NBER WORKING PAPER SERIES

HOW EXPORTERS GROW

Doireann Fitzgerald
Stefanie Haller
Yaniv Yedid-Levi

Working Paper 21935
http://www.nber.org/papers/w21935

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2016

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 that 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 Adrian Corcoran, Matt Shapiro and Anthony Priolo. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. Yaniv Yedid-Levi is grateful for financial support from the Social Sciences and Humanities Research Council of Canada. We thank George Alessandria, Costas Arkolakis, Kim Ruhl, Ana Maria Santacreu, James Tybout and Daniel Xu for comments and suggestions. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Doireann Fitzgerald, Stefanie Haller, and Yaniv Yedid-Levi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Exporters Grow
Doireann Fitzgerald, Stefanie Haller, and Yaniv Yedid-Levi
NBER Working Paper No. 21935
January 2016, Revised January 2017
JEL No. E2,F1,L1

ABSTRACT

We document how export quantities and prices evolve after entry to a market. Controlling for marginal cost, and taking account of selection on idiosyncratic demand, there are economically and statistically significant dynamics of quantities, but no dynamics of prices. To match these facts, we estimate a model where firms invest in customer base through non-price actions (e.g. marketing and advertising), and learn gradually about their idiosyncratic demand. The model matches quantity, price and exit moments. Parameter estimates imply costs of adjusting investment in customer base, and slow learning about demand, both of which generate sluggish responses of sales to shocks.

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

Yaniv Yedid-Levi
University of British Columbia
yaniv@mail.ubc.ca

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

A online appendix is available at http://www.nber.org/data-appendix/w21935

How Exporters Grow*

Doireann Fitzgerald† Stefanie Haller‡ Yaniv Yedid-Levi§

December 2016

Abstract

We document how export quantities and prices evolve after entry to a market. Controlling for marginal cost, and taking account of selection on idiosyncratic demand, there are economically and statistically significant dynamics of quantities, but no dynamics of prices. To match these facts, we estimate a model where firms invest in customer base through non-price actions (e.g. marketing and advertising), and learn gradually about their idiosyncratic demand. The model matches quantity, price and exit moments. Parameter estimates imply costs of adjusting investment in customer base, and slow learning about demand, both of which generate sluggish responses of sales to shocks.

1 Introduction

There is a substantial literature showing that firms face costs of adjusting inputs such as capital and labor in response to shocks. We use export micro-data to provide evidence that firms also face *demand-side* costs of adjustment. Export micro-data allows us to track a firm's history of participation in different markets, and to measure quantities and prices in

*Online appendix available at www.doireann.com. 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 that 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 Adrian Corcoran, Matt Shapiro and Anthony Priolo. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. Yaniv Yedid-Levi is grateful for financial support from the Social Sciences and Humanities Research Council of Canada. We thank George Alessandria, Costas Arkolakis, Kim Ruhl, Ana Maria Santacreu, James Tybout and Daniel Xu for comments and suggestions. 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 and NBER, doireann.fitzgerald@gmail.com

‡School of Economics, University College Dublin, stefanie.haller@ucd.ie

§Vancouver School of Economics, University of British Columbia, yaniv@mail.ubc.ca

each, and this turns out to be very informative about the nature of these frictions. Viewed through the lens of a model that fits the data very well, we find that both costs of adjusting investment in customer base and slow learning about idiosyncratic demand play an important role in explaining post-entry exporter dynamics. Together, these two frictions imply sluggish adjustment of sales in response to shocks.

Our first contribution is to use matched firm and customs data for Ireland to document a novel set of facts about the dynamics of quantities and prices after export entry. In order to isolate the market-specific demand channel, we focus on variation in quantities and prices across different markets (i.e. countries) within a given firm-product-year triplet. Selection on unobserved heteregeneity in demand can give rise to the appearance of dynamics when there are none. We isolate true dynamics by differentiating between what happens in export episodes that turn out to be brief, and those that end up lasting many years.[1]

We show that (1) conditional on survival, export quantities grow substantially in a firm's first five years in a market; (2) there are no statistically or economically significant dynamics of prices with respect to number of years in a market; (3) there is an increasing relationship between quantities sold in the first year in a market, and the number of years the firm eventually survives in that market; (4) prices in the initial year in a market do not forecast survival. In sum, there are statistically and economically significant post-entry dynamics in quantities, but no dynamics in prices. In addition, as others have shown,[2] conditioning on marginal cost, the probability of exit is decreasing with the number of years a firm has been in a market. These patterns are robust, and hold across different industries, across multinationals and domestic firms, across firms of different sizes, and across different export markets.

Since we focus on variation across markets within a firm, product and year, the post-entry dynamics of quantities cannot be generated by dynamics in productivity, capital, labor or financial constraints. Supply-side factors that are market-specific could rationalize quantity dynamics, but have a hard time doing so without generating price dynamics. This suggests that dynamics arise from the demand side.[3] However the joint behavior of quantities and

[1] We build on Roberts and Tybout (1997), Das, Roberts, and Tybout (2007), Eaton et al. (2008), Eaton et al. (2014), and Ruhl and Willis (2016), who use data for Colombia to document facts about export entry, exit and post-entry export dynamics, but not about quantities and prices separately. Our findings on revenues and exit are similar to theirs. Contemporaneous work by Bastos et al. (2015), Piveteau (2016), and Berman et al. (2015) document some aspects of quantity and price dynamics using customs data for Portugal and France.

[2] See e.g. Ruhl and Willis (2016).

[3] Recent years have seen increased interest in the role of demand in explaining firm dynamics. See e.g. Foster, Haltiwanger, and Syverson (2008, 2016), Gourio and Rudanko (2014).

prices is inconsistent with a number of popular demand-based explanations for post-entry exporter dynamics and related models of the role of demand in firm dynamics. In particular, the absence of post-entry price dynamics is inconsistent with models where current customer base, and hence demand, is determined by past sales, as these models imply an increasing path of prices as firms acquire customers.[4] Meanwhile, the fact that initial prices do not forecast survival in an export market is inconsistent with models where quantity growth is generated by firms learning about uncertain idiosyncratic demand, and setting quantities rather than prices in the face of that uncertainty.[5]

Our second contribution is to propose and estimate a model of post-entry export dynamics that can match the joint behavior of quantities, prices and exit in the data. The model has two key features. First, firms accumulate customer base by undertaking *non-price* actions, e.g. investing in marketing and advertising. They face adjustment costs on this investment.[6] Second, firms learn about their idiosyncratic demand, which has exogenous dynamics.[7] Firms set prices rather than quantities in the face of demand uncertainty. They face constant price elasticity of demand, and current prices do not affect future demand, so markups are constant.[8] This model has the ability to match qualitatively all the facts we document. The combination of learning about idiosyncratic demand and price setting generates a positive correlation between initial quantities and survival conditional on marginal cost, but no relationship between initial prices and survival. Learning also generates a downward-sloping exit hazard. The combination of customer base accumulation, adjustment costs of investment, and learning about demand generates quantity growth conditional on marginal costs and survival. Price setting, independence of future demand from current prices and constant price elasticity of demand deliver the absence of price dynamics accompanying quantity dynamics.

[4]This channel operates in Ravn, Schmitt-Grohe and Uribe (2006) and Foster, Haltiwanger and Syverson (2016) and is applied to exporter dynamics by Piveteau (2016).

[5]See Jovanovic (1982), where firms learn about marginal cost. Albornoz et al. (2012), Berman et al. (2015), Timoshenko (2015a, 2015b) and Bastos et al. (2016) propose that post entry export dynamics are due to learning about idiosyncratic demand in an environment where firms set quantities rather than prices. Arkolakis et al. (2015) have a model of firm dynamics with this feature.

[6]This extends Arkolakis (2010), who has a static model with convex costs of acquiring customers in export markets, to a dynamic setting. Arkolakis (2010) and Gourio and Rudanko (2014) report direct evidence on the importance of marketing and advertising expenditures.

