an important role. Our results imply that for a model to successfully account for the puzzle, it must capture the fact that export revenue conditional on participation responds more elastically to tariffs than to real exchange rates. While further work is clearly necessary, 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), 289-336.

[2] Arkolakis, C. (2010), "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Politicial Economy* 118 (6), 1151-1199.

[3] Arkolakis, C., J. Eaton and S. Kortum (2012), "Staggered Adjustment and Trade Dynamics," mimeo.

[4] Atkeson, A. and A. Burstein (2008), "Pricing to Market, Trade Costs and International Relative Prices," *American Economic Review* 98 (5), 1998-2031.

[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 (2012), "How do Different Exporters React to Exchange Rate Changes?" *Quarterly Journal of Economics* 127 (1), 437-492.

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

[9] Bernard, A., B. Jensen, S. Redding and P. Schott (2009), "The Margins of US Trade," *American Economic Review: Papers & Proceedings*, 99 (2), 487-493.

[10] Bernard, A., R. Massari, J.-D. Reyes and D. Taglioni (2014), "Exporter Dynamics, Firm Size and Growth, and Partial Year Effects," mimeo.

[11] Bernard, A., S. Redding and P. Schott (2011), "Multi-Product Firms and Trade Liberalization," *Quarterly Journal of Economics* 111 (4), 681-709.

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

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

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

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

[16] Drozd, L. and J. Nosal (2012), "Understanding International Prices: Customers as Capital," *American Economic Review* 102 (1), 364-395.

[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] Feenstra, R., M. Obstfeld and K. Russ (2012), "In Search of the Armington Elasticity," mimeo.

[19] Foster, L., J. Haltiwanger and C. Syverson (2013) "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 (2013), "Customer Capital," *Review of Economic Studies*, forthcoming.

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

[23] Lileeva, A. and D. Trefler (2010), "Improved Access to Foreign Markets Raises Plant-level Productivity... For Some Plants," *Quarterly Journal of Economics* 125 (3), 1051-1099.

[24] Morales, E., G. Sheu and A. Zahler (2014), "Gravity and Extended Gravity: Using Moment Inequalities to Estimate Model of Export Entry," mimeo.

[25] Pierce, J. and P. Schott (2012), "A Concordance Between Ten-Digit US Harmonized System Codes and SIC/NAICS Product Classes and Industries," *Journal of Economic and Social Measurement* 37 (1-2), 61-96.

[26] 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.

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

[28] Van Beveren, I., A. Bernard and H. Vandenbussche (2012), "Concording EU Trade and Production Data over Time," mimeo.

Table 1: Aggregate revenue equation

	coeff	s.e.
rer_t^k	0.42	(0.15)**
dem_t^k	0.46	(0.11)**
tau_t^{jk}	-1.57	(0.25)**
Market f.e.	yes	
Prod-year f.e.	yes	
# prod-mkt-years	271924	
# prod-years	50852	
# products	5721	
R^2	0.58	
R^2-adj	0.48	

Notes: Estimation method is OLS. Dependent variable is log US dollar export revenue at the level of the product and the market. Included markets are Australia, Austria, Belgium, Brazil, Bulgaria, Canada, China, Czech Republic, Denmark, Estonia, Finland, France, Germany, Hong Kong, Hungary, India, Italy, Japan, Latvia, Lithuania, Malaysia, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Romania, Russia, Saudi Arabia, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland, Thailand, Turkey, United Arab Emirates, United Kingdom, United States. For tariff-subject markets we include only HS6 categories where there are only ad valorem tariffs and no tariff variation within the HS6 category. HS6s are concorded between 1996-2001 and 2002-2006 or between 2007-2009 and 2002-2006 only if there is a 1-1 concordance. Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 2: Shares of exports in our baseline sample by destination

