

Exporters and Shocks

Doireann Fitzgerald
Federal Reserve Bank of Minneapolis

Stefanie Haller
School of Economics, University College Dublin

Staff Report 549
May 2017

DOI: <https://doi.org/10.21034/sr.549>

Keywords: Tariffs; Real exchange rates; International elasticity puzzle

JEL classification: F14, F41

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Exporters and Shocks*

Doireann Fitzgerald[†] Stefanie Haller[‡]

May 2017

Abstract

We use micro data for Ireland to estimate how export participation and the export revenue of incumbent exporters respond to tariffs and real exchange rates. Both participation and revenue, but especially revenue, are more responsive to tariffs than to real exchange rates. Our estimates translate into an elasticity of aggregate exports with respect to tariffs of between -3.8 and -5.4, and with respect to real exchange rates of between 0.45 and 0.6, consistent with estimates in the literature based on aggregate data. We argue that forward-looking investment in customer base combined with the fact that tariffs are much more predictable than real exchange rates can explain why export revenue responds so much more to tariffs.

1 Introduction

Many important economic policy questions require an explicit comparison of the response of exports to trade policy and the response of exports to real exchange rates. The debate on whether “currency manipulation” is equivalent to trade protection through tariffs or trade

*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 data sets 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, Adrian Corcoran and Matt Shapiro. 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.” This paper has benefited from extensive discussions with Yaniv Yedid-Levi. We are also grateful for comments and suggestions from Costas Arkolakis, Ina Jäkel, Richard Rogerson, Kim Ruhl and Jim Tybout. An online appendix to the paper is available at www.doireann.com. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[†]Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, Minneapolis MN 55401, doireann.fitzgerald@gmail.com

[‡]School of Economics, University College Dublin, Dublin 4, Ireland, stefanie.haller@ucd.ie

promotion through export subsidies is one example.¹ The debate on whether exchange rate movements might offset the impact of border adjustment of corporate taxation on U.S. trade is another.²

Motivated by these policy questions, we use firm and customs micro data to document how export entry, exit and the export revenue of incumbent exporters respond to tariffs and real exchange rates. We find that export entry and especially export revenue are much more responsive to tariffs than to real exchange rates. This runs contrary to the prediction of workhorse models of international trade and international business cycles, which assume that changes in tariffs and changes in real exchange rates have similar effects on incentives to engage in international trade. We argue that the key feature missing from these models is forward-looking investment in customer base. Since tariffs are more predictable than real exchange rates, firms optimally increase investment in customer base more in response to favorable changes in tariffs than in response to similar favorable movements in real exchange rates, thus making export revenue more responsive to tariffs. Our findings suggest that policy makers should be very cautious about assuming an equivalence between trade and exchange rate policy.

To estimate how export entry, export exit and export revenue conditional on participation respond to tariffs and real exchange rates, we make use of fourteen years (1996-2009) of panel data on firms, products and exports for Ireland. We estimate entry, exit and revenue responses conditional on a rich set of fixed effects which control for marginal cost. We identify responses to both tariffs and real exchange rates using only time-series variation in these variables over our sample period. This allows us to eliminate the possibility that different responses are due to focusing on long-run responses in the case of tariffs, and short-run responses in the case of real exchange rates.

We find that although entry into export markets is four times more responsive to tariffs than it is to real exchange rates, the absolute level of entry responses to both variables is very modest. We do not find statistically significant responses of exit to either variable. On the other hand, the export revenue of incumbent exporters is six times more elastic with respect to tariffs than it is to real exchange rates. Taking account of the dynamics of entry, exit and export revenue, our estimates translate into an elasticity of aggregate exports with respect to tariffs of between -3.8 and -5.4, and with respect to real exchange rates of between 0.45 and 0.6. These aggregate elasticities are in line with estimates based on aggregate data summarized in Ruhl (2008). The first-order difference between these responses is due to the

¹See Staiger and Sykes (2010) and Gagnon (2012) for contrasting positions on this issue.

²See Freund and Gagnon (2017) for a discussion.

behavior of export revenue, not entry and exit.

To shed light on why the export revenue of incumbents responds so differently to tariffs and real exchange rates, we decompose revenue responses into the responses of quantities and the responses of markups (we hold marginal cost fixed throughout). We use the estimated responses to show that *conditional on markups*, quantities respond a lot to tariffs, but very little to real exchange rates. This implies that differential stickiness of prices in response to tariff and real exchange rate changes cannot explain why revenue responses are so different.

Instead, we hypothesize that the forward-looking nature of investment in customer base is likely to be key. This hypothesis builds on a growing literature in macroeconomics and international trade on the importance of customer base,³ as well as evidence in our data that export revenue is the outcome of forward-looking decisions. If customer base is something firms invest in, then decisions today affect not just revenue and profits today, but also revenue and profits in the future. As a result, firms are more likely to adjust investment in customer base (and hence export revenue) in response to variables like tariffs, which are predictable, than variables like real exchange rates, which are unpredictable.⁴

We investigate this hypothesis by simulating the responses of firms in a model with customer base to processes designed to resemble the tariff and real exchange rate processes faced by firms in our data. The main source of tariff variation in our data is tariff reductions agreed under the Uruguay Round in 1994, and phased in over the period 1995-2005. We mimic this in the model with a pre-announced path of tariff reductions. To mimic the real exchange rate, we use an AR(1) with parameters estimated using data on real exchange rates. We find that although the real exchange rate shock is very persistent, the elasticity of model export revenue with respect to the model tariff shock is nearly four times bigger than its elasticity with respect to the model real exchange rate shock.

While we are the first to provide an integrated empirical treatment of exporter responses to both tariffs and real exchange rates, micro-level 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 real exchange rates. Our findings are similar to his. They are also similar to those of Berman, Martin and Mayer (2012), who use French data to estimate participation, quantity, and price elasticities to exchange rates. In the trade literature, empirical work estimating micro responses to tariffs

³See e.g. Foster, Haltiwanger and Syverson (2008, 2016), Arkolakis (2010), Eaton, Kortum and Kramarz (2011), Drozd and Nosal (2012), Gourio and Rudanko (2014) and Fitzgerald, Haller and Yedid-Levi (2016).

⁴The intuition for this explanation is similar to that of Ruhl (2008), who focuses on sunk costs rather than customer base as the source of forward-looking behavior.

has focused on export participation. Lileeva and Trebler (2010) find a statistically significant impact of CAFTA tariff reductions on participation of Canadian firms in the US. Bustos (2011) estimates the response of export participation of Argentine firms to reductions in tariffs associated with Mercosur, and finds substantially bigger responses than we do.⁵ We are not aware of any prior work that uses micro data to estimate export revenue responses to tariffs.

As already noted, our work is also very closely related to Ruhl (2008), who coins the term “International Elasticity Puzzle” to describe the fact that workhorse models of international trade and business cycles cannot simultaneously match responses of aggregate exports to trade liberalizations and business cycle comovements of real exchange rates and relative quantities. A more recent contribution to this literature 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. Feenstra, Luck, Obstfeld and Russ (2014) also try to rationalize different elasticities. They estimate that for some sectors, the elasticity of substitution between home and foreign goods is different from the elasticity of substitution between goods from pairs of foreign countries. However, this cannot explain why, as we document, the elasticity of relative exports across pairs of destination markets at the firm level responds differently to movements in relative prices depending on the source of price variation. In a related paper, Imbs and Mejean (2015) argue that different elasticities are due to estimation at different levels of aggregation. However, we find evidence of different elasticities with respect to tariffs and real exchange rates at the (same) most disaggregated level possible.

In the next section of the paper, we frame our analysis using a model of the exporter’s problem. In the third section, we describe our data. In the fourth section, we document stylized facts about export dynamics in our data. In the fifth section, we describe our empirical strategy for estimating micro responses to tariffs and real exchange rates. In the sixth section, we describe our results, and combine them with the stylized facts about export dynamics to calculate elasticities of aggregate exports with respect to tariffs and real exchange rates. In the seventh section, we report our simulation results. The final section concludes.

⁵In both cases, it is cross-firm variation in tariffs that identifies participation responses. We only use variation over time within a given firm-product-market.

2 Model

To frame our analysis, we describe a state-of-the-art model of the exporter problem. We build on a substantial empirical literature, both reduced form and structural, that shows the importance of dynamic considerations for exporter behavior. The first generation of this literature emphasizes sunk costs as the source of dynamic decision making. More recent work downplays the quantitative importance of sunk costs, and highlights instead customer base accumulation and learning.⁶

The particular model we describe is based on Fitzgerald, Haller and Yedid-Levi (2016), who show that it can match a rich set of facts about exporter dynamics. The key elements of the model are as follows. Firms face the same cost of production irrespective of the market they are serving. They face sunk costs of export entry at the market level. In addition, they are uncertain about their idiosyncratic demand in each market, and must participate in a market in order to learn whether their demand is high or low. Once in a market, firms can attract customers in two ways: by charging low prices, and by investing in market-specific customer base through expenditures on marketing and advertising. These investments may be subject to adjustment costs. The combination of sunk costs of entry, learning about demand, and accumulation of customer base makes the firm’s problem forward-looking.

In what follows, i indexes firms, j indexes products, and k indexes markets. Firm i producing good j faces marginal cost of production $W_t C_t^{ij}$ expressed in some numeraire. This is the same for all markets the firm serves. W_t is an aggregate variable (the “wage”) which shifts all home firms’ marginal cost. C_t^{ij} follows an exogenous process known to the firm. C_t^{ij} may be autocorrelated, and marginal cost draws for different products produced by the same firm may be correlated with each other. The marginal cost of actually *delivering* a unit of the good to market k is given by $d^{jk} W_t C_t^{ij}$, where $d^{jk} \geq 1$ is a product-market-specific iceberg trade cost common to all firms. Firms observe W_t and C_t^{ij} at the time they take decisions, and they also know d^{jk} . We treat the set of firms in existence and the products each firm can produce as exogenous, and focus purely on decisions related to exporting.