[7]We model learning as follows: Firms can be uninformed (know nothing about the history of idiosyncratic demand) or informed (observe the full history). By participating in a market, a firm gets access to a Poisson probability of becoming informed. Being informed is an absorbing state. The probability of transitioning from being uninformed to informed governs the speed of learning, which does not depend on the process for idiosyncratic demand.

[8]This model shares many of the features of Eaton et al. (2014) but does not require transaction-level data to estimate.

We estimate the parameters of the model using simulated method of moments. We target the quantity and exit moments from the data, since price moments are matched by construction. The model fits all moments remarkably well. It matches the relationship between initial quantities and survival; the growth of quantities in short and long export episodes; and the behavior of the exit hazard. The estimated parameters imply nontrivial costs of adjusting investment in customer base as well as gradual learning about idiosyncratic demand, both of which slow down demand-side adjustment to shocks.

We illustrate the role of different features of the model by shutting them down one-by-one and re-estimating. Customer base, adjustment costs of investing in customer base, and slow learning about demand are all key to matching the growth of quantities conditional on survival. Slow learning also helps to match the relationship between initial quantities and survival and the exit hazard. Exogenous dynamics of idiosyncratic demand are necessary to match the fact that quantities grow faster in ultimately successful export episodes than in ultimately unsuccessful episodes. In the baseline model, learning is independent of the process for idiosyncratic demand. We also estimate a model where firms are Bayesian and use the Kalman filter to learn about their idiosyncratic demand based on the history of realized quantities. The parameters of the idiosyncratic demand process necessary to match the behavior of quantities imply very rapid learning. As a result, this model cannot match the behavior of the exit hazard, and has a poor fit overall relative to the baseline model.

Finally, we simulate export responses to tariff changes using our estimated model, examining in particular the difference between short run and long run responses. In doing so we intersect with an empirical literature that documents differences between short and long run elasticities of exports with respect to prices,[9] and a quantitative literature that examines whether sunk costs of export participation can generate these differences.[10] In contrast to the case where sunk costs of export entry are the only adjustment friction, our estimated model implies that large incumbents as well as potential entrants are subject to frictions, due to costs of adjusting customer base. In our simulations, holding constant the rate of entry, and general equilibrium factors such as costs, foreign demand and prices (important channels for generating dynamics in the literature just cited), long run export elasticities are 1.3 times larger than short run elasticities. To put this in context, Gallaway, McDaniel and Rivera (2003) find that long-run elasticities are twice as big as short-run elasticities.

Our findings are of interest to the literature on adjustment costs and business cycles, as we provide evidence of an additional source of adjustment costs beyond the standard costs of

[9]See e.g. Hooper, Johnson and Marquez (2000) and Gallaway, McDaniel and Rivera (2003).

[10]See e.g. Alessandria and Choi (2014) and Alessandria, Choi and Ruhl (2015).

adjusting capital and labor. In particular, our findings have important implications for open economy business cycles, as market-specific adjustment costs have much greater scope to affect the joint dynamics of international relative prices and quantities than adjustment costs at the level of the plant or firm. To date, this field has typically focused on the implications of costs of adjusting prices.[11] We provide evidence that costs of adjusting quantities at the market level are quantitatively important. Our findings have important implications for dynamic adjustment to changes in trade policy, as we illustrate in our simulation exercises. We develop the implications of our findings for the exchange rate disconnect puzzle in a follow-up paper (Fitzgerald, Yedid-Levi and Haller (2016)).

The paper is organized as follows. In the second section, we describe our data. In the third section, we outline our empirical strategy. In the fourth section, we present our results on the post-entry dynamics of quantities, prices, products, and exit, and relate them to findings in the literature. In the fifth section, we lay out our model. In the sixth section, we describe how we estimate this model, report our estimation results, and the results from estimating alternative models. In the seventh section, we report the results of our simulation exercises. The final section concludes.

2 Data description

We make use of two sources of confidential micro data made available to us by the Central Statistics Office (CSO) in Ireland: the Irish Census of Industrial Production (CIP) and Irish customs records. Here, we note the key points about each data set. The data are described in detail in the online appendix; the appendix also describes a third data set (the PRODCOM survey), which we use to obtain the number of products produced at the firm level and firm-product prices used in robustness checks.

2.1 Census of Industrial Production

The CIP is an annual census of manufacturing, mining, and utilities. Firms with three or more persons engaged are required to file returns.[12] We make use of data for the years 1996-2009 and for NACE Revision 1.1 sectors 10-40 (manufacturing, mining, and utilities). Of the variables collected in the CIP, those we make use of in this paper are the country of

[11] A notable exception is Drozd and Nosal (2012).

[12] Multiplant firms also fill in returns at the level of individual plants. We work with the firm-level data, since this is the level at which the match with customs records can be performed.

ownership, total revenue, employment, and an indicator for whether the firm has multiple plants in Ireland.

In constructing our sample for 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, for example, in cases in which ownership changes.

2.2 Customs records

Our second source of data is customs records of Irish merchandise exports for the years 1996-2014. The value (euros) and quantity (tonnes)[13] of exports are available at the level of the VAT number, the Combined Nomenclature (CN) eight-digit product, and the destination market (country), aggregated to an annual frequency. These data are matched by the CSO to CIP firms using a correspondence between VAT numbers and CIP firm identifiers, along with other confidential information. The online appendix provides additional information on this match.

A key feature of customs in the European Union is that data for intra-EU and extra-EU trade are collected separately, using two different systems called Intrastat and Extrastat. The threshold for mandatory reporting of intra-EU exports (635,000 euro per year in total shipments within the EU) is different from the threshold for extra-EU exports (254 euro per transaction).[14] The high threshold for intra-EU exports likely leads to censoring of exports by small exporters to the EU. However it applies not at the market level, but to exports to the EU as a whole, and we observe many firms exporting amounts below the 635,000 euro threshold to individual EU markets.

An important feature of the customs data is that the eight-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, and Vandenbussche (2012).[15] For our baseline analysis, we restrict attention to the period 1996-2009, for which we have CIP data in addition to customs data, and we make use only of customs data that matches to a CIP firm. In some robustness checks, we make use of the full sample period, 1996-2014, and all of the customs data irrespective of a CIP match. We

[13] The value is always available, but the quantity is missing for about 10% of export records.

[14] Intra-EU exports below the threshold are recovered based on VAT returns. The destination market within the EU is not recorded for these returns.

[15] Van Beveren, Bernard, and Vandenbussche (2012) show that once the data are appropriately concorded, there is less product churn than naive calculations based on raw data would suggest.

Table 1: Summary statistics: Firms and exports, averages 1996-2009

Mean number of firms per year	4748
Mean employees	50
Mean age (years)	17
Share of firms foreign owned	0.12
Share of multi-plant firms	0.03
Mean number of concorded products per firm	4
Share of firms exporting	0.44
Exporter size premium (employees, mean)	1.65
Exporter size premium (revenue, mean)	1.85
Mean export share conditional on exporting	0.32
Mean number of markets per exporter	6.6

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. Source: CSO and authors' calculations.

perform the product concordance separately for the two different sample periods, as dictated by the Pierce and Schott approach.

As a result, we have annual data on value and quantity of exports at the firm-product-market level, where the product is defined at the eight-digit (concorded) level, and the market refers to the destination country. We use this to construct a price (unit value) by dividing value by quantity, where available. In aggregate trade statistics, unit value data at the product level are notoriously noisy. However, conditioning on the exporting firm as well as the product considerably reduces this noise.

2.3 Summary statistics

Table 1 shows summary statistics on the firms in our data, focusing in particular on exporting behavior. Export participation is high, export intensity conditional on participation is high, and at least half of exporters participate in multiple markets (we observe 140 distinct export markets over the course of the panel). These facts are typical of small open European economies. Apart from the relatively high rate of export participation and the high intensity of exporting conditional on participation, the broad facts about exporting are also similar to those documented for large developed countries such as the United States and France and for developing countries such as Colombia.

Entry and exit are not synchronized across different export markets within a given firm.[16] This is illustrated in Table 2, which reports summary statistics on churn in the number of

[16]This is consistent with Lawless (2009), who uses a different data set on Irish firms.