	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
Australia	0	1	1	2	1	2	2	1	2	2	2	1	1	1
Canada	0	0	0	0	0	0	0	0	0	0	1	1	1	1
China	0	0	0	0	0	0	0	0	1	0	0	0	1	1
Denmark	2	2	1	1	1	1	1	1	1	1	1	1	1	1
Japan	2	3	3	3	5	4	2	2	2	2	2	1	1	1
Norway	1	1	1	1	1	1	1	1	1	1	1	0	0	0
Sweden	2	2	2	2	2	2	2	2	1	2	2	1	1	1
Switzerland	1	1	2	2	1	1	1	3	10	1	1	1	13	1
UK	33	32	28	30	29	29	34	30	26	32	31	35	31	34
US	7	7	9	8	8	8	8	8	11	7	9	18	11	13
Euro 9	34	33	38	39	37	37	38	41	34	40	38	30	28	32
E. Europe	1	1	1	1	2	1	1	1	1	1	2	2	2	3

Notes: Source: CSO and authors' calculations. Our baseline CIP sample excludes firms which report less than 0.5% of revenue from the Irish market. We use only customs export records which match to a CIP firm in our baseline sample. The Euro 9 includes Austria, Belgium, Finland, France, Germany, Italy, Netherlands, Portugal and Spain (Greece and Luxembourg are excluded because Irish trade with these countries is negligible). E. Europe refers to Eastern European countries which joined the EU in 2004 and 2007: Bulgaria, Czech Republic, Estonia, Latvia, Lithuania, Poland, Romania, Slovakia and Slovenia.

Table 3: Summary statistics on exporters and non-exporters

	# firms		Avg # employees		Avg revenue		Avg export %	Avg # dest
	All	Exporters	Nonex.	Exporters	Nonex.	Exporters	Exporters	Exporters
1996	4101	909	92	161	5877	22074	42	8
1997	4234	1533	94	132	7586	16518	41	6
1998	4222	1587	94	133	7396	18738	40	6
1999	4304	1659	92	132	7411	18408	38	6
2000	4582	1914	91	124	7967	17219	35	6
2001	4517	1957	90	123	8125	16296	36	6
2002	4706	1967	89	119	10004	16752	35	6
2003	4704	1990	88	120	10321	20416	36	6
2004	4411	1976	88	121	12297	22977	36	6
2005	4126	1824	89	121	14809	23805	37	6
2006	4305	1871	86	121	14663	25382	36	6
2007	5130	1897	83	118	12964	22569	33	5
2008	5221	1868	79	116	8733	21266	32	5
2009	4807	1827	77	115	8940	18943	33	5

Notes: Statistics are for our cleaned dataset of CIP firms, which excludes firms reporting less than 0.5% of revenue from sales in the Irish market. Firms are defined as exporters if they are matched to positive exports from the customs data. Export revenue is export revenue from Customs data. Export share is calculated as total exports from the customs data divided by sales reported in the CIP. Values greater than 100 are replaced by 100. Source: CSO and authors' calculations.

Table 4: Transitions of firms into and out of exporting

Year-t status	exports			no exports			"born global"
Year-t+1 status	no exports	exports	exit	no exports	exports	exit	share of entrants
1996-97	0.05	0.93	0.02	0.72	0.18	0.09	0.15
1997-98	0.07	0.91	0.02	0.86	0.04	0.10	0.41
1998-99	0.09	0.88	0.02	0.81	0.06	0.13	0.44
1999-00	0.05	0.91	0.05	0.77	0.09	0.14	0.42
2000-01	0.05	0.90	0.05	0.87	0.06	0.07	0.31
2001-02	0.06	0.88	0.06	0.86	0.04	0.10	0.55
2002-03	0.07	0.85	0.07	0.83	0.06	0.10	0.43
2003-04	0.06	0.84	0.10	0.75	0.07	0.18	0.33
2004-05	0.07	0.82	0.11	0.81	0.05	0.14	0.42
2005-06	0.08	0.82	0.10	0.84	0.05	0.11	0.66
2006-07	0.09	0.82	0.09	0.85	0.03	0.11	0.77
2007-08	0.06	0.82	0.12	0.78	0.04	0.18	0.55
2008-09	0.08	0.84	0.09	0.80	0.05	0.14	0.30
average	0.07	0.86	0.07	0.81	0.07	0.12	0.44

Notes: Table reports share of firms of year-t status ending up in each category of year-t+1 status. Statistics are for our cleaned dataset of CIP firms, which excludes firms reporting less than 0.5% of revenue from sales in the Irish market. Firms are defined as exporters if they are matched to positive exports from the customs data. Exiters are firms which do not appear in the CIP under the same firm id in the next period. Firms are born when they do not appear in the CIP under the same firm id in the previous period. Source: CSO and authors' calculations.