In order to sell product j in market k , firm i must first pay a sunk cost of entry, $W_t S_t^{ijk}$, expressed in the same units as marginal cost. S_t^{ijk} is drawn i.i.d. from a fixed distribution, the same for all firms. Modeling sunk cost in this way is consistent with the fact that entry is rare, synchronization of entry across firms within a market or across markets within a firm is limited, and there is considerable overlap in the size distribution of exporters and non-

⁶See Roberts and Tybout (1997), Das, Roberts and Tybout (2007), Eaton et al. (2014), Ruhl and Willis (2017), and Fitzgerald, Haller and Yedid-Levi (2016).

exporters. There is also a per-period fixed cost of participating in a market, $W_t F_t^{ijk}$, expressed in the same units as marginal cost. F_t^{ijk} is also drawn i.i.d. from a fixed distribution, the same for all firms. It is this cost which generates export exit in the model.

A firm's demand in market k has four components. It depends on aggregate demand in market k , and on the consumer price of its good relative to the aggregate price level in market k . In addition, it depends on the fraction of customers it reaches, which is a function of its "customer base."⁷ Finally, there is an idiosyncratic component to demand. To learn about this component of demand, firms must actually sell in the market. Fitzgerald, Haller and Yedid-Levi (2016) show that when own-price elasticity of demand is constant, these assumptions are consistent with stylized facts about the post-entry behavior of export quantities and prices. Based on their formulation, we write demand for firm i 's product j in market k as follows:

$$Q_t^{ijk} = Q_t^k \left(\frac{(1 + T_t^{jk}) P_t^{ijk}}{P_t^k} \right)^{-\theta} \Phi \left(D_t^{ijk} \right) \exp \left(\varepsilon_t^{ijk} \right). \quad (1)$$

Here, Q_t^k is aggregate real demand in market k , P_t^k is the aggregate price level in market k , and P_t^{ijk} is the price the firm charges to customers from market k , net of tariffs. Both prices are expressed in the same units as marginal cost. T_t^{jk} is the ad valorem tariff on product j in market k .⁸ D_t^{ijk} is a state variable of the firm's problem which we refer to as customer base, and $\Phi \left(D_t^{ijk} \right)$ is the fraction of customers in market k reached by the firm. ε_t^{ijk} is idiosyncratic demand.

We assume that $\Phi(0) = 0$, $\Phi'(\cdot) > 0$, and $\Phi''(\cdot) < 0$, which guarantees that the optimal customer base conditional on export participation is finite and positive. Customer base accumulates as follows:⁹

$$D_t^{ijk} = d \left(D_{t-1}^{ijk}, X_t^{ijk} \right) + A_t^{ijk}, \quad (2)$$

where $X_t^{ijk} \in \{0, 1\}$ is an indicator for participation in market k by firm i selling product j at date t , and A_t^{ijk} is the increment to customer base due to marketing and advertising activities undertaken by the firm. We assume $d(\cdot, \cdot)$ is increasing in its first argument, but $d \left(D_{t-1}^{ijk}, X_t^{ijk} \right) \leq D_{t-1}^{ijk}$. The rate of depreciation may depend on whether or not the firm actually sells in the market in the current period. Expressed in the same units as marginal

⁷This follows Arkolakis (2010) and Eaton, Kortum and Kramarz (2011).

⁸The impact of specific tariffs on the incentives to export is quite different from that of ad valorem tariffs. We restrict attention to ad valorem tariffs in the empirical analysis.

⁹This is a dynamic extension of Arkolakis (2010).

cost, expenditure on investment in customer base is given by $W_t c(A_t^{ijk}, D_t^{ijk})$, with $c(\cdot, \cdot)$ increasing in its first argument. This formulation allows for the possibility of irreversibility and other costs of adjustment.

When making choices about participation, investment, and prices or quantities, firms observe the current values of individual state variables $\{C_t^{ij}, F_t^{ijk}, S_t^{ijk}\}$ and aggregate state variables $\{W_t, Q_t^k, P_t^k, T_t^{jk}\}$, as well as knowing the processes from which these are drawn. They do not observe current idiosyncratic demand ε_t^{ijk} at the time choices are made. They know the process for ε_t^{ijk} , and may have some additional information, \mathbf{I}_t^{ijk} , that they use to form conditional expectations of ε_t^{ijk} . This information may evolve over time and depend on past realizations of idiosyncratic demand.¹⁰

Firms set prices rather than quantities in the face of uncertainty about the idiosyncratic component of current demand.¹¹ Because demand is CES and there are no strategic interactions with other firms, the optimal price is equal to the statically optimal markup $(\frac{\theta}{\theta-1})$ over the marginal cost of delivery, independent of the firm's participation history, information set, or customer base.

Let $\tilde{\theta} = (\theta - 1)^{\theta-1} / \theta^\theta$ and let $\mathbf{Z}_t^{jk} = \{W_t, Q_t^k, P_t^k, T_t^{jk}\}$ be the vector of aggregate state variables, which follows a joint process known to all firms. Note that P_t^k can be interpreted as the real exchange rate between the home market and market k , and W_t can be interpreted as real wages in the home market. The intertemporal optimization problem for firm i , product j and market k (assuming discounting at rate β) is as follows:

$$\begin{aligned}
 & V\left(X_{t-1}^{ijk}, D_{t-1}^{ijk}, \mathbf{I}_t^{ijk}, F_t^{ijk}, S_t^{ijk}, C_t^{ij}, \mathbf{Z}_t^{jk}\right) = \\
 & \max_{\substack{X_t^{ijk} \in \{0, 1\} \\ A_t^{ijk}}} \left\{ \begin{aligned}
 & X_t^{ijk} \tilde{\theta} Q_t^k \left(\frac{(1+T_t^{jk})d^{jk}}{P_t^k}\right)^{-\theta} (W_t C_t^{ij})^{1-\theta} \Phi(D_t^{ijk}) \mathbb{E}\left(\exp(\varepsilon_t^{ijk}) \mid \mathbf{I}_t^{ijk}\right) \\
 & - X_t^{ijk} W_t \left(F_t^{ijk} + (1 - X_{t-1}^{ijk}) S_t^{ijk}\right) - W_t c\left(A_t^{ijk}, D_t^{ijk}\right) \\
 & + \beta \mathbb{E}\left(V\left(X_t^{ijk}, D_t^{ijk}, \mathbf{I}_{t+1}^{ijk}, F_{t+1}^{ijk}, S_{t+1}^{ijk}, C_{t+1}^{ij}, \mathbf{Z}_{t+1}^{jk}\right) \mid \mathbf{I}_t^{ijk}\right)
 \end{aligned} \right\} \quad (3)
 \end{aligned}$$

subject to (a) the accumulation of customer base D_t^{ijk} and (b) the evolution of information \mathbf{I}_t^{ijk} about idiosyncratic demand. The solution to this problem yields two policy functions

¹⁰Fitzgerald, Haller and Yedid-Levi (2016) examine the performance of specific assumptions about how this information evolves in matching the behavior of exports.

¹¹This assumption allows the model to match quantity and price as well as revenue dynamics.

for the firm:

$$X_t^{ijk} = X \left(X_{t-1}^{ijk}, D_{t-1}^{ijk}, \mathbf{I}_t^{ijk}, F_t^{ijk}, S_t^{ijk}, C_t^{ij}, \mathbf{Z}_t^{jk} \right) \quad (4)$$

$$A_t^{ijk} = A \left(X_{t-1}^{ijk}, D_{t-1}^{ijk}, \mathbf{I}_t^{ijk}, F_t^{ijk}, S_t^{ijk}, C_t^{ij}, \mathbf{Z}_t^{jk} \right) \quad (5)$$

Given these policy functions, export revenue is:

$$R_t^{ijk} = X_t^{ijk} \tilde{\theta} Q_t^k \left(\frac{(1 + T_t^{jk}) d^{jk}}{P_t^k} \right)^{-\theta} (W_t C_t^{ij})^{1-\theta} \Phi \left(d \left(D_{t-1}^{ijk}, X_t^{ijk} \right) + A_t^{ijk} \right) \exp \left(\varepsilon_t^{ijk} \right) \quad (6)$$

There are two main points to take away from the model. First, gross ad valorem tariffs $1 + T_t^{jk}$ and the real exchange rate P_t^k enter the firm's current payoff as a ratio, just as in workhorse models of international trade and international business cycles. In the workhorse models, firms' decisions are static, so optimal responses to identical changes in tariffs and the (inverse of the) real exchange rate are identical. But here, participation and revenue are the outcome of forward-looking decisions. This means that firms' responses to current realizations of tariffs and real exchange rates depend on their beliefs about future realizations. And the impact of current realizations on future beliefs depends on the stochastic processes followed by these variables. This implies that firms in this model may respond differently to changes in tariffs and real exchange rates if they follow processes which differ in their forecastability, persistence, or volatility. Moreover, it is not just participation which responds differently. Investment in customer base also responds differently, and this carries over to export revenue.

Second, though direct evidence on the sunk costs, learning, and investment in consumer base that are the source of forward-looking decision-making is hard to come by, the forward-looking nature of participation and investment decisions is reflected in history dependence of participation and revenue. Through D_{t-1}^{ijk} and \mathbf{I}_{t-1}^{ijk} , both participation and revenue may depend on the entire history of past participation. Identifying history dependence in the data is tricky, because persistent unobserved heterogeneity (in marginal cost C_t^{ij} and idiosyncratic demand ε_t^{ijk}), and selection on this unobserved heterogeneity, may generate patterns that look like history dependence. But if we can successfully control for the unobserved heterogeneity, history dependence of both participation and revenue is suggestive evidence of the forward-looking nature of firms' decisions.