Table 2: Summary statistics: percentage of exporters by change in number of markets year to year

Change	<-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	>6
%	2	1	2	3	5	11	51	12	5	3	2	1	3

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 revenue is concorded product export revenue from customs data. There are 140 export markets. Source: CSO and authors' calculations.

export markets from year to year. This is a lower bound on churn, as some firms may keep the total number of export markets constant, while switching markets. In any given year, on average 49% of exporters change the *number* of markets they participate in. This churn induces within-firm-year variation in market tenure and completed export spell length, which we exploit in our empirical strategy.

3 Empirical strategy

The goal of our empirical analysis is to identify dynamics in quantities and prices that are systematically related to how long a firm has been selling a product in a market. To isolate dynamics associated with demand, we control for time-varying heterogeneity at the firm-product level, e.g. in marginal cost[17] using firm-product-year fixed effects. In order to identify true dynamics, we must also deal with selection within the firm-product-year on unobserved heterogeneity in idiosyncratic demand. We do this by differentiating between export episodes according to how long they last, and estimating dynamics separately for episodes of different ex-post duration. This generalizes the approach of Ruhl and Willis (2016), and is similar to the approach of Altonji and Shakotko (1987) to dealing with selection in estimating the effect of job tenure on wages.

3.1 Product revenue, quantity, price, and product-market exit

First, we construct two key variables which we call *market tenure* and *spell length*. This is illustrated in Table 3. The top panel gives a (fictitious) example of the pattern of participation of a firm-product pair in markets A through G over a period of six years. We define an export *spell* as a continuous episode of market participation.

[17]Heterogeneity in marginal cost could arise from differences in productivity, or constraints such as capital adjustment costs and financial constraints, all of which have been posited as potential sources of firm dynamics. Our empirical approach also controls for heterogeneity in demand due to quality or product appeal that affects all markets equally, though for brevity, we may refer only to marginal cost.

Table 3: Illustration of identifying variation in market tenure and spell length

Year	1	2	3	4	5	6
Market	I. Participation					
A	X	X	X	X	X	X
B		X	X	X		
C			X	X	X	
D		X	X			
E		X		X		
F				X	X	X
G	X					
Market	II. Market tenure					
A	cens	cens	cens	cens	cens	cens
B		1	2	3		
C			1	2	3	
D		1	2			
E		1		1		
F				1	2	3
G	cens					
Market	III. Spell length					
A	cens	cens	cens	cens	cens	cens
B		3	3	3		
C			3	3	3	
D		2	2			
E		1		1		
F				cens	cens	cens
G	cens					
Market	IV. Spell length, topcoded at 3					
A	cens	cens	cens	cens	cens	cens
B		3	3	3		
C			3	3	3	
D		2	2			
E		1		1		
F				3	3	3
G	cens					

In the second panel, we show how market tenure is constructed. We set market tenure equal to 1 in the first year a firm exports a given product to a given market after not exporting in the previous period. Note that we do not observe market tenure if entry is censored (e.g., markets A and G in Table 3). Tenure is incremented by 1 in each subsequent year of continuous participation. If the firm-product exits a market for some period, market tenure is reset to 1 in the first subsequent year of participation (e.g., market E in year 4).

The third panel shows how we construct spell length. If the spell is neither left- nor right-censored, we observe completed spell length (markets B, C, D, E). If we observe zero exports for one or more years after some positive exports, any reentry is counted as part of a distinct export spell (e.g., market E).[18] The fourth panel shows that by top-coding spell length at some number, we can assign a spell length to some right-censored spells (e.g., market F, where completed spell length is at least 3).

As Table 3 illustrates, there can be cross-sectional variation in both spell length and market tenure within a firm-product-year. This allows us to use fixed effects to control for time-varying heterogeneity at the firm-product level e.g. in marginal cost. By further focusing only on variation in market tenure within spells of a given length, we can isolate true dynamics from the effect of selection on unobserved idiosyncratic demand. Finally, variation across spells of different length conditional on market tenure of 1 gives us some information on the degree of unobserved heterogeneity in idiosyncratic demand that may be driving selection.

Mechanically, we implement this as follows. Let w_t^{ijk} be log revenue, log quantity, or log price. Let $\mathbf{a}_t^{ijk}$ be a vector of indicator variables for firm i's tenure in market k with product j. Let $\mathbf{s}_t^{ijk}$ be a vector of indicators for the length of the relevant spell. This indicator does not vary within a spell, but is indexed by t to capture the fact that we may observe multiple export spells of different length for firm i, product j, and market k over the period of our panel (e.g., market E in Table 3). We top-code both market tenure and spell length at seven years in our baseline specification. We drop spells whose length is right-censored at a level below the top-code.[19] We also include a separate indicator ($cens^{ijk}$) for spells that are both left- and right-censored, as including these spells helps control for marginal cost. We then

[18] In our baseline analysis we treat these "reentry" spells the same as "first entry" spells. In robustness checks, we relax this and treat them differently.

[19] Allowing the full range of market tenures and spell lengths would force us to throw out all right-censored spells, would not allow us to separately identify the impact of market tenure and spell length for the longest spells, and would also confound cohort effects with the impact of these variables. Using our full panel of customs data, which lasts for 19 years, we show that our key results are robust to top-coding at 10 years.

estimate:

$$w_t^{ijk} = \delta^k + c_t^{ij} + \boldsymbol{\beta}' \left(\mathbf{a}_t^{ijk} \otimes \mathbf{s}_t^{ijk} \right) + cens^{ijk} + \varepsilon_t^{ijk}. \tag{1}$$

Here, δ^k is a set of market dummy variables (our baseline results are robust to generalizing this to market-year or product-market-year fixed effects). These control for differences across markets in aggregate demand and average prices. The term c_t^{ij} indicates the firm-product-year fixed effects that control for marginal cost. The symbol $\otimes$ indicates the Kronecker product. Of course, we do not observe tenures of greater than s for a spell that lasts exactly s years, so the redundant interactions are dropped.

The vector $\boldsymbol{\beta}$ contains the coefficients of interest. Exponentiated, appropriate linear combinations of the elements of $\boldsymbol{\beta}$ allow us to characterize variation in initial revenue, quantity, and price with completed spell length, and the evolution of revenue, quantity, and price with market tenure over the lifetime of spells of different length.

Our second empirical exercise examines the hazard of exit conditional on marginal cost. This allows us to characterize the distribution of spell lengths. We adopt a similar strategy to the above to show how exit varies with market tenure, exploiting again only variation within a firm-product-year. Let X_t^{ijk} be an indicator for participation of firm i with product j in market k at date t. We then estimate the linear probability model:

$$\Pr \left[X_{t+1}^{ijk} = 0 | X_t^{ijk} = 1 \right] = \delta^k + c_t^{ij} + \boldsymbol{\beta}' \mathbf{a}_t^{ijk} + \varepsilon_t^{ijk}. \tag{2}$$

The terms δ^k, c_t^{ij} and $\mathbf{a}_t^{ijk}$ are as above, and $\boldsymbol{\beta}$ is again the vector of coefficients of interest.[20]

3.2 Market revenue, number of products, and market exit

At the firm-market level, we observe revenue and the number of products a firm sells to a destination. This allows us to characterize the extent to which overall revenue dynamics depend on dynamics in the number of products. As described above, we exploit variation in market tenure and spell length, this time within a firm-year, to characterize how these variables evolve with market tenure, controlling for marginal cost and selection.

The construction of market tenure and spell length at the firm-market level is analogous to the approach at the firm-product-market level. Let w_t^{ik} be log revenue or log number of products. Let $\mathbf{a}_t^{ik}$ be a vector of indicator variables for firm i's tenure in market k. Let $\mathbf{s}_t^{ik}$ be a vector of indicators for the total length of the relevant spell. Let $cens^{ik}$ be an indicator

[20] Linearity is clearly less defensible for exit than log-linearity for revenue, quantity and price. Our estimates do not purport to be structural, but will be used as moments that we will ask our structural model to match.

for spells that are both left- and right-censored. We then estimate:

$$w_t^{ik} = \delta^k + c_t^i + \boldsymbol{\beta}' \left(\mathbf{a}_t^{ik} \otimes \mathbf{s}_t^{ik}\right) + cens^{ik} + \varepsilon_t^{ik}. \tag{3}$$

As above, δ^k is a set of market dummy variables, and c_t^i is a set of firm-year fixed effects.