Table 5: Share of matched exports accounted for by entering and exiting firms and products

Firm status	Entrant	Exiter	Incumbent	Continuer
Prod. status	All	All	Entering	Exiting
1997-98	0.17	0.04	0.03	0.04
1998-99	0.08	0.24	0.04	0.03
1999-00	0.06	0.11	0.03	0.03
2000-01	0.10	0.08	0.03	0.04
2001-02	0.07	0.15	0.03	0.02
2002-03	0.11	0.04	0.03	0.03
2003-04	0.15	0.09	0.04	0.02
2004-05	0.06	0.23	0.01	0.02
2005-06	0.09	0.05	0.03	0.03
2006-07	0.04	0.29	0.03	0.04
2007-08	0.24	0.15	0.03	0.03
2008-09	0.08	0.24	0.02	0.02
avg	0.10	0.14	0.03	0.03
2000-09	0.39	0.59	0.12	0.16

Notes: The first two columns report the share of matched exports accounted for by entering and exiting firms, where entry and exit are defined as entry into exporting and exit from exporting. For the period 2000-09, an entrant is a firm which exported in 2009 but not in 2000. Conversely, an exiter is a firm which exported in 2000 but not in 2009. In this case, entrant shares are as a share of 2000 export revenue, and exiter shares are as a share of 2009 export revenue. The second two columns report for incumbent exporters the share of their export revenue accounted for by entering products (products they did not export in the previous period), and for continuing exporters the share of their exports accounted for by exiting products (products they do not export in the subsequent period). For the period 2000-09, incumbents and continuing firms are firms that export both in 2000 and in 2009. For these firms, an entering product is a product the firm does not export in 2000 but does export in 2009, while an exiting product is a product the firm exports in 2000 but not in 2009. The sample for product entry and exit is restricted to incumbent and continuing exporters, because by definition, all exports of entering firms are entering products, and all exports of exiting firms are exiting products. Source: CSO and authors' calculations.

Table 6: Size of entering and exiting firms and products relative to incumbents and continuers

Firm status	Entrant	Exiter	Incumbent	Continuer
Prod. status	All	All	Entering	Exiting
1997-98	1.09	0.32	0.04	0.04
1998-99	0.37	1.62	0.03	0.03
1999-00	0.20	0.78	0.03	0.03
2000-01	0.65	0.60	0.03	0.05
2001-02	0.44	1.04	0.03	0.03
2002-03	0.53	0.22	0.03	0.03
2003-04	0.81	0.44	0.05	0.02
2004-05	0.41	1.20	0.02	0.03
2005-06	0.36	0.19	0.03	0.04
2006-07	0.16	1.60	0.04	0.05
2007-08	1.34	0.71	0.05	0.04
2008-09	0.39	1.30	0.03	0.04
avg	0.56	0.84	0.03	0.03
2000-09	0.60	1.25	0.09	0.11

Notes: The first two columns report the ratio of average entrant export revenue to average incumbent export revenue and the ratio of average exiter export revenue to average continuer export revenue, where entry and exit are defined as entry into exporting and exit from exporting. For the period 2000-09, an entrant is a firm which exported in 2009 but not in 2000. Conversely, an exiter is a firm which exported in 2000 but not in 2009. In this case, average entrant export revenue and average incumbent export revenue refer to 2009, while average exiter export revenue and average continuer export revenue refer to 2000. The second two columns report for incumbent firms the ratio of average export revenue per entering product to average export revenue per incumbent product, and for continuing firms, the ratio of average export revenue per exiting product to average export revenue per continuing product. For the period 2000-09, incumbents and continuing firms are firms that export in both 2000 and in 2009. For these firms, an entering product is a product the firm does not export in 2000 but does export in 2009, while an exiting product is a product the firm exports in 2000 but not in 2009. The sample for product entry and exit is restricted to incumbent and continuing exporters, because by definition, all exports of entering firms are entering products, and all exports of exiting firms are exiting products. Source: CSO and authors' calculations.