3 Data

We now turn to the data, and describe in turn the three micro data sets we use, as well as the tariff and macro data we obtain from publicly available data sets. We make use of three sources of confidential micro data made available to us by the Central Statistics Office (CSO) in Ireland: the Census of Industrial Production (CIP), customs records, and the Prodcum survey.

3.1 Micro data

3.1.1 Census of Industrial Production

The CIP is an annual census of firms and establishments in manufacturing, mining and utilities. Firms with 3 or more persons engaged are required to file returns.¹² We make use of data for the years 1996 through 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 total revenue, employment, NACE sector, and country of ownership. We use these data to construct the set of potential export participants for our analysis of export entry, and also to condition on firm characteristics in our analysis of export entry, exit, and revenue.

In constructing the sample for our analysis, we drop firms with 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. Further details on the data and how we have cleaned them are provided in the Appendix.

3.1.2 Customs records

Our second source of data is customs records of Irish merchandise exports for the years 1996 through 2009. These records consist of the value (in Euros) and volume (in tonnes) of exports at the level of the VAT number, the Combined Nomenclature (CN) 8-digit product, and the export market (country). For most of our analysis we use only the value data. We make use of the volume data (for which coverage is more limited) only in robustness checks. These data are aggregated by the CSO to an annual frequency, and then matched by them to CIP firms using a correspondence between VAT numbers and CIP firm identifiers, along

¹²Multi-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 the match with the customs data can be performed.

with other confidential information. In the Appendix, we provide summary statistics on this match. We classify firms as exporters if they are matched with positive exports from the customs data. Match quality appears to be poor in 1996, so in our baseline empirical analysis we assume that export participation is censored in 1996.

A key feature of customs 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 in total shipments within the EU) is much higher than the reporting threshold for extra-European exports (254 Euro per transaction).¹³ However the intra-European threshold applies to *total* intra-European exports, not at the individual country level, and it is common for firms to report exports to individual EU countries which are well below the 635,000 Euro threshold. In our baseline analysis we pool both Intrastat and Extrastat exports, ignoring the fact that export participation may be censored for Intrastat markets. We check that our results are robust to restricting the sample in various ways to take account of potential censoring in Intrastat markets.

In our analysis of responses to tariffs and real exchange rates, we restrict attention to 30 of Ireland’s largest export markets, with the exact choice of markets governed by tariff availability. These markets account for 94% of total exports in the sample period. Table 1 reports the breakdown by market. The main export markets for our firms are the US, the UK and the large economies of the Eurozone.

Table 1: Percentage of exports by destination

	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
Australia	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Canada	0	0	0	0	0	1	0	1	0	0	0	0	0	1
China	0	0	0	0	0	0	1	1	1	1	1	2	2	2
Denmark	1	1	1	1	1	1	1	1	1	1	1	1	1	0
Japan	2	3	3	3	4	4	3	3	3	3	3	2	2	2
Norway	1	1	1	1	1	1	1	1	1	1	1	1	0	0
Sweden	2	2	2	2	2	2	1	1	1	1	1	1	1	1
Switzerland	2	2	2	2	3	3	4	4	4	5	4	4	4	4
UK	24	24	21	18	18	17	16	15	16	16	17	17	17	15
US	11	12	16	16	20	19	21	23	21	16	16	17	18	22
Euro 9	43	41	43	47	42	41	43	42	44	47	45	43	43	44
Other	5	4	3	4	4	5	4	4	4	3	4	4	5	4
Total in-sample	93	92	93	95	93	94	95	95	94	94	93	93	94	95

Notes: We use only customs export records which match to our cleaned data set of CIP firms. The Euro 9 includes Austria, Belgium, Finland, France, Germany, Italy, Netherlands, Portugal and Spain. Other includes Brazil, Hong Kong, India, Malaysia, Mexico, New Zealand, Saudi Arabia, South Africa, Thailand, Turkey and the United Arab Emirates. Source: CSO and authors’ calculations.

¹³Intra-European exports below the threshold are recovered based on VAT returns. The destination within the EU is not recorded for these returns. Since we do not know the exact destination, we do not make use of these exports other than in our summary statistics on exporters vs. non-exporters.

A final important feature of the customs data is that the 8-digit CN classification system changes every year. We concord the product-level export data over time at the most disaggregated level possible following the approach of Pierce and Schott (2012) as implemented by van Beveren, Bernard and Vandebussche (2012). We make use of these concorded data in our analysis of export revenue. For our analysis of export participation we further concord the trade data with production data, as described below.

3.1.3 Prodcom survey

Our third source of data is the Prodcom survey for the years 1996 to 2009. This is an annual survey of the value (in Euros) and volume (in tonnes) of all products manufactured by the firm and sold in the relevant year. We use only the value data. The survey basis is all firms in the CIP, excluding some mining sectors. Not all CIP firms respond to the survey (the share of firms participating is 64% on average over the sample), but it covers 94% of total CIP revenue. Products are classified at the 8-digit level according to the Prodcom classification. These data allow us to construct the set of products a firm might potentially export.

As with the CN 8-digit classification, the 8-digit Prodcom classification system changes from year to year. We first concord the Prodcom data over time at the most disaggregated level possible. We then concord the time-concorded Prodcom data to the time-concorded CN data. In doing so, we follow van Beveren, Bernard and Vandebussche (2012). There are fewer Prodcom 8-digit products than there are CN 8-digit products, so this involves some aggregation at the CN8 level.

To obtain the set of (concorded) products a firm might potentially export in any given year, we collapse across years to find the set of products that it produces at least once during its lifetime in the CIP. If the firm does not export one of these products to a particular destination market in a particular year, we code this as non-participation at the firm-product-market-year level.

3.2 Tariff data

Our analysis requires data on the tariffs faced by Irish firms in the 30 destination markets we focus on. As Ireland is a member of the EU, Irish exporters do not face tariffs in EU markets. In countries for which there is no preferential trade agreement with the EU, Irish exporters face the Most Favored Nation (MFN) tariff. We obtain MFN tariff data from the WTO for the following countries: Australia, Brazil, Canada, China, Hong Kong, India, Japan, Malaysia, Mexico, New Zealand, Saudi Arabia, South Africa, Thailand, United

Arab Emirates and the US. The EU has preferential trade agreements with Switzerland, Norway, Turkey, South Africa and Mexico for some or all of our sample period. We obtain information on the relevant preferential tariffs from the WTO and other sources (detailed in the Appendix). For several of the countries in which Irish exporters face tariffs, especially developing countries, tariff data are not available for all of the sample years. Under limited circumstances we interpolate tariffs for the years for which data are not available. Full details of the sources, construction and coverage of the tariff data are in the Appendix.

Since the EU is a customs union, in all of the countries where Irish exporters do face tariffs, the actual level of tariffs is determined by bargaining between the EU and the relevant counterparties. Given that Ireland is one small member of the EU, this makes us fairly confident that reverse causality is not a problem in estimating export participation and revenue responses to destination market tariffs.

The tariff data are reported using the Harmonized System (HS) 6-digit classification. This is used by all countries as the basis for their tariff lines, and is the most disaggregated level at which tariffs, export, and production data can be matched across countries. We restrict attention to HS6 product-market-years for which there are no non-ad-valorem tariffs,¹⁴ and for which there is no sub-HS6 variation in ad valorem tariffs.¹⁵ The HS6 classification changes in 2002 and 2007. We concord the classification over the period 1996-2009 following the approach of Pierce and Schott (2012). When this implies joining multiple HS6 categories together in a given period, we take the simple average of tariffs to construct tariffs at the concorded product-market-year level.

Finally, to make use of the tariff data, we concord it with our export and production data. At a 6-digit level, the CN (export) and Prodcom (production) classifications correspond to the HS classification. In some cases, our concordance of the CN and Prodcom classifications results in “products” that cover multiple HS6 categories. In our analysis of export entry and exit, we take the simple average of tariffs across the relevant HS6 categories to obtain a tariff at the concorded product-market-year level in these cases. In our analysis of export revenue, we use revenue at the firm-HS6-market-year level to construct a weighted average of tariffs across the relevant HS6 categories.

To identify responses to tariffs, in our baseline analysis we rely on time-series variation only.¹⁶ The main source of time-series variation in tariffs faced by Irish exporters in this

¹⁴Unlike ad-valorem tariffs, non-ad-valorem tariffs affect incentives to export differently depending on the firm’s export price.

¹⁵We test robustness to relaxing this requirement.

¹⁶There is considerable cross-sectional variation in tariff levels, but we restrict ourselves to using similar sources of variation to identify responses to tariffs and real exchange rates.

period 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 1994. We illustrate the nature of this variation by performing country-by-country regressions of the log gross ad valorem tariff, i.e. $\ln(1 + T_t^{jk})$, on HS6 fixed effects and year dummies for some important markets, and plotting the coefficients on the year dummies in Figure 1. Additional summary statistics on the tariff data are reported in the Appendix.




Figure 1: Average evolution of tariffs for five rich destination markets

Notes: Figure shows coefficients on year dummies in country-by-country regression of $\ln(1 + T_t^{jk})$ on HS6 fixed effects and year dummies. Source: WTO and authors' calculations.