Exponentiated, linear combinations of the elements of $\boldsymbol{\beta}$ allow us to characterize both variation in initial revenue and number of products with completed spell length, and the evolution of revenue and number of products with market tenure over the lifetime of spells of different length.

To characterize the distribution of spell length at the firm-market level, we adopt a similar strategy. Let X_t^{ik} be an indicator for participation of firm i in market k at date t. We then estimate:

$$\Pr\left[X_{t+1}^{ik} = 0 | X_t^{ik} = 1\right] = \delta^k + c_t^i + \boldsymbol{\beta}' \mathbf{a}_t^{ik} + \varepsilon_t^{ik}. \tag{4}$$

4 Empirical findings

4.1 Product revenue, quantity, price, and product-market exit

The first three columns of Table 4 report results for the baseline estimation of equation (1), with log revenue, log quantity, and log price in turn as the dependent variable.[21] The omitted category in all regressions is export spells which last exactly one year. The log of the dependent variable for each of these spells is hence normalized to 0. In the table, we organize our results into initial values conditional on spell length and within-spell trajectories normalizing the start of each spell to 0. Figures 1, 2, and 3, graph the trajectories of revenues, quantities and prices implied by taking the exponential of the relevant sums of coefficients from Table 4.[22] Table A.5 in the online appendix compares summary statistics on the firm-years included in this analysis with summary statistics for all firm-years in our data. Exporters included in the analysis of product revenue, quantity, and price are bigger, more export-intensive, and export to more destinations than the average exporter.

[21] In the revenue equation, we include only firm-product-market-years for which quantity data are available, so the sample is identical in the first three columns.

[22] We graph the standard errors for all revenue and quantity trajectories, but for price trajectories, we graph only the standard errors on the longest spell to make the figure easier to read. None of the points on the price trajectories are significantly different from 1.

Table 4: Dynamics of revenue, quantity, price, and number of products

Obs. level	Firm-product-market						Firm-market			
Dep. var. (ln)	Revenue		Quantity		Price		Revenue		# Products	
Spell lgth	Spell intercept									
2 years	0.51	(0.02)**	0.52	(0.02)**	-0.01	(0.01)	0.40	(0.04)**	0.10	(0.01)**
3 years	0.76	(0.03)**	0.76	(0.04)**	0.00	(0.02)	0.74	(0.06)**	0.15	(0.01)**
4 years	0.95	(0.05)**	0.95	(0.05)**	0.00	(0.02)	0.84	(0.07)**	0.18	(0.02)**
5 years	1.07	(0.06)**	1.08	(0.07)**	-0.01	(0.03)	1.09	(0.09)**	0.19	(0.02)**
6 years	1.13	(0.08)**	1.09	(0.08)**	0.04	(0.03)	1.15	(0.11)**	0.25	(0.03)**
7+ years	1.39	(0.05)**	1.39	(0.05)**	0.01	(0.02)	1.32	(0.05)**	0.28	(0.01)**
cens	3.66	(0.03)**	3.70	(0.03)**	-0.04	(0.01)**	3.98	(0.03)**	0.91	(0.01)**
Mkt tenure	2-year spell									
2 years	-0.03	(0.03)	-0.03	(0.03)	0.00	(0.02)	-0.02	(0.05)	-0.00	(0.01)
Mkt tenure	3-year spell									
2 years	0.44	(0.04)**	0.45	(0.05)**	-0.01	(0.02)	0.48	(0.07)**	0.11	(0.02)**
3 years	-0.05	(0.05)	-0.05	(0.05)	0.00	(0.02)	0.02	(0.07)	0.01	(0.02)
Mkt tenure	4-year spell									
2 years	0.53	(0.06)**	0.55	(0.06)**	-0.02	(0.03)	0.61	(0.09)**	0.13	(0.02)**
3 years	0.55	(0.06)**	0.60	(0.06)**	-0.05	(0.03)*	0.57	(0.09)**	0.12	(0.02)**
4 years	-0.02	(0.07)	-0.01	(0.07)	-0.01	(0.03)	0.19	(0.10)*	0.01	(0.02)
Mkt tenure	5-year spell									
2 years	0.63	(0.09)**	0.62	(0.09)**	0.01	(0.04)	0.71	(0.12)**	0.16	(0.03)**
3 years	0.70	(0.09)**	0.69	(0.09)**	0.01	(0.04)	0.74	(0.12)**	0.19	(0.03)**
4 years	0.57	(0.09)**	0.61	(0.09)**	-0.04	(0.04)	0.59	(0.12)**	0.19	(0.03)**
5 years	-0.01	(0.09)	0.01	(0.09)	-0.02	(0.04)	0.05	(0.12)	0.05	(0.03)
Mkt tenure	6-year spell									
2 years	0.74	(0.11)**	0.78	(0.11)**	-0.04	(0.05)	0.68	(0.14)**	0.21	(0.04)**
3 years	0.87	(0.11)**	0.95	(0.11)**	-0.07	(0.05)	0.90	(0.14)**	0.21	(0.04)**
4 years	0.85	(0.11)**	0.92	(0.11)**	-0.07	(0.05)	1.03	(0.14)**	0.24	(0.04)**
5 years	0.71	(0.11)**	0.75	(0.11)**	-0.04	(0.05)	0.75	(0.14)**	0.14	(0.04)**
6 years	0.12	(0.11)	0.14	(0.11)	-0.02	(0.05)	0.11	(0.15)	0.00	(0.04)
Mkt tenure	7+ year spell									
2 years	0.85	(0.06)**	0.88	(0.06)**	-0.03	(0.03)	1.01	(0.07)**	0.21	(0.02)**
3 years	1.16	(0.06)**	1.20	(0.06)**	-0.03	(0.03)	1.35	(0.07)**	0.28	(0.02)**
4 years	1.31	(0.06)**	1.34	(0.06)**	-0.03	(0.03)	1.51	(0.07)**	0.32	(0.02)**
5 years	1.34	(0.06)**	1.37	(0.06)**	-0.04	(0.03)	1.60	(0.07)**	0.33	(0.02)**
6 years	1.30	(0.06)**	1.33	(0.07)**	-0.03	(0.03)	1.59	(0.07)**	0.32	(0.02)**
7+ years	1.28	(0.06)**	1.35	(0.06)**	-0.07	(0.03)**	1.64	(0.06)**	0.33	(0.02)**
	Fixed effects									
Firm-prod-yr	Yes		Yes		Yes		No		No	
Firm-yr	No		No		No		Yes		Yes	
Market	Yes		Yes		Yes		Yes		Yes	
N	312952		312952		312952		113912		113912	
rsq	0.76		0.82		0.90		0.65		0.56	
rsq-adj	0.58		0.69		0.82		0.58		0.47	

Notes: Dependent variable is in turn log revenue, log quantity, and log unit value at the firm-product-market-year level, and log revenue and log number of products at the firm-market-year level. Full set of firm-product-year and market effects included in firm-product-market-year regressions. Full set of firm-year and market effects included in firm-market-year regressions. Omitted category is spells that last one year. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

Figure 1: Firm-product-market revenue by completed spell length and market tenure