Table 7: Participation

	(1)		(2)		(3)		(4)		(5)	
	coeff	s.e.	coeff	s.e.	coeff	s.e.	coeff	s.e.	coeff	s.e.
X^{ik}_{t-1}	0.59	(0.01)**	0.58	(0.01)**	0.13	(0.05)**	0.20	(0.05)**	0.01	(0.07)
$X^{ik}_{t-1} * rev^{ik}_{t-1}$									0.14	(0.01)**
rer^k_t			0.01	(0.00)**	0.01	(0.00)**	0.00	(0.00)	0.00	(0.00)
$emp2^i_{t-1} * rer^k_t$							0.01	(0.00)**	0.01	(0.00)**
$emp3^i_{t-1} * rer^k_t$							0.02	(0.00)**	0.02	(0.00)**
$X^{ik}_{t-1} * rer^k_t$					0.03	(0.00)**	0.02	(0.00)**	0.03	(0.01)**
$X^{ik}_{t-1} * emp2^i_{t-1} * rer^k_t$							0.01	(0.01)*	-0.00	(0.00)
$X^{ik}_{t-1} * emp3^i_{t-1} * rer^k_t$							-0.00	(0.00)	-0.01	(0.01)*
$X^{ik}_{t-1} * rev^{ik}_{t-1} * rer^k_t$									-0.01	(0.00)**
dem^k_t			0.02	(0.00)**	0.02	(0.00)**	0.01	(0.00)**	0.01	(0.00)**
$emp2^i_{t-1} * dem^k_t$							0.00	(0.00)**	0.00	(0.00)**
$emp3^i_{t-1} * dem^k_t$							0.01	(0.00)**	0.01	(0.00)**
$X^{ik}_{t-1} * dem^k_t$					0.02	(0.00)**	0.02	(0.00)**	0.02	(0.00)**
$X^{ik}_{t-1} * emp2^i_{t-1} * dem^k_t$							-0.00	(0.00)	-0.00	(0.00)
$X^{ik}_{t-1} * emp3^i_{t-1} * dem^k_t$							-0.00	(0.00)	-0.01	(0.00)**
$X^{ik}_{t-1} * rev^{ik}_{t-1} * dem^k_t$									-0.00	(0.00)**
tau^{ik}_t			0.01	(0.01)	0.01	(0.01)	0.03	(0.00)**	0.02	(0.01)**
$emp2^i_{t-1} * tau^{ik}_t$							-0.05	(0.01)**	-0.06	(0.01)**
$emp3^i_{t-1} * tau^{ik}_t$							-0.04	(0.02)*	-0.05	(0.02)**
$X^{ik}_{t-1} * tau^{ik}_t$					-0.25	(0.07)**	0.06	(0.17)	0.14	(0.17)
$X^{ik}_{t-1} * emp2^i_{t-1} * tau^{ik}_t$							-0.33	(0.18)*	-0.50	(0.18)**
$X^{ik}_{t-1} * emp3^i_{t-1} * tau^{ik}_t$							-0.39	(0.19)**	-0.36	(0.20)**
$X^{ik}_{t-1} * rev^{ik}_{t-1} * tau^{ik}_t$									0.05	(0.02)**
Market f.e.	yes		yes		yes		yes		yes	
Firm-year f.e.	yes		yes		yes		yes		yes	
R^2	0.63		0.63		0.64		0.64		0.65	
R^2-adj	0.59		0.59		0.59		0.59		0.61	
N	334880		334880		334902		334880		334880	