3.3 Macro data

Our analysis requires measures of the real consumption exchange rate and real local currency demand in the 30 destination markets we focus on. We construct real exchange rates between Ireland and these markets using data on annual average nominal exchange rates and CPIs from the IMF's *International Financial Statistics* (IFS). We proxy real demand using GDP less exports plus imports, all measured in current local currency, with this aggregate deflated by the relevant country's CPI. The National Accounts data are taken from the OECD's *National Accounts Statistics* where available, and otherwise from the World Bank's *World Development Indicators*. The CPIs are taken from IFS.

The bulk of the variation in real exchange rates is driven by nominal exchange rates. This is illustrated in Figure 2, which contrasts the evolution of real exchange rates with selected

non-Euro markets with the evolution of real exchange rates with the nine Eurozone markets we work with.

For most of the sample period (i.e. 1998-2009), monetary policy is decided at the Eurozone level by the ECB, mitigating concerns about reverse causality. In a robustness check, we exclude markets which use the Euro so as to further mitigate concerns about reverse causality.




Figure 2: Evolution of real exchange rates in selected non-Euro and Eurozone markets

Notes: Figure shows log deviation of annual average real exchange rate from 1996 level for selected Non-Euro and Eurozone countries. The Non-Euro countries are Australia, Brazil, China, Japan, Norway, Sweden, Switzerland, UK and USA. The Eurozone countries are Austria, Belgium, Finland, France, Germany, Italy, Netherlands, Spain and Portugal. Source: IMF and authors' calculations.

4 Summary statistics on exporter dynamics

In this section, we show how in our data, export behavior depends on firm characteristics and past export history.¹⁷ There are three reasons to do this. First, we lay the foundation for the empirical strategy we use to estimate responses of export participation and export revenue to tariffs and real exchange rates. Second, our estimates of the post-entry dynamics of exports will allow us to translate these estimates into aggregate responses. Third, by providing evidence that is consistent with the predictions of the model about history dependence (albeit without structurally estimating the model), we lay the foundation for our explanation for why firms respond differently to tariffs and real exchange rates.

¹⁷For export exit and revenue, our approach closely follows Fitzgerald, Haller and Yedid-Levi (2016).

Table 2: Summary statistics: Firms and exports, averages 1996-2009

Number of firms per year	4729
Employees per firm	51
Firm age (years)	17
Share of firms foreign owned	0.13
Share of multi-plant firms	0.03
Export participation rate	0.42
Export entry rate	0.09
Export exit rate	0.12
Exporter size premium (employees)	3.29
Exporter size premium (revenue)	5.20
Export share conditional on exporting	0.33
Number of markets per exporter	6.9

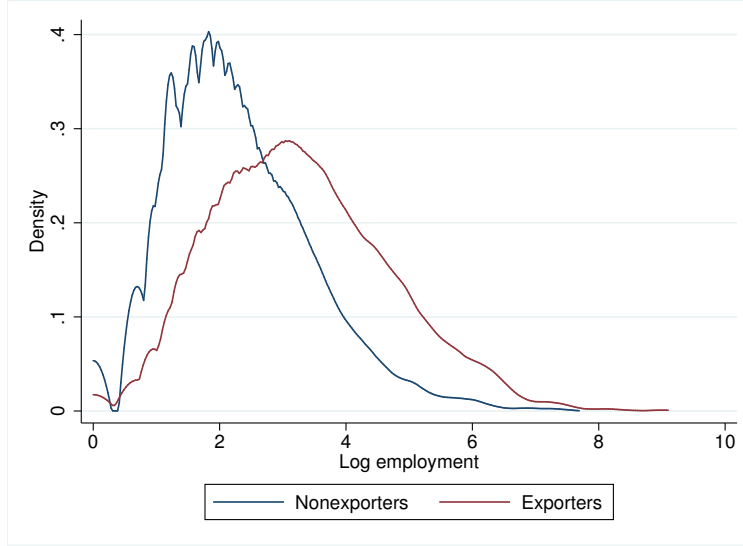
Notes: Statistics are for our cleaned data set of CIP firms. Firms are defined as exporters if they are matched to positive concorded product exports from customs data. Export intensity is calculated as total concorded product exports from customs divided by sales reported in the CIP. Values greater than 1 are replaced by 1. The set of potential entrants used to calculate the entry rate in year t includes firms born in year t . The set of potential exiters used to construct the exit rate in year t includes firms present in the CIP in year $t - 1$ but not in year t . Entry and exit rate averages are calculated over 1998-2009. Source: CSO and authors' calculations.

We start by confirming that Irish exporters look like exporters in other countries. Table 2 shows that exporters are on average bigger than firms which sell exclusively to the domestic market, though there is substantial overlap in the size distributions, as illustrated in Figure 3. Export participation and export intensity for exporters are both relatively high, as is typical of small open economies in Europe (see ISGEP (2008)). Export participation is quite persistent, a fact which is again typical of other countries for which micro data on exporting are available.

Next, we turn to export behavior at the firm-market and firm-product-market level. From this point forward, we restrict attention to the 30 export markets for which we collect tariff and macro data. Table 3 reports average participation, entry, and exit rates across all 30 markets. It also reports average participation, entry, and exit rates by firm size, along with rates for a few selected markets. Participation and entry rates are increasing in firm size. Participation and entry rates at the level of individual markets are low. Persistence of exporting overall conceals considerable churn at the level of individual markets: exit rates in individual markets above 30% are not unusual.

We now document history dependence in export participation and revenue, using the firm-product-market-year as the unit of analysis. Because the average probability of export participation differs dramatically based on lagged participation, we examine entry and exit separately. We take a reduced form approach, estimating linear probability models for

Figure 3: Distribution of log employment for non-exporters and exporters



Notes: Firms are defined as exporters if they are matched to positive exports from customs data. Figure plots kernel density estimates of the distribution of log employment for non-exporters and exporters, pooling across 1996-2009. Source: CSO and authors' calculations.

Table 3: Average participation, entry and exit rates, 1998-2009

Rate	All firms	Small [†]	Med [†]	Large [†]	USA	Japan	China	UK*	Germany*	France*
Firm-market level										
Particip.	0.051	0.03	0.18	0.29	0.11	0.05	0.03	0.17	0.10	0.10
Entry	0.012	0.01	0.04	0.07	0.03	0.01	0.01	0.03	0.02	0.02
Exit	0.226	0.26	0.21	0.17	0.25	0.26	0.33	0.12	0.17	0.16
Firm-product-market level										
Particip.	0.028	0.01	0.07	0.12	0.05	0.02	0.02	0.12	0.06	0.06
Entry	0.008	0.00	0.02	0.03	0.02	0.01	0.01	0.02	0.01	0.01
Exit	0.261	0.29	0.25	0.22	0.30	0.34	0.41	0.16	0.20	0.20

Notes: Products are defined based on the concordance of Prodcom and CN product definitions as described in Sections 3.1.2 and 3.1.3. The set of potential entrants used to calculate the entry rate in year t includes firms and firm-products born in year t . The set of potential exiters used to construct the exit rate in year t includes firms present in the CIP in year $t - 1$ but not in year t . [†] Small firms are those with fewer than 100 employees. Medium firms have 100-249 employees. Large firms have 250+ employees. * Participation in EU markets is censored by the higher reporting threshold. Source: CSO and authors' calculations.

Table 4: Entry and past participation

	coeff.	s.e.
$evpart_t^{ijk}$	0.091	(0.002)**
firm-prod-yr f.e.		yes
prod-mkt-yr f.e.		yes
N	2,785,844	
R ²	0.33	
R ² -adj	0.21	

Notes: Products are defined based on the concordance of Prodcom and CN product definitions as described in Sections 3.1.2 and 3.1.3. Only potential entrants are included in the regression. The set of potential entrants includes firms born at date t . Dependent variable is an indicator for entry at the firm-product-market level. Sample period is 1998-2009. Omitted category is observations with no past export experience. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

entry and exit,¹⁸ controlling for marginal cost using firm-product-year fixed effects, and using indicator variables to capture past export history, rather than imposing a particular functional form. In the case of export revenue, we are careful to separate history dependence from selection on idiosyncratic demand. Since our focus at this point is history dependence, not responses to tariffs and real exchange rates, we control for tariffs and real exchange rates (and other shocks) using product-market-year fixed effects.

We start with export entry. We define a firm-product-market observation as a potential entrant at date t if the firm is present in the CIP at date t , and did not export that product to that market in $t - 1$.¹⁹ As in Section 2, X_t^{ijk} is an indicator variable for participation by firm i in market k with product j at time t . For each potential entrant, we also construct an indicator variable for participation prior to $t - 1$. Our estimating equation is then:

$$\Pr \left[X_t^{ijk} = 1 | X_{t-1}^{ijk} = 0 \right] = c_t^{ij} + \gamma_t^{jk} + \lambda' evpart_t^{ijk} + \varepsilon_t^{ijk}, \quad (7)$$

where c_t^{ij} is a firm-product-year fixed effect, γ_t^{jk} is a product-market-year fixed effect, and $evpart_t^{ijk}$ is the indicator variable for past participation. The results are reported in Table 4. Even conditional on the fixed effects, a history of participation is a strong predictor of entry, increasing the probability of entry by 9%.²⁰

We turn next to exit. We define a firm-product-market observation as a potential exiter at date t if it had positive exports in year $t - 1$.²¹ For each observation currently participating

¹⁸This is a convenient way of characterizing average marginal effects conditional on a rich set of fixed effects.

¹⁹Entering firms are categorized as potential entrants in all markets.

²⁰In results reported in the Appendix, we show that a more recent history of participation is a stronger predictor of entry than a more distant history of participation.

²¹Exiting firms are included.