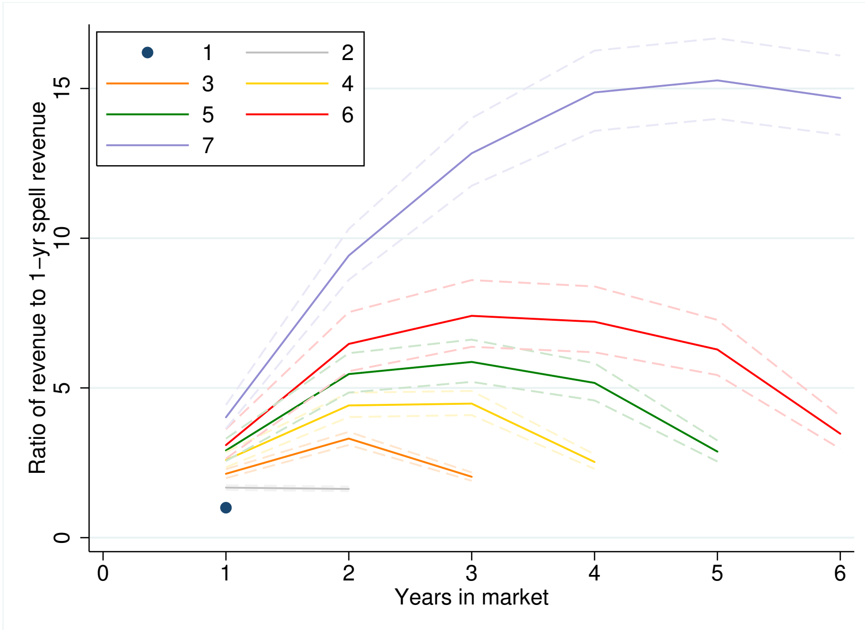

Notes: Figure shows evolution of revenue at the firm-product-market level with market tenure, allowing trajectories to differ by export spell length. Trajectories are conditional on firm-product-year and market effects. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

Figure 2: Firm-product-market quantity by completed spell length and market tenure

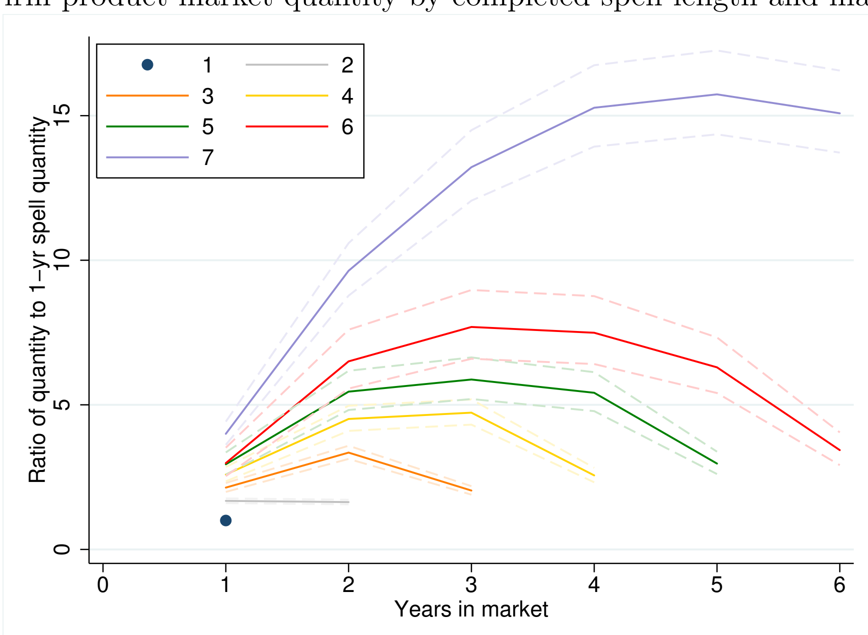

Notes: Figure shows evolution of quantities at the firm-product-market level with market tenure, allowing trajectories to differ by export spell length. Trajectories are conditional on firm-product-year and market effects. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

Figure 3: Firm-product-market price by completed spell length and market tenure, different scale

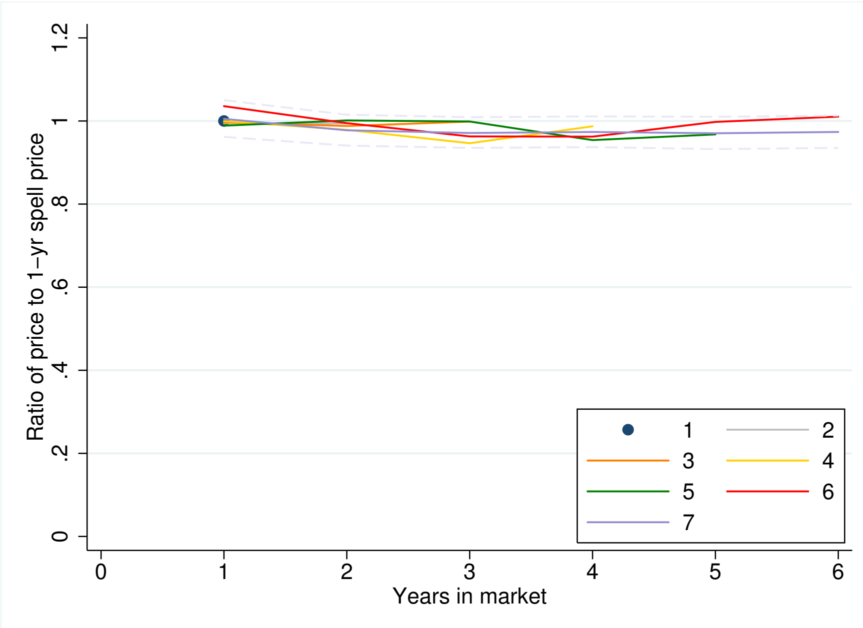

Notes: Figure shows evolution of prices at the firm-product-market level with market tenure, allowing trajectories to differ by export spell length. Trajectories are conditional on firm-product-year and market effects. 95% confidence interval for spells of 7+ years is plotted. Source: CSO and authors' calculations.

There are four key findings on quantities and prices: (1) Quantities grow dramatically in the first five years of successful export spells, defined as spells that last at least seven years. This growth is statistically significant up to a horizon of four years and is not driven purely by part-year effects in the first year (i.e. there is economically and statistically significant growth between years 2 and 4). (2) Within successful export spells, up to a horizon of six years, there are no statistically or economically significant dynamics in prices. (3) Higher initial quantities predict longer export spells: for spells lasting between one and four years, all pairwise comparisons of initial quantities are statistically different. (4) Initial prices do not predict export spell length.

Additional findings on quantities and prices are as follows. In "unsuccessful" spells, quantities initially rise and subsequently fall, though they are never observed to fall much below quantities in the first year in the market: the difference is economically small and never statistically significant. With one exception (the third year of four-year spells) there are no statistically significant dynamics of prices with tenure in these spells.

In spells that are both left- and right-censored, quantities are an order of magnitude larger (40 times larger) than in one-year spells, whereas prices are marginally lower (4% lower) than in one-year spells (both of these differences are statistically significant). These

Table 5: Exit hazard

Market tenure	Firm-prod-mkt		Firm-mkt	
2 years	-0.13	(0.00)**	-0.16	(0.00)**
3 years	-0.20	(0.00)**	-0.22	(0.01)**
4 years	-0.24	(0.00)**	-0.25	(0.01)**
5 years	-0.25	(0.01)**	-0.27	(0.01)**
6 years	-0.24	(0.01)**	-0.27	(0.01)**
7+ years	-0.24	(0.00)**	-0.26	(0.01)**
	Fixed effects			
Firm-prod-yr	Yes		No	
Firm-yr	No		Yes	
Market	Yes		Yes	
N	381452		103297	
rsq	0.70		0.47	
rsq-adj	0.47		0.34	

Notes: Dependent variable is an indicator for exit in the next period. Full set of firm-product-year and market effects included at the firm-product-market-year level. Full set of firm-year and market effects included at the firm-market-year level. Omitted category is market tenure equal to one year. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

differences combine the effect of market tenure with selection. In years 7-13 of successful export spells, prices are on average 7% lower than at the beginning of these spells, and this difference is statistically significant. Again, this result is driven by a combination of market tenure and selection.

The probability of exit is initially decreasing in market tenure and then flattens out after four years in a market. The first column of Table 5 reports results for the baseline estimation of equation (2), while Figure 4 illustrates these findings.