Notes: Estimation method is OLS. Dependent variable is an indicator for participation. Tariff measure is constructed using average production weights. Sample consists of all firm-mkt-years in the Extrastat sample 2000-2009 with positive current and lagged sales in the Irish market, and for which the independent variables are available (Eastern European Accession countries are excluded because they switch from Extrastat to Intrastat). The omitted category is potential entrants with 1-99 employees in the previous period. $emp2^i_{t-1}$ indicates that the firm had 100-249 employees and $emp3^i_{t-1}$ that the firm had 250+ employees in the previous period. Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 8: Marginal effects from participation equation

		rer		demand		tariff		participation	participation
Status	Employees	coeff	se	coeff	se	coeff	se	all mkts	US
Potential entrants	1-99	0.00	(0.00)	0.01	(0.00)**	0.02	(0.01)**	0.01	0.02
	100-249	0.01	(0.00)**	0.02	(0.00)**	-0.03	(0.01)**	0.03	0.07
	250+	0.02	(0.00)**	0.03	(0.00)**	-0.02	(0.02)	0.06	0.15
Incumbents	1-99	0.03	(0.01)**	0.03	(0.00)**	0.26	(0.16)	0.69	0.74
	100-249	0.03	(0.01)**	0.03	(0.01)**	-0.30	(0.11)**	0.77	0.84
	250+	0.04	(0.01)**	0.04	(0.01)**	-0.14	(0.11)	0.79	0.88

Notes: Marginal effects are calculated based on the estimates reported in column (5) of Table 7. For incumbents, marginal effects are evaluated at the median lagged revenue for the relevant size group. The final columns report the average in-sample participation rate across all firms, markets and years for the relevant category of firms, and the average in-sample participation rate in the US market across all firms and years for the relevant category of firms. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 9: Revenue

	coeff	s.e.
rer_t^k	0.64	(0.15)**
dem_t^k	0.32	(0.16)**
tau_t^{ik}	-2.42	(0.46)**
Market f.e.	yes	
Firm-year f.e.	yes	
# firm-mkt-years	49037	
# firm-years	7858	
# firms	931	
R^2	0.57	
R^2-adj	0.49	

Notes: Estimation method is OLS. Dependent variable is log Euro revenue at the level of the firm and the market deflated by the Irish CPI. Tariff measure is constructed using average export weights. Sample consists of all firm-mkt pairs in Intrastat and Extrastat destinations (except Eastern Europe) for which continuous participation is observed throughout the sample. Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 10: Revenue: EU Accession

	coeff	s.e.
rer_t^k	0.96	(0.18)**
dem_t^k	0.42	(0.16)**
eu_t^k	0.13	(0.08)
Market f.e.	yes	
Firm-year f.e.	yes	
# firm-mkt-years	41331	
# firm-years	7066	
# firms	868	
R^2	0.62	
R^2-adj	0.54	

Notes: Estimation method is OLS. Dependent variable is log Euro revenue at the level of the firm and the market deflated by the Irish CPI. The EU dummy is set equal to 1 for Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia in 2004, and for Bulgaria and Romania in 2007. Sample consists of all firm-mkt pairs in EFTA for which continuous participation is observed throughout the sample. Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 11: Product revenue

	coeff	s.e.
rer_t^k	0.80	(0.16)**
dem_t^k	0.15	(0.17)
tau_t^{ijk}	-2.65	(0.74)**
Market f.e.	yes	
Firm-prod-year f.e.	yes	
# firm-mkt-years	77483	
# firm-years	6258	
# firms	794	
R^2	0.76	
R^2-adj	0.59	

Notes: Estimation method is OLS. Dependent variable is log Euro revenue at the level of the firm, the concorded 8-digit product and the market deflated by the Irish CPI. Tariff measure is constructed using export weights where the product concordance is across HS6 categories. Sample consists of all continuing firm-product-mkt triplets for firm-market pairs included in the baseline estimation of the firm-market revenue equation. Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 12: Response of product entry and exit to macro variables and tariffs