Table 5: Exit and past participation

Market tenure	coeff.	s.e.
2 years	-0.147	(0.010)**
3 years	-0.206	(0.012)**
4 years	-0.241	(0.013)**
5 years	-0.218	(0.014)**
6 years	-0.263	(0.015)**
7+ years	-0.265	(0.013)**
censored	-0.272	(0.009)**
firm-prod-yr f.e.	yes	
prod-mkt-yr f.e.	yes	
N	47,592	
R ²	0.65	
R ² -adj	0.47	

Notes: Products are defined based on the concordance of Prodcom and CN product definitions as described in Sections 3.1.2 and 3.1.3. Only potential exiters are included in the regression. This includes exiting firms. Dependent variable is an indicator for exit at the firm-product-market level in the next year. Omitted category is observations with market tenure equal to 1 year. Censored indicates that market tenure is censored because exporting was ongoing at the start of the sample. This is extended also to observations observed to “enter” in 1997, because of the poor quality of the match between CIP and customs data in 1996. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors’ calculations.

for which entry is observed in the sample, we calculate *market tenure*, where market tenure is defined as the number of years of continuous exporting for the firm-product-market triplet in question. We top-code market tenure at 7 years.²² We also create an indicator variable for market tenure being censored because exporting was already ongoing at the beginning of the sample. Our estimating equation is:

$$\Pr \left[X_t^{ijk} = 0 | X_{t-1}^{ijk} = 1 \right] = c_t^{ij} + \gamma_t^{jk} + \boldsymbol{\lambda}' \mathbf{a}_{t-1}^{ijk} + cens_{t-1}^{ijk} + \varepsilon_t^{ijk}, \quad (8)$$

where c_t^{ij} is a firm-product-year fixed effect, γ_t^{jk} is a product-market-year fixed effect, and \mathbf{a}_{t-1}^{ijk} is a vector of indicator variables for lagged market tenure, i.e. firm i ’s tenure with product j in market k as of period $t - 1$. The excluded category is observations with lagged market tenure of 1. $cens_{t-1}^{ijk}$ is the indicator for lagged market tenure being censored. The results are reported in Table 5. Exit rates are broadly declining with tenure in a market. Exit rates for observations where tenure is censored are comparable to those for observations with market tenure of 6 years or more.

Finally, we examine export revenue conditional on participation. To separate history dependence from selection on idiosyncratic demand, we define an additional concept: *export*

²²We top-code because the behavior of the exit hazard at long horizons cannot be distinguished from cohort effects. The Appendix to Fitzgerald, Haller and Yedid-Levi (2016) shows that results are robust to varying the level at which market tenure is top-coded.

spell length. An *export spell* is a continuous episode of exporting by the same firm-product-market triplet. *Export spell length* is the number of years from entry to exit in such an episode of continuous export participation. For each observation currently participating, we calculate the length of the relevant export spell, top-coding at 7 years.²³ We then allow the evolution of revenue with market tenure to differ across export spells of different length. Let r_t^{ijk} be the log of export revenue of firm i selling product j to market k at time t . Our estimating equation is:

$$r_t^{ijk} = c_t^{ij} + \gamma_t^{jk} + \boldsymbol{\delta}' \left(\mathbf{a}_t^{ijk} \otimes \mathbf{s}_t^{ijk} \right) + \mathbf{cens}_t^{ijk} + \varepsilon_t^{ijk}, \quad (9)$$

where c_t^{ij} is a firm-product-year fixed effect and γ_t^{jk} is a product-market-year fixed effect. \mathbf{a}_t^{ijk} is a vector of indicator variables for firm i 's tenure with product j in market k at time t , while \mathbf{s}_t^{ijk} is a vector of indicator variables for the length of the relevant export spell. The symbol \otimes indicates the Kronecker product, but redundant interactions between \mathbf{a}_t^{ijk} and \mathbf{s}_t^{ijk} are dropped. Meanwhile, \mathbf{cens}_t^{ijk} contains indicators for censored entry and censored exit.


Full results from estimating equation (9) are reported in the Appendix. Figure 4 graphs the export trajectories obtained by taking the exponents of our estimates of $\boldsymbol{\delta}$. Broadly speaking, export spells that start bigger last longer, consistent with selection on idiosyncratic demand. Because we condition on export spell length, the growth patterns in this figure are not due to selection, but provide evidence of history dependence. There is substantial growth of export revenue in the first five years of the longest-lasting export spells, while short-lived export spells grow more slowly, and exhibit hump-shaped dynamics.

The three take-aways from Tables 4 and 5, and Figure 4 are as follows. First, there are robust patterns in entry, exit, and export revenue that are not due to changes in tariffs and real exchange rates. To obtain clean estimates of responses to tariffs and real exchange rates, we must control for these patterns. Second, the precise patterns we document imply that in the short run, responses of export entry and exit to tariffs and exchange rates will have only a modest impact on aggregate exports, because as Figure 4 shows, entrants and exiters are small compared to long-time participants. In the longer term, entry responses can have a bigger effect, as a few successful entrants grow rapidly in new markets. However, the “longer term” is quite long. Observations with market tenure of less than 6 years account for just over 50% of observations, but only 25% of total exports in the sample.²⁴ Third,

²³Spell length may be censored in two ways: because the spell is ongoing at the beginning of the sample (censored entry), or because it is still ongoing at the end of the sample (censored exit). In spells where exit is censored, but market tenure is at least 7 by the end of the sample, we assign the topcoded spell length.

²⁴See the Appendix for details.

Figure 4: Post-entry dynamics of revenue



Notes: Figure shows evolution of revenue at the firm-product-market level with market tenure, allowing trajectories to differ by export spell length. Ongoing spells are spells for which entry is censored by the beginning of the sample. Trajectories are conditional on firm-product-year and product-market-year effects. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

these patterns confirm the prediction of the model that export history matters, and matters not just for export participation, but also for export revenue. This is evidence that firms are forward-looking, which is key to our explanation of responses to tariffs and real exchange rates.

5 Empirical strategy

In this section we describe our empirical strategy for estimating entry, exit and revenue responses to tariffs and real exchange rates.

5.1 Entry and exit

We start with the linear probability models for entry and exit equations from the previous section. But instead of using product-market-year fixed effects to absorb the effects of shocks, we include the log of tariffs, real exchange rates and foreign demand, along with product-market fixed effects. This means that the coefficients on the variables of interest can be interpreted as the marginal effects of percentage deviations from average levels over the sample for the relevant product-market. Since tariffs vary at the product-market-year level, and real exchange rates and foreign demand at the market-year level, only time-series variation is used to identify the coefficients on all three variables.

The baseline specification for entry is as follows:

$$\Pr \left[X_t^{ijk} = 1 | X_{t-1}^{ijk} = 0 \right] = c_t^{ij} + \gamma^{jk} + \lambda evpart_{t-1}^{ijk} + \beta' \mathbf{z}_t^{ijk} + \eta_t^{ijk} \quad (10)$$

As in the previous section, $X_t^{ijk} = 1$ is an indicator for participation, c_t^{ij} is a firm-product-year fixed effect, and γ^{jk} is a product-market fixed effect. The variable $evpart_{t-1}^{ijk}$ is, as in the previous section, an indicator variable equal to one if firm i participated in market k with product j in some year prior to $t - 1$. Finally, \mathbf{z}_t^{ijk} is a vector, the elements of which are $\tau_t^{jk} = \ln(1 + T_t^{jk})$, which is the log gross ad valorem tariff faced by exporter i in market k , p_t^k , which is the log of an index of the real exchange rate between the home market and market k , and q_t^k , which is the log of an index of real aggregate demand in market k . The construction of tariffs and macro variables is described above in sections 3.3 and 3.2. We use robust standard errors to deal with the standard heteroskedasticity issue in estimating linear probability models.

The baseline specification for exit is as follows:

$$\Pr \left[X_t^{ijk} = 0 | X_{t-1}^{ijk} = 1 \right] = c_t^{ij} + \gamma^{jk} + \lambda' \mathbf{a}_{t-1}^{ijk} + \beta' \mathbf{z}_t^{ijk} + \nu_t^{ijk} \quad (11)$$

Firm-product-year and product-market fixed effects are exactly as in the entry equation. As in the previous section, \mathbf{a}_{t-1}^{ijk} is a vector of indicator variables for firm i 's topcoded tenure in market k with product j at time $t - 1$. Finally, \mathbf{z}_t^{ijk} is as in the entry equation.

The model suggests that the sensitivity of entry and exit rates to the variables of interest may depend on the density of firms at the margin. As shown in Tables 3, 4 and 5, entry and exit rates do differ systematically across firms of different sizes, across different markets, and across participation histories. By averaging marginal effects across all observations, we may run the risk of mistakenly inferring insensitivity to shocks. In a set of robustness checks, we allow for variation in marginal effects, either by splitting the sample, or by allowing for interactions between \mathbf{z}_t^{ijk} and an appropriate variable.²⁵

5.2 Export revenue

For export revenue we start with the log-linear specification from the previous section. As with entry and exit, instead of using product-market-year fixed effects to absorb the effects of shocks, we include the log of tariffs, real exchange rates and foreign demand, along with

²⁵Whether to split the sample or allow for interactions depends on whether the variable of interest varies across firm-product-years (split) or within firm-product-years (interact).

product-market fixed effects. This means that the coefficients of interest can be interpreted as the elasticities of revenue with respect to these variables. Similar to our analysis of entry and exit, since tariffs vary at the product-market-year level, and real exchange rates and foreign demand at the market-year level, only time-series variation is used to identify the coefficients on all three variables.

A key issue we must address is the possibility of selection bias. As noted in the previous section, higher initial export revenue forecasts longer survival in a market, which is direct evidence of selection on unobserved idiosyncratic demand. This type of selection may lead to biased estimates of tariff and real exchange rate elasticities, because the expectation of the error term conditional on participation may be correlated with the realizations of these variables.