4.2 Market revenue, number of products, and market exit

We now report results at the firm-market level. The fourth and fifth columns of Table 4 report the results from the baseline estimation of equation (3), with log revenue and log number of products as the dependent variable in turn.[23] These results are illustrated in Figures 5 and 6.The evolution of revenue at the firm-market level is qualitatively very similar to the evolution of revenue at the firm-product-market level, though the trajectories are somewhat steeper, reflecting the fact that the number of products per market also evolves with market tenure. Focusing on the longest spells, 70-80% of the growth of revenue at the

[23]Note that the sample of firms included in columns 4 and 5 includes some firms not present in column 1, as the revenue equation at the product level drops the 10% of observations for which quantity is not available.

Figure 4: Exit probability and market tenure: Firm-product-market and firm-market

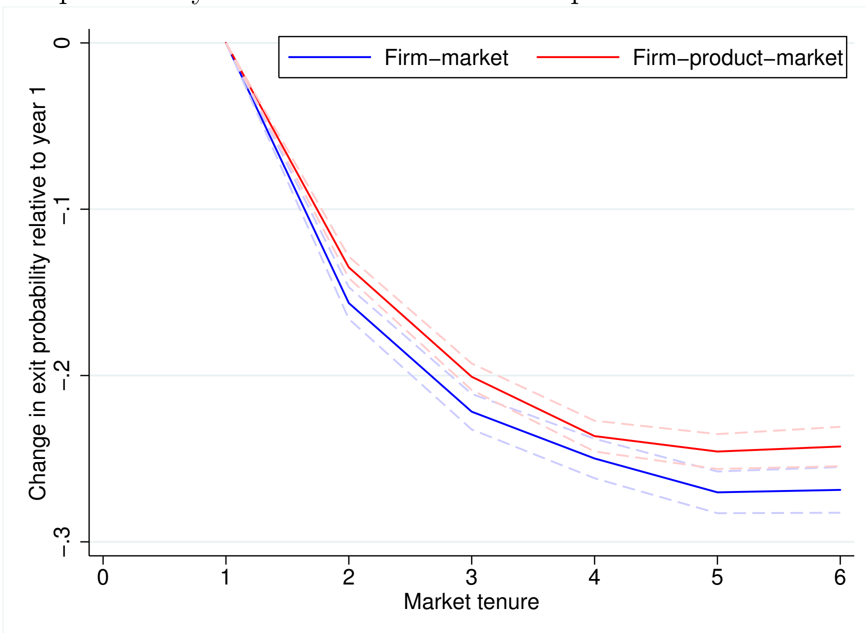

Notes: Figure shows reduction in probability of exit at the firm-market and firm-product-market levels with compared to probability of exit in the first year in a market. Trajectories are conditional on firm-year and market and firm-product-year and market effects, respectively. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

Figure 5: Firm-market revenue by completed spell length and market tenure

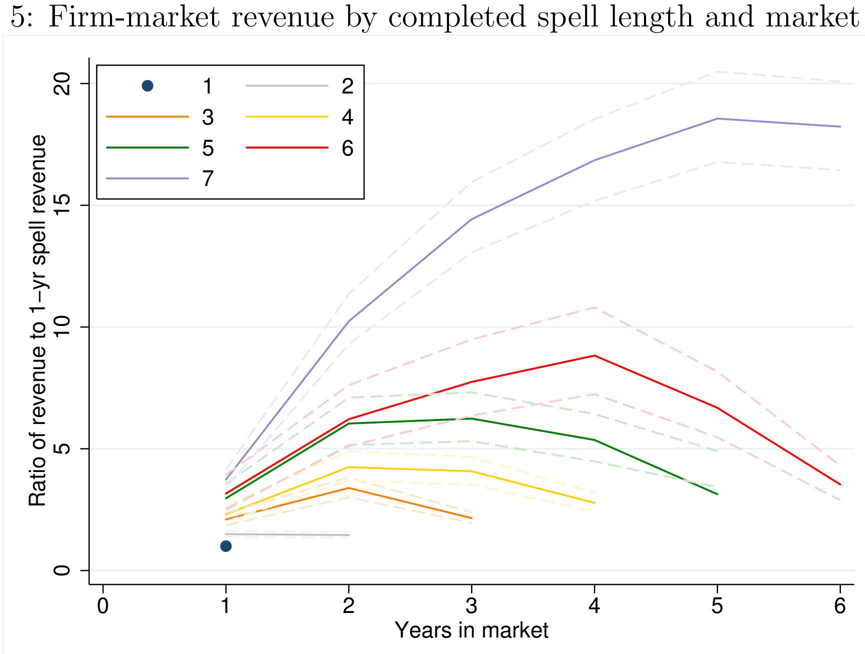

Notes: Figure shows evolution of revenue at the firm-market level with market tenure, allowing trajectories to differ by export spell length. Trajectories are conditional on firm-year and market effects. Source: CSO and authors' calculations.

Figure 6: Number of products per market by completed spell length and market tenure

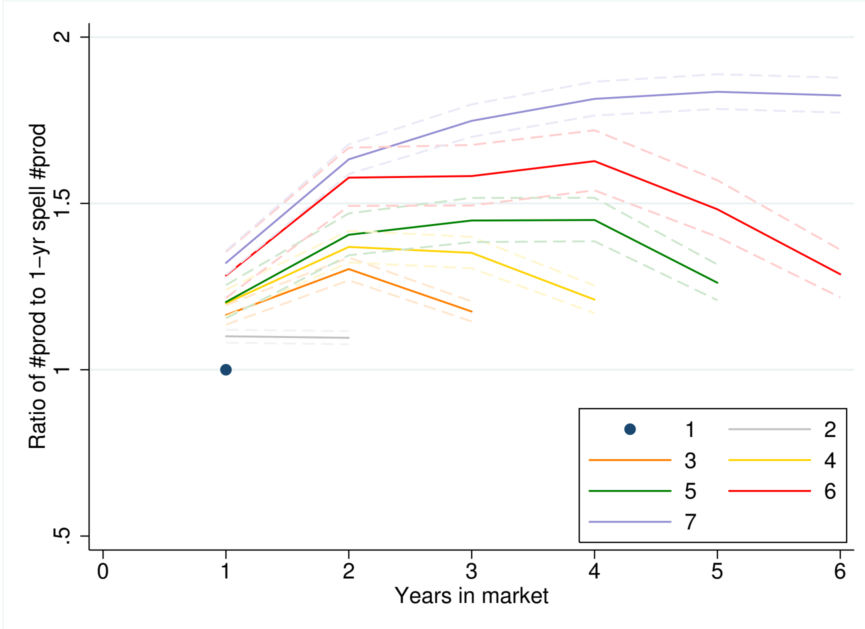

Notes: Figure shows evolution of number of products at the firm-market level with market tenure, allowing trajectories to differ by export spell length. Trajectories are conditional on firm-year and market effects. Source: CSO and authors' calculations.

market level along the growth path is accounted for by within-product growth in revenue, indicating that the within-product margin is of first-order importance in explaining export growth.

The second column of Table 5 reports the results from the baseline estimation of the firm-market exit equation (4). The evolution of exit at the market level is very similar to the evolution of exit at the product-market level, though the probability of exit continues falling until the firm is five years in the market.[24] Figure 4 illustrates the evolution of the probability of exit with market tenure at the market level, with the corresponding evolution at the product-market level for comparison.

4.3 Robustness

We focus on the robustness of the four key results about prices and quantities at the firm-product-market level, and on the finding that at both the product-market level and the market level, the exit hazard initially declines rapidly with market tenure, flattening out after four to five years in a market. We examine both specification robustness and robustness to various cuts of the data. We describe key findings here, and refer to specific tables in the online appendix where the full set of results (including some not described here) is reported.

[24]The probability of exit in the first year is substantially higher at the product-market level than at the market level (62% vs. 45%, unconditional probability of exit from the nonparametric hazard).

4.3.1 Specification robustness

When we include market-year fixed effects (Tables A.6 and A.7 and Figures A.4, A.5 and A.6) or product-market-year fixed effects (Tables A.8 and A.9 and Figures A.6, A.7 and A.8) as appropriate, rather than just market dummies, our results are qualitatively and quantitatively almost unchanged along the key dimensions.