Dependent variable	$shenter_t^{ik}$		$shexit_t^{ik}$		$shexit_{t-1}^{ik}$	
	coeff	s.e.	coeff	s.e.	coeff	s.e.
rer_t^k	-0.10	(0.03)**	-0.07	(0.04)*	-0.07	(0.03)**
dem_t^k	-0.01	(0.04)	-0.02	(0.04)	-0.05	(0.03)
tau_t^{ik}	0.21	(0.10)**	0.36	(0.11)**	0.08	(0.12)
Market f.e.	yes		yes		yes	
Firm-year f.e.	yes		yes		yes	
# firm-mkt-years	39040		37110		32904	
# firm-years	6509		6251		6026	
# firms	846		823		767	
R^2	0.44		0.42		0.43	
R^2-adj	0.33		0.30		0.30	

Notes: Estimation method is OLS. $shenter_t^{ik}$ is the share of entering products in export revenue, where revenue and entry are defined at the firm-market level. $shexit_t^{ik}$ is the share of exiting products in export revenue, where exit is defined at the firm-market level. $shexit_{t-1}^{ik}$ is the lag of $shexit_t^{ik}$. Sample consists of all firm-market pairs included in the baseline estimation of the firm-market revenue equation for which the dependent variable can be calculated (one year is lost relative to the baseline estimation in the first two columns; two years are lost in the final column). Standard errors are clustered at the level of the year and the market. ** indicates significance at the 5% level. * indicates significance at the 10% level.

Table 13: Elasticity of aggregate exports with respect to shocks implied by micro responses

Participation Revenue	RER response RER response	Tariff response Tariff response	None RER response	None Tariff response
1 year	0.75	2.96	0.64	2.42
2 years	0.83	3.36	0.64	2.42
3 years	0.89	3.65	0.64	2.42
4 years	0.94	3.87	0.64	2.42
Long run	1.07	4.57	0.64	2.42

Notes: This table reports the elasticity of aggregate exports with respect to real exchange rates and tariffs implied by a back-of-the-envelope calculation based on the estimated responses from the data and average export participation and churn in the sample period. Details of the calculation are in the text.

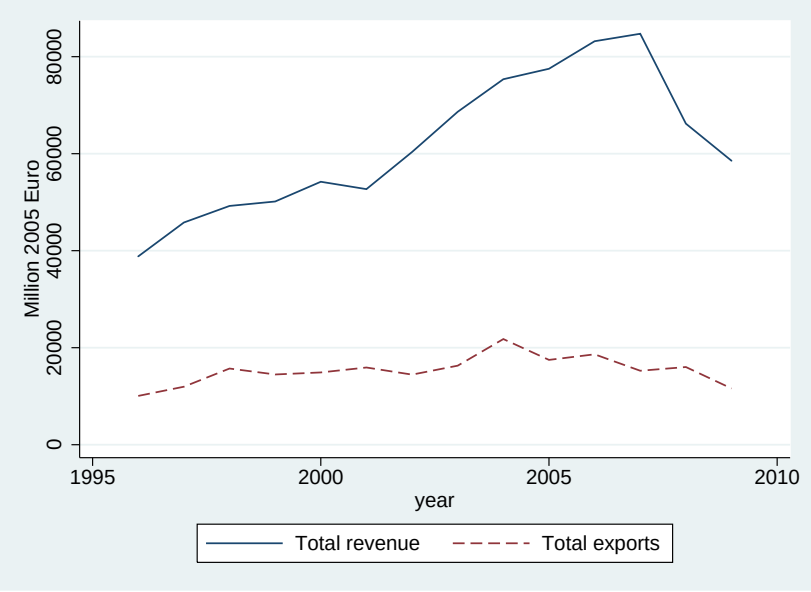

Figure 1: Total revenue and total exports

Notes: Figure shows total real revenue for CIP firms in our baseline sample, and total real exports matched to these firms from customs data. Source: CSO and authors' calculations.