Our environment is not conducive to a Heckman-style selection correction, because we do not have instruments which affect participation but not revenue conditional on participation.²⁶ However, we do know, based on the observed history of participation, that some observations are more likely to be marginal than others. In particular, as Table 5 shows, exit rates decline systematically with market tenure. Observations with tenure of 6 or more years, or which were ongoing at the start of the sample, have the lowest exit rates, on the order of 14-16% unconditionally.²⁷ Correlation of unobserved idiosyncratic demand with tariffs and real exchange rates is less likely to be an issue for observations that are infra-marginal than for observations that are marginal, and the bias for these observations should be less severe.²⁸ Based on this, we allow elasticities with respect to tariffs and macro variables differ across observations according to their market tenure, and we interpret the coefficients on the low exit probability observations as being less subject to selection bias.

Our baseline estimating equation for export revenue conditional on participation is then:

$$r_t^{ijk} = c_t^{ij} + \gamma^{jk} + \boldsymbol{\delta}' \left(\mathbf{a}_t^{ijk} \otimes \mathbf{s}_t^{ijk} \right) + \mathbf{cens}_t^{ijk} + \boldsymbol{\beta}' \mathbf{z}_t^{ijk} + \boldsymbol{\phi}' \left(\mathbf{z}_t^{ik} * low_t^{ijk} \right) + \varepsilon_t^{ijk} \quad (12)$$

Here, r_t^{ijk} is the log of export revenue of firm i from selling product j to market k at date t , c_t^{ij} is a firm-product-year fixed effect, and γ^{jk} is a product-market fixed effect. As in the previous section, \mathbf{a}_t^{ijk} is a vector of indicator variables for firm i 's topcoded tenure in market k with product j at time t , while \mathbf{s}_t^{ijk} is a vector of indicator variables for the topcoded length of the relevant export spell. Redundant interactions between these variables are dropped.

²⁶Somewhat less crucially, other dimensions of our identification strategy rely on a rich set of fixed effects which would be tricky to incorporate in the standard Heckman setup.

²⁷See the Appendix for details.

²⁸Loosely speaking, this idea resembles Chamberlain's (1986) "identification at infinity."

\mathbf{cens}_t^{ijk} is a vector of indicators for censored entry or exit. \mathbf{z}_t^{ik} is as described above. Finally, low_t^{ijk} (for low exit probability) is an indicator variable, set equal to 1 for observations where market tenure is at least 6 years as well as observations in export spells whose entry is censored, and equal to zero otherwise.²⁹

In robustness analysis, we split export revenue into its quantity and price (unit value) components, and estimate equation (12) with log quantity and log price in turn as the dependent variable.

6 Results

We now describe our main results.

6.1 Export entry and exit

The first column of Table 6 reports the results from estimating our baseline entry equation, equation (10). We find that the probability of entry is higher in periods where tariffs in the destination market are low, and in periods where the Irish real exchange rate is weak vis-a-vis the destination market relative to its sample average (i.e. higher p_t^k). The signs of these effects are as predicted by theory, and both coefficients are statistically significant at the 5% level.

Remembering that the average in-sample entry rate is 0.8%, the estimated coefficients can be interpreted as follows: A reduction in tariffs from 10% to 0 would imply an increase in the entry rate from 0.8%, to 0.86%, while a 10% depreciation of the home real exchange rate against the destination market would imply an increase in the entry rate from 0.8% to 0.81%. The absolute value of the coefficient on the tariff is 4.7 times that on the real exchange rate, so entry is more responsive to tariffs than to real exchange rates. But both responses are economically modest, in a sense we will make precise at the end of this section.

The second column of Table 6 reports the results from estimating our baseline exit equation, equation (11). We do not find evidence of statistically significant relationships between the rate of exit and the level of tariffs or real exchange rates. This contrasts with our findings on entry.

Although the point estimates in the exit equation are not statistically significant, it is still useful to explain how to interpret the coefficient signs and magnitudes. A positive coefficient

²⁹The level effect of low_t^{ijk} is captured by the inclusion of \mathbf{a}_t^{ijk} and \mathbf{cens}_t^{ijk} in the regression.

Table 6: Export entry and exit responses to shocks

	(1)		(2)	
	Entry		Exit	
	coeff	s.e.	coeff	s.e.
τ_t^{jk}	-0.006	(0.002)**	0.176	(0.115)
p_t^k	0.001	(0.001)**	0.012	(0.023)
q_t^k	0.005	(0.001)**	0.009	(0.024)
Export history controls	yes		yes	
Firm-prod-yr f.e.	yes		yes	
Prod-mkt f.e.	yes		yes	
N	2,380,829		70,198	
R ²	0.28		0.60	
R ² -adjusted	0.23		0.49	
	In-sample entry rate		In-sample exit rate	
	0.008		0.218	

Notes: Products are defined based on the concordance of Prodcom and CN product definitions as described in Sections 3.1.2 and 3.1.3. Sample period is 1998-2009. Only potential entrants are included in the entry regression. Only potential exiters are included in the exit regression. Dependent variable is an indicator for entry or exit at the firm-product-market-year level. Export history controls in column (1) include $evpar_{t-1}^{ijk}$. Export history controls in column (2) include \mathbf{a}_{t-1}^{ijk} and $\mathbf{cens}_{t-1}^{ijk}$. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

on tariffs implies that exit rates are lower when tariffs are lower, consistent with what theory predicts. The point estimate implies that if the exit rate is 21.8% (the average exit rate in-sample) under a tariff of 10%, it will be 20.04% under a zero tariff. A positive coefficient on the real exchange rate implies that depreciations of the Irish real exchange rate against a destination market increase exit from that market, contrary to what theory predicts. The point estimate implies that a 10% depreciation of the home real exchange rate against a destination market would imply an increase in the exit rate from 21.8% to 21.92%.

We perform a large number of robustness tests, varying the specification, the estimation sample, and measurement of the independent variables. Our findings are broadly robust to these variations. Full results are reported in the Appendix. Here we highlight a few points of interest.

First, in order to investigate whether averaging of marginal effects across infra-marginal and marginal observations leads us to conclude incorrectly that entry and exit responses are small, or not present, we allow marginal effects to differ across export histories and firm size. We find that export entry is indeed more sensitive to both tariffs and real exchange rates in observations with a history of past participation than in observations with no such history. But the absolute size of the response is in proportion to the higher unconditional entry rate for these observations, and thus still modest. Entry sensitivities vary with firm size, but remain modest. Our baseline responses appear to be driven by entry sensitivity to tariffs for large firms and to real exchange rates for small firms.

As regards exit, we find weak evidence that pooling across observations may mask exit

sensitivity to tariffs. Lower tariffs do reduce the probability of export exit for observations with longer market tenure, with the coefficient statistically significant at the 10% level, and similar in magnitude to the baseline estimate. But conditioning on firm size does not affect the baseline finding that there are no statistically significant responses of exit to either variable.

Second, in our baseline specification, we identify responses to tariffs and real exchange rates based on deviations from product-market averages. This approach is helpful in estimating responses to tariffs, where there is relatively little year-on-year variation. But estimation in first differences has some advantages. It allows us to investigate whether firms respond to future tariff changes, which could be an indicator of forward-looking behavior, and to look for asymmetries in response to positive and negative shocks.

When we estimate in first differences, entry responses to both tariffs and real exchange rates are magnified, consistent with short-run overshooting of entry relative to the long run. We find that future tariff reductions reduce entry today, while entry responses to real exchange rates appear to be driven by adverse movements rather than positive movements. As regards exit, in first differences, real exchange rate depreciations *increase* the rate of exit, rather than decreasing it as theory would suggest. But this finding is not robust to allowing for asymmetric responses to positive and negative shocks. Moreover, we do not find any evidence that exit responds to future tariff changes.

Third, there are various reasons to test whether our results are robust to changing the sample of countries included in the estimation. Restricting the sample to non-Eurozone markets does not change our baseline results on entry and exit, though when we restrict further to non-EU markets, we do not find a statistically significant response of entry to real exchange rates.

Finally, it is of interest to know whether responses to the price and nominal exchange rate components of the real exchange rate are different. We find that entry is more responsive to the price component than to the nominal exchange rate component, while exit is not responsive to either.

6.2 Revenue

The first column of Table 7 reports the results from estimating our baseline export revenue equation, equation (12). We focus on the elasticities for observations with low predicted exit probability for which selection bias is less likely to be an issue. For these observations, we estimate that the elasticity of revenue with respect to tariffs is equal to -3.02, and the

Table 7: Export revenue, quantity and price responses to shocks

		(1)		(2)		(3)	
		Revenue		Quantity		Price	
		coeff	s.e.	coeff	s.e.	coeff	s.e.
Low exit prob	τ_t^{jk}	-3.02	(0.65)**	-2.90	(0.66)**	-0.03	(0.35)
	p_t^k	0.46	(0.08)**	0.24	(0.09)**	0.18	(0.04)**
	q_t^k	0.24	(0.09)**	0.23	(0.09)**	0.08	(0.05)*
High exit prob	τ_t^{jk}	0.89	(0.58)	0.83	(0.54)	0.18	(0.31)
	p_t^k	0.42	(0.08)**	0.21	(0.09)**	0.18	(0.04)**
	q_t^k	0.17	(0.09)*	0.17	(0.09)*	0.08	(0.05)*
Export history controls		yes		yes		yes	
Firm-prod-yr f.e.		yes		yes		yes	
Prod-mkt f.e.		yes		yes		yes	
N		191,780		184,890		184,890	
R ²		0.77		0.83		0.91	
R ² -adjusted		0.68		0.77		0.87	

Notes: Products are defined based on the concordance of CN product definitions as described in Section 3.1.2. Dependent variable is log revenue in Euro, log quantity in tonnes, or log unit value at the firm-product-market level. Export history controls include $\mathbf{a}_t^{ijk} \otimes \mathbf{s}_t^{ijk}$ and \mathbf{cens}_t^{ijk} . Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

elasticity of revenue with respect to the real exchange rate is equal to 0.46. Both elasticities are statistically different from zero. Their signs are as predicted by theory: revenue increases when tariffs in an export market fall, and when the Irish real exchange rate depreciates against a destination market. The two elasticities are statistically different from each other in absolute value. The absolute value of the elasticity of revenue with respect to tariffs is more than six times greater than the elasticity with respect to real exchange rates.