We check what happens when we make use of a subsample of products and firms for which a second measure of quantity (other than tonnes) is reported, constructing quantities and unit values using this alternative measure (Table A.10 and Figures A.9 and A.10). For this subsample, which is 1/6 of the baseline sample, we find that lower initial prices predict that export spells last longer than one year. However, all other results are unchanged.

We vary the level at which spell lengths and market tenure are top-coded, in the range 5 to 8 in our 14-year baseline sample, and in the range 7 to 10 in the 19-year sample that does not require a match between a firm in the customs data and a firm in the Census of Industrial Production. The key results are qualitatively unchanged (a subset of these results are reported in the appendix in Tables A.11 through A.17 and Figures A.11 through A.16). However, when we top-code spell length and market tenure at 10 in the 19-year sample, we cannot reject that quantities in the longest spells reach steady state after about 6 years (rather than 4 years in the baseline analysis). We find that in years 6-9 of the longest spells, prices are statistically different from the beginning of those spells (lower, by 6-8%). We also find that the probability of exit keeps falling after year 5 in a market, and though quantitatively small, this reduction in the probability of exit is statistically significant.

One concern about our baseline specification is that there may be some dimension of idiosyncratic demand that is observable to firms, but not to us, and firms may choose to enter markets that are more attractive along this dimension earlier than less attractive markets. This could bias us towards finding within-spell growth in quantities where there is none. To address this, we estimate two alternative specifications. First we augment the baseline specification with firm-product-*cohort* fixed effects (Tables A.18 and A.19 and Figures A.17, A.18 and A.19). Second, we transform the dependent variable by subtracting the value in the first year of the relevant spell (Table A.20 and Figures A.20 and A.21). The first approach allows us to estimate differences in initial quantities and prices across spells of different length, while the second does not. We find that the growth of quantities in the longest spells is marginally lower under these specifications (in the first 5 years of the longest spells, quantities grow by a factor of 3.3-3.6 rather than 3.9 in the baseline), but otherwise our findings are both qualitatively and quantitatively unchanged.

4.3.2 First and subsequent markets, products and spells

Prompted by the possibility of a role for learning about demand in explaining post-entry export dynamics, several papers use micro data on exports to examine the difference in dynamics between "firsts" (first markets, first products, first spells) and "subsequents" (subsequent markets, subsequent products, subsequent spells). We perform similar cuts of our data.

We allow trajectories to differ across export spells based on the number of markets the firm exported to at the beginning of the spell: a total of three or fewer markets versus four or more markets (Tables A.21, A.22 and A.23 and Figures A.22, A.23 and A.24). Identification of the coefficients of interest comes from within-firm-product-year or within-firm-market-year variation across markets, so restricting to the case where there are few markets reduces the precision of the estimates. However, the key stylized facts are qualitatively replicated for both sets of spells. The one statistically significant difference is that the probability of exit falls more with market tenure in first markets than in subsequent markets.

We allow trajectories to differ between first products and subsequent products, where a product is "first" if on entry, the firm does not export any other products to that market, and is "subsequent" if on entry, the firm already exports at least one product to that market (Tables A.24, A.25 and A.26 and Figures A.25, A.26 and A.27). The key stylized facts are qualitatively replicated for both sets of spells. Quantitatively, the only difference is that in successful spells, the growth of quantities is somewhat steeper for first products than subsequent products.

We allow trajectories to differ between first spells in a firm-product-market and subsequent or reentry spells in the same firm-product-market (Tables A.27, A.28 and A.29 and Figures A.28, A.29 and A.30). Note that spells we classify as first may include some subsequent spells where first spells are censored (i.e., took place before our sample begins). The estimates for subsequent spells are noisy, as there are relatively few of these spells. However, the key stylized facts are qualitatively replicated for both sets of spells. The only statistically significant difference is that the probability of exit in year 1 is lower (by about 7%) in subsequent spells than in first spells.

4.3.3 Firm and product characteristics

In a second set of sample robustness checks, we split the sample by firm and product characteristics. We first estimate separate trajectories for domestic-owned and foreign-owned firms (Tables A.30, A.31, A.32 and A.33 and Figures A.31, A.32 and A.33). Although they account for only 10% of firms in the CIP, more than half of the observations in our baseline sample come from foreign-owned firms, as they are bigger and more export-oriented than domestic-owned firms. The key stylized facts are qualitatively replicated for both groups. The only statistically significant difference is that growth in quantities in the initial years of successful export spells is higher in foreign-owned than domestic-owned firms.

We estimate different trajectories based on firm size (as measured by employment) at the time of firm-product-market or firm-market entry (Tables A.34 through A.42 and Figures A.34 through A.42). The key stylized facts are qualitatively replicated for small and large firms. The only difference is that growth in quantities in the initial years of successful export spells is higher in large than small firms. This does not depend on the threshold for classifying a firm as “large.”

We estimate separate sets of trajectories for different industry groups: consumer food products; consumer nonfood nondurables; consumer durables; intermediates and capital goods (Tables A.43 through A.52 and Figures A.43, A.44 and A.45). This categorization is based on the NACE Revision 1.1 three-digit sector of the firm.[25] Estimates for consumer nonfood nondurables and consumer durables are noisy, as there are relatively few firms in these industries.[26] The key stylized facts are qualitatively replicated for all industry groups.

We use a concordance between the Rauch (1999) classification of goods as homogeneous, reference-priced, or differentiated, and the HS six-digit product classification, to apply the Rauch classification at the product level in our data. This allows us to classify products for 89% of our baseline estimation sample. Of these, about 5% are classified as homogeneous, 16% as reference-priced, and the remainder as differentiated. We then estimate separate sets of trajectories for the three groups of products (Tables A.53 through A.58 and Figures A.46, A.47 and A.48). The key stylized facts are qualitatively replicated for all product types.

4.4 Relation to the literature

A key question is whether our dramatic findings on the absence of price dynamics are due to some special feature of our data, or hold more broadly. Since we do not have access to data for other countries, we cannot test this directly. However there are several kinds of evidence that make us confident that our findings have general applicability.

First, using our data, we can replicate the findings of a large body of literature working with firm and customs micro data for other countries. Summary statistics on the cross-

[25] The assignment of three-digit sectors to industry groups is detailed in the online appendix.

[26] Pharmaceuticals, a key industry for Ireland in terms of export value, though not employment, is categorized as a consumer nonfood nondurable.

sectional dimension of exporting in our data are in line with those for other small open economies (see ISGEP (2008)). Our findings on the post-entry dynamics of revenues and exit are similar to those in the previous literature, (e.g., Eaton et al. (2008), Eaton et al. (2014), Ruhl and Willis (2016)). In Fitzgerald, Yedid-Levi and Haller (2016), we show that there is pricing-to-market in our data, as has been shown for other countries (e.g. Berman et al. (2012) who use French data). In the online appendix to this paper (Table A.63), we show that prices vary with destination market characteristics just as in the literature surveyed in Harrigan, Ma and Shlychkov (2015).

Second, we can replicate the findings of several papers contemporaneous with ours which explicitly address the relationship between quantity, price, and market tenure using data for other countries. Berman et al. (2015) and Piveteau (2016) use customs data for France, while Bastos et al. (2016) use data for Portugal. The empirical specifications used by these authors differ from our baseline, and from each other. Berman et al. and Bastos et al control for marginal cost, but not for selection, and find that prices fall marginally with market tenure, while Piveteau conditions on spell length, but does not control for marginal cost, and finds a positive correlation between prices and market tenure. When we estimate the Berman et al. and Piveteau baseline specifications using our data, we find results similar to theirs (Tables A.64 and A.65). Moreover, in robustness anaysis, Berman et al. and Piveteau estimate specifications that do resemble ours. In these exercises, they find that prices are flat with respect to market tenure, exactly as we find.[27] Unlike us, none of these authors show how initial prices are related to export spell length.