In contrast, for observations with high predicted exit probability, the coefficient on tariffs is not significantly different from zero, though the coefficient on the real exchange rate is very similar to that for the low exit probability observations. Our interpretation of this is that sample selection is an issue with respect to tariffs, but not with respect to exchange rates, consistent with the fact that entry is three times more responsive to tariffs than to exchange rates, and with the large (though insignificant) point estimate of the response of exit to tariffs.

As with entry and exit, we perform a large number of robustness tests, varying the specification, the sample, and measurement of the independent variables. Our findings are broadly robust to these variations. Full results are reported in the Appendix. Here we highlight a few points of interest.

We try a number of variations on our approach to dealing with potential selection bias. When we pool across low and high exit probability observations, unsurprisingly, the estimated coefficients are an average of the coefficients for the two subgroups. When we use spell length rather than market tenure to identify low exit probability observations, the re-

sults are very similar to the baseline. Results are also very similar when we include only indicators for market tenure as controls, rather than market tenure interacted with spell length.

As with entry and exit, estimation in first differences allows us to investigate the nuances of how revenue responds to shocks. However, lack of identifying variation in annual differences in tariffs seems to be an issue in this specification. We do not find a statistically significant response to current or future tariff changes. The coefficient on the real exchange is very similar to the baseline, though less precisely estimated. As with entry, responses to real exchange rates appear to be driven by appreciations.

We examine whether revenue elasticities differ systematically by firm size. As with export entry, export revenue in big firms is sensitive to tariffs, but export revenue in small firms is not. Elasticities with respect to real exchange rates differ (nonmonotonically) across firms of different sizes, but the differences are not statistically significant. It is not possible to separate out whether these differences are due to underlying differences in parameters for firms of different size, or to differences in the severity of selection bias.

As with entry and exit, there are various reasons to test whether our results are robust to changing the sample of countries included in the estimation. When we restrict the sample to non-Eurozone markets, our findings are consistent with the baseline. The same is true when we restrict the sample to non-EU markets.

Finally, we do not find evidence that export revenue responds differently to the nominal exchange rate and price components of real exchange rates.

6.3 Quantity and price

Can markup adjustment (due to, e.g. sticky prices) explain why the elasticity of export revenue with respect to tariffs is so different from the elasticity with respect to real exchange rates? We now show that it cannot.³⁰

Suppose that demand can be written:

$$Q_t^{ijk} = d\left(\frac{(1 + T_t^{jk}) \mu_t^{ijk} C_t^{ij}}{P_t^k}\right) = d(P_t^{ijk*}) \quad (13)$$

where $d(\cdot)$ is some decreasing function, μ_t^{ijk} is the endogenous markup over marginal cost for this firm-product-market triplet, and P_t^{ijk*} is the price customers in market k must pay for

³⁰Markup adjustment and sticky prices may be very important in other contexts, just not here.

good j sold by firm i in terms of the numeraire. The key assumption here is that a firm can only affect quantity demanded through changing its (relative) price, not through any other variable. Let θ_t^{ijk} be the (negative of the) price elasticity of demand. Then the following relationships must hold:

$$\frac{\partial \ln Q_t^{ijk}}{\partial \ln (1 + T_t^{jk})} = \underbrace{\left(-\frac{\partial \ln Q_t^{ijk}}{\partial \ln P_t^{ijk*}} \right)}_{\theta_t^{ijk}} \left(-1 - \frac{\partial \ln \mu_t^{ijk}}{\partial \ln (1 + T_t^{jk})} \right) \quad (14)$$

$$\frac{\partial \ln Q_t^{ijk}}{\partial \ln P_t^k} = \underbrace{\left(-\frac{\partial \ln Q_t^{ijk}}{\partial \ln P_t^{ijk*}} \right)}_{\theta_t^{ijk}} \left(1 - \frac{\partial \ln \mu_t^{ijk}}{\partial \ln P_t^k} \right) \quad (15)$$

If markup adjustment is to explain why revenue elasticities with respect to tariffs and real exchange rates are so different, it must be that quantity and markup elasticities with respect to tariffs and real exchange rates both deliver the same value for the price elasticity of demand, θ_t^{ijk} .

To obtain the quantity and markup elasticities from equations (14) and (15), we estimate equation (12), but with log quantity (i.e. tonnes) and log unit value (i.e. value/tonnes) in turn as the dependent variable. The results are reported in Columns 2 and 3 of Table 7. Focusing on the estimates for low exit probability observations, we find evidence of markup adjustment with respect to real exchange rates, but not with respect to tariffs.³¹ However substituting the estimates from the second two columns of Table 7 into equations (14) and (15), we see that quantity and markup elasticities with respect to tariffs imply a price elasticity of demand of -2.99, while quantity and markup elasticities with respect to real exchange rates imply a price elasticity of demand of -0.29. These two values are inconsistent. The reason they are inconsistent is because there is an order of magnitude difference in the responses of quantities to tariffs and real exchange rates *conditional on markup adjustment*. This implies that differences in markup responses to tariffs and real exchange rates cannot explain the differences we document in revenue elasticities.

This inconsistency points to an additional channel beyond their own price through which firms can affect their quantity demanded. Customer base is precisely such a channel. If

³¹The quantity and price elasticities do not add up to the revenue elasticity because the estimation samples are slightly different: quantities are not available for all observations. When we estimate the revenue equation on the sample for which quantities are available, we obtain the sum of the coefficients in the second two columns of Table 7. These results are reported in the Appendix.

demand is given by equation (1) from our model rather than (13), and if investment in customer base responds differently to tariffs than to real exchange rates, there will be an extra term in equations (14) and (15) which can explain why export quantity, and hence also export revenue respond differently to the two variables. We explore this further in Section 7 below.

6.4 Adding up: implications for aggregate exports

We now use our baseline estimates of entry, exit and revenue responses to tariffs and real exchange rates to calculate the implied elasticities of aggregate exports with respect to these variables. To do so, we take into account the estimated pattern of exporter dynamics from Section 4. The exercise we carry out is as follows. We consider two initially identical hypothetical export markets, and suppose that domestic firms face a reduction in the tariff from 10% to 0 in one of the markets, but not the other, or a 10% depreciation in the home currency against one of the markets, but not the other. We calculate what happens to aggregate exports to the affected market relative to the unaffected market, holding fixed domestic costs. We do this for each of the first six years after the shock. We also calculate long run responses. The elasticity of aggregate exports at horizon t with respect to the shock is then:

$$elasticity_{agg} = \frac{\ln EX_t^{change} - \ln EX_t^{nochange}}{\ln(1) - \ln(1.1)},$$

where *change* refers to the market where the variable of interest changes, and *nochange* refers to the market where it does not change.

The full details of the implementation, including how we make use of the estimated pattern of export dynamics, are reported in the Appendix. Table 8 reports the results. For each shock, the first column reports the elasticity of aggregate exports setting entry and revenue responses to their baseline point estimates, but setting exit responses to zero (as a reminder, the estimated exit responses are not significantly different from zero). The second column reports the elasticity of aggregate exports, with all responses set equal to their point estimates. The final two columns in each panel report the aggregate elasticities allowing only revenue responses, and only entry and exit responses respectively.

The elasticities reported in Table 8 are broadly in line with those estimated using aggregate data (see Ruhl (2008) for a survey). Whether or not we allow for exit responses, on impact, aggregate behavior is governed mainly by the response of export revenue conditional on participation. Over time, the contribution of participation responses grows. Setting exit

responses to zero, the elasticity of aggregate exports with respect to tariffs is six times bigger than the elasticity with respect to real exchange rates. This is mainly driven by the difference in revenue responses, as both entry responses and entrants are small.

Using the point estimates of exit responses rather than setting them equal to zero, the difference between tariff and real exchange rate elasticities is ten-fold. For tariffs, the point estimate of the exit response is as predicted by theory, so adding this response magnifies the elasticity of overall exports with respect to a shock. However for exchange rates, entry and exit responses go in opposite directions, and the two margins effectively cancel each other out.

Overall, whether or not we allow for exit responses, we find that the responses of aggregate exports to tariffs and real exchange rates owe more to the responses of incumbents on the intensive margin than to entry and exit responses. It is this which leads us to say that entry and exit responses are “modest.”

Table 8: Elasticities of aggregate exports with respect to shocks

	Tariff				Real ex. rate			
	Margins							
Horizon	entry, rev	entry, exit, rev	rev	entry, exit	entry, rev	entry, exit, rev	rev	entry, exit
1	-3.08	-3.33	-3.02	-0.31	0.47	0.45	0.46	-0.01
2	-3.16	-3.65	-3.02	-0.63	0.48	0.45	0.46	-0.01
3	-3.24	-3.94	-3.02	-0.92	0.50	0.45	0.46	-0.01
4	-3.32	-4.19	-3.02	-1.17	0.51	0.45	0.46	-0.01
5	-3.39	-4.41	-3.02	-1.39	0.52	0.45	0.46	-0.01
6	-3.46	-4.59	-3.02	-1.57	0.53	0.46	0.46	0.00
steady state	-3.76	-5.38	-3.02	-2.36	0.59	0.48	0.46	0.02

Notes: First column of each panel reports aggregate elasticities allowing entry and export revenue to respond. Second column reports aggregate elasticities allowing entry, exit and revenue to respond. Third column reports elasticities allowing only revenue to respond. Fourth column reports elasticities allowing only entry and exit to respond. For details of construction of aggregate elasticities, see text.

7 Simulation

As we have just illustrated, the response of export revenue is key to the behavior of aggregate exports. Moreover, from the behavior of quantities and prices, the differences in export revenue responses to tariffs and real exchange rates cannot be accounted for by differences in markup adjustment. So why is export revenue so much more responsive to tariffs than it is to real exchange rates? We now show that forward-looking behavior due to customer base, combined with the very different processes followed by tariffs and real exchange rates, has the potential to explain this behavior. To do so, we make use of the model from Section 2.

We use an off-the-shelf parameterization of the model, taken from Fitzgerald, Haller and Yedid-Levi (2016). Briefly, marginal cost is assumed constant, the same for all firms, while the process for the sunk cost is such that the rate of entry into exporting is exogenous.³² In addition, the following assumptions are made about customer base and learning: (1) there is a quadratic cost of adjusting investment in customer base, (2) investment in customer base is irreversible, and (3), firms learn probabilistically about the nature of their idiosyncratic demand. Fitzgerald, Haller and Yedid-Levi (2016) estimate this model to match annual moments of the post-entry dynamics of export quantities, prices and exit (similar to the moments we document in Section 4) and we work with their baseline parameter estimates.³³ Full details about functional forms and parameter values are reported in the Appendix.

Most of the identifying variation in tariffs in our data comes from reductions agreed in the Uruguay Round in 1994, and phased in gradually over the period 1995-2005.³⁴ To mimic the forecastability of tariff declines, in our simulation we assume that initially exporters face a tariff rate of 5%, and that they expect this to continue forever. However, at some date, it is announced unexpectedly that in one year, the tariff will fall to 4.5%, the year after it will fall to 4%, the year after that to 3.5% and so on and so on, until after 10 years, the tariff reaches 0, where it will remain forever. For the real exchange rate, we use an AR(1) process. We choose the persistence and innovation variance of this process by estimating an AR(1) using data on the real exchange rate between Ireland and the United States.³⁵

We then use the model to solve numerically for the policy functions in response to each of these two shocks. Based on these policy functions and the assumed processes, we simulate panels of data to mimic the actual data. The set of firms is assumed exogenous (equal to 5000) in these simulations. All firms start out as non-exporters, and we first simulate for 100 years to arrive at the ergodic distribution of export participation and exports. In year 101 we start the “trade agreement” exercise. We then extract the simulated data for periods 102 through 115, giving us a panel of 14 years, just as in the actual data. Similarly, in the exchange rate case, we extract the simulated data for years 102 through 115. We use the simulated data to construct the same variables and to estimate the same revenue equations as in the actual data. We implement 100 simulations for the tariff exercise, and 100 simulations for the real exchange rate exercise.

Table 9 reports the average elasticities across simulations, where the elasticities are the

³²Note that we find entry is relatively unresponsive to shocks.

³³The price elasticity of demand θ is not identified by the moments they target, but set equal to 2.

³⁴See Figure 1.

³⁵The model operates at a bi-annual frequency, so we estimate the AR(1) using bi-annual data constructed using quarterly data from IFS. The sample period for this estimation is 1980-2014.

Shock	Elasticity
Pre-announced tariff reduction	-11.71
Real exchange rate: AR(1), $\rho = 0.911$, $\sigma = 0.063$	3.09

Notes: Table reports average across 100 simulated data sets of estimated coefficients on low exit probability observations in revenue equations similar to equation (12).

estimated coefficients for low exit probability observations, defined as in the actual data. The average elasticity in response to the tariff shock is nearly four times as big in absolute value as the elasticity in response to the exchange rate shock. The level of both elasticities is much greater than in the actual data.³⁶ However, it is worth noting that the model ignores a lot of the richness in the data, and moreover, we tie our hands by taking parameter values direct from Fitzgerald, Haller and Yedid-Levi (2016), who do not target responses to shocks.

We conclude from this exercise that forward-looking behavior due to customer base has the potential to rationalize our empirical finding that export revenue is more responsive to tariffs than to real exchange rates. We think it is worth noting that this is despite the fact that the “real exchange rate” process we feed into the model has considerable persistence. Relative to Ruhl (2008), this draws attention not just to the persistence of shocks, but to their forecastability.

8 Conclusion

The main contribution of this paper is to use firm and customs micro data for Ireland to pinpoint exactly how firm behavior diverges from the prediction of workhorse models that exporters should respond similarly to tariffs and real exchange rates. We find that export revenue is six times more responsive to tariffs than it is to real exchange rates, and it is this, rather than differences in entry and exit responses, which accounts for most of the difference in aggregate responses.

We provide evidence that the difference in export revenue responsiveness to tariffs and real exchange rates is likely not due to sticky prices and variable markups, which have been the focus of a very large literature in international macroeconomics. Based on estimating elasticities of quantities and markups with respect to these two variables, we can see that the difference in export revenue responses is driven by much greater responsiveness of quantities to tariffs than real exchange rates, *conditional on markups*.

³⁶As in the actual data, the tariff coefficient differs substantially between low and high exit probability observations, but the real exchange rate coefficient does not.

We argue instead that the behavior of export revenue is likely due to a combination of two factors. First, demand depends on customer base as well as prices, customer base can be accumulated through non-price actions, and accumulation of customer base makes firms' decisions about how to respond to tariffs and real exchange rates forward-looking, rather than static, as assumed in the workhorse models. As a result, the process followed by shocks matters for how firms adjust their investment in customer base, and hence their export revenue. Second, tariffs and exchange rates follow very different processes. Tariffs are predictable, while real exchange rates are not. This implies that forward-looking firms increase customer base (and hence revenue) much more in response to favorable tariff changes than in response to favorable real exchange rate changes.

Returning to the policy questions that motivate our work, our findings suggest that "currency manipulation" is likely to have only a limited effect on international trade, because irrespective of current policy, firms see future real exchange rates as uncertain. In contrast, policy changes such as new trade agreements (or border adjustment of corporate taxes) where implementation is credible, and the bar for future policy change is very high can have much larger effects.

References

- [1] Arkolakis, C. (2010), "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy* 118 (6), 1151-1199.
- [2] Arkolakis, C., J. Eaton and S. Kortum (2012), "Staggered Adjustment and Trade Dynamics," mimeo.
- [3] 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.
- [4] Bustos, P. (2011), "Trade Liberalization, Exports and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review* 101 (1), 304-340.
- [5] Campa, J. (2004), "Exchange Rates and Trade: How Important is Hysteresis in Trade?" *European Economic Review* 48, 527-548.
- [6] Chamberlain, G. (1986), "Asymptotic efficiency in semiparametric model with censoring," *Journal of Econometrics* 32, 189-218.

- [7] Das, M., M. Roberts and J. Tybout (2007), "Market Entry Costs, Producer Heterogeneity and Export Dynamics," *Econometrica* 75 (3), 837-873.
- [8] Drozd, L. and J. Nosal (2012), "Understanding International Prices: Customers as Capital," *American Economic Review* 102 (1), 364-395.
- [9] Eaton, J., M. Eslava, M. Kugler and J. Tybout (2014), "A Search and Learning Model of Export Dynamics," mimeo.
- [10] Eaton, J., S. Kortum and F. Kramarz (2011), "An Anatomy of International Trade: Evidence from French Firms," *Econometrica* 1453-1498.
- [11] Feenstra, R., P. Luck, M. Obstfeld and K. Russ (2014), "In Search of the Armington Elasticity," NBER Working Paper 20063.
- [12] Fitzgerald, D., S. Haller and Y. Yedid-Levi (2016), "How Exporters Grow," NBER Working Paper 21935.
- [13] Foster, L., J. Haltiwanger and C. Syverson (2008), "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98 (1), 394-425.
- [14] Foster, L., J. Haltiwanger and C. Syverson (2016) "The Slow Growth of New Plants: Learning About Demand?" *Economica* 83, 91-129.
- [15] Freund, C. and J. Gagnon (2017), "Effects of Consumption Taxes on Real Exchange Rates and Trade Balances," Peterson Institute for International Economics Working Paper 2017-5.
- [16] Gourio, F. and L. Rudanko (2014), "Customer Capital," *Review of Economic Studies*, 81(3),1102-1136.
- [17] Imbs, J. and I. Mejean (2015), "Elasticity Optimism," *American Economic Journal: Macroeconomics* 7 (3), 43-83.
- [18] International Monetary Fund: *International Financial Statistics*
- [19] ISGEP (International Study Group on Exports and Productivity) (2008), "Understanding Cross-Country Differences in Exporter Premia - Comparable Evidence for 14 Countries," *Review of World Economics* 144 (4), 596-635.

- [20] 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.
- [21] Organisation for Economic Co-operation and Development: *National Accounts Statistics*
- [22] 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.
- [23] 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.
- [24] Ruhl, K. (2008), “The International Elasticity Puzzle,” mimeo.
- [25] Ruhl, K. and J. Willis (2017), “New Exporter Dynamics,” *International Economic Review*, forthcoming.
- [26] Staiger, R. and A. Sykes (2010), “‘Currency Manipulation’ and World Trade,” *World Trade Review* 9(4), 583-627.
- [27] van Beveren, I., A. Bernard and H. Vandenbussche (2012), “Concording EU Trade and Production Data over Time,” NBER Working Paper 18604.
- [28] World Bank: *World Development Indicators*
- [29] World Trade Organization: *Tariff Download Facility*