Third, when we use firm-product level prices for our firms from the PRODCOM survey, we can replicate the findings of Foster, Haltiwanger, and Syverson (2008). These authors use data for a select set of US manufacturing industries to show that if one does not condition on survival, older firms charge higher prices. Meanwhile, exiting firms charge lower prices than continuing firms, suggesting that selection can account at least partially for the apparent relationship between prices and age. When we implement a specification that resembles that of Foster et al., prices increase with firm-product age, and that exiting firm-products have lower prices than continuing firm-products (Table A.66).[28]

Finally, our results on the importance of the product extensive margin are quantitatively very consistent with those of Hottman, Redding, and Weinstein (2016) on the contribution

[27]See column (6) of Table 6 in Berman et al. (2015) and especially column (4) of Appendix Table 9 in Piveteau (2016).

[28]This finding does not conflict with our baseline results, as there could be dynamics of cost and quality at the firm level that we condition out using firm-product-year fixed effects in our baseline specification.

of the extensive margin of products to explaining variation in firm size.

Taking stock, we conclude that our results are not driven by any special features of our data, and any differences between our findings and those of the related literature are due to specification differences.

5 A model of post-entry export dynamics

In what follows, we focus on within-product dynamics, as 70-80% of revenue dynamics at the firm-market level are due to the within-product dimension. Supply-side factors (e.g., productivity, costs of adjusting physical capital, capacity constraints, or financial constraints) that affect all markets equally cannot explain the dynamics in quantities and revenues that we document, since our estimates are conditional on firm-product-year fixed effects. Supply-side factors that are market-specific could rationalize quantity dynamics, but have a hard time doing so without generating price dynamics. For example, market-specific product quality that is systematically related to market tenure could generate quantity dynamics. But if the production of quality is costly, this would imply price dynamics which we do not observe. As a result, we turn our attention to the demand side.[29]

The literature has proposed several ways to generate post-entry export dynamics arising from the demand side. We briefly describe two of the most popular mechanisms, and argue that they cannot match the joint behavior of quantities and prices. To do this it will be useful to have some notation. Let firm i's demand in market k at time t be given by[30]

$$Q_t^{ik} = d\left(P_t^{ik}, \varepsilon_t^{ik}, D_t^{ik}\right). \tag{5}$$

Q_t^{ik} is the quantity firm i sells in market k at date t. P_t^{ik} is the price the firm charges to buyers from k. ε_t^{ik} is a shock, idiosyncratic to the firm and the market, which shifts demand conditional on price and is exogenous to the firm. D_t^{ik} is a variable that is idiosyncratic to the firm and the market, and which shifts demand conditional on price, but which depends on actions taken by the firm at t or in previous periods. We refer to it as "customer base."

Ravn et al. (2006), Foster et al. (2016) and Piveteau (2016) assume that customer base

[29] It is possible to construct an alternative rationalization of the facts we document based on a price-taking firm that faces marginal costs of distributing goods that are increasing in the quantity sold. This explanation has a very similar flavor to the one we pursue.

[30] We abstract from factors such as aggregate demand and the aggregate price level in market k that are controlled for by market, market-year, or product-market-year fixed effects in our regressions. We also abstract from the product dimension for the reasons noted above.

D_t^{ik} is increasing in past sales. If this is how customer base is accumulated, entrants have an incentive to distort markups below their long-run level in order to increase future demand. Hence, quantities, markups, and prices are predicted to rise with tenure in a market in successful export spells. This story runs counter to the facts we document, since we do not observe dynamics in prices.

A growing literature models post-entry export dynamics as arising from learning about the process for idiosyncratic demand, following Jovanovic (1982).[31] Informally, this works as follows.[32] Entrants have a prior belief about the distribution of ε_t^{ik}. They can learn about their realization of ε_t^{ik} by selling in market k. Under uncertainty about demand, it matters whether firms maximize expected profits by choosing prices or quantities. Suppose they choose quantities. Ex post, by observing realized prices, they learn about ε_t^{ik}. A high price implies high idiosyncratic demand, whereas a low price implies low idiosyncratic demand. Given updated beliefs, next period firms choose whether or not to participate, and conditional on participation, they adjust quantities to maximize expected profit given the new beliefs. High initial prices induce firms to increase quantities (and hence realized prices fall), whereas low initial prices induce them to reduce quantities (and hence realized prices rise) or, alternatively, to exit. The converse is true when the firm sets prices instead of quantities.[33]

The behavior of revenue is consistent with dynamics being driven by learning about idiosyncratic demand, but the joint behavior of quantities and prices is not. The fact that initial quantities are positively correlated with spell length, whereas initial prices do not help forecast spell length, suggests that firms set prices rather than quantities. But if learning about idiosyncratic demand is the only source of dynamics, and if firms set prices rather than quantities, prices should (weakly) rise and quantities should (weakly) fall with market tenure in successful spells. This is not what we observe.

We propose instead a model of the firm's decision problem which features both customer base and learning, but which is consistent with the behavior of quantities and prices (the behavior of prices is hardwired into the model). We make several simplifying assumptions to keep the model computationally tractable and tailored to the moments of quantities, prices and exit documented in the previous section. These assumptions can easily be relaxed to allow for more richness in applications where speed of computation is not a concern, or richer

[31] See, for example, Albornoz et al. (2012), Berman et al. (2015), Timoshenko (2015a, 2015b) and Arkolakis et al. (2015).

[32] We show this formally in section 3 of the online appendix.

[33] Under its belief, the firm does not anticipate price dynamics; but the econometrician observes realized prices, not firms' expectations of prices.

data moments are available. We now describe the model in detail.

5.1 Model description

We start with a simple model of export entry. As illustrated by Table 2, entry is not perfectly synchronized across markets within the firm, and moreover, there is a good deal of steady state churn in the number of markets a firm participates in. This suggests an idiosyncratic dimension to entry at the firm-market level. To capture this in the simplest possible way, we assume that firm i faces a stochastic sunk cost S_t^{ik} (in terms of some numeraire) of participating in market k. With probability λ, independent across firms, markets, and over time, $S_t^{ik} = S < \infty$, and entry is possible. With probability $1-\lambda$, the sunk cost is infinity, and entry is not possible.[34]

We also assume that the firm faces a stochastic per-period fixed cost of participating in market k given by F_t^{ik}. With probability $1-\omega$, independent across firms, markets, and over time, $F_t^{ik} = F < \infty$, and the firm may choose to participate. With probability ω, the fixed cost is equal to infinity, and the firm must exit (or remain out of the market). These costs stand in for entry and exit that is triggered by macro factors that we do not model, as well as entry and exit that is not driven by macro factors, but is idiosyncratic to the firm and the market. We make use of the following notation: $X_t^{ik} \in \{0,1\}$ is an indicator for participation in market k by firm i at date t.

Firm i is characterized by marginal cost C^i, assumed the same in all markets.

Demand for firm i in market k depends on its own price P_t^{ik}, on customer base D_t^{ik}, and on idiosyncratic demand ε_t^{ik}. Demand takes the following form:[35]

$$Q_t^{ik} = \left(P_t^{ik}\right)^{-\theta} \left(D_t^{ik}\right)^{\alpha} \exp\left(\varepsilon_t^{ik}\right). \tag{6}$$

Note that customer base and idiosyncratic demand do not affect the price elasticity of demand (θ) in this formulation. If $\alpha \in (0,1)$, demand is increasing in customer base, but at a diminishing rate. In this case there is a finite positive steady state for D_t^{ik}.

Customer base accumulates as follows:

$$D_t^{ik} = (1-\delta) X_{t-1}^{ik} D_{t-1}^{ik} + A_t^{ik}, \tag{7}$$

[34]The choice of two-point distributions for sunk and fixed costs is motivated by the moments we use to identify these distributions.

[35]We abstract from factors such as aggregate demand and the aggregate price level in market k, as these are cleaned out of our target moments.

where A_t^{ik} is investment in customer base, and δ is the depreciation rate of customer base.[36] Investment in customer base is assumed to be subject to both convex costs of adjustment governed by ϕ, and irreversibility:

$$c\left(D_t^{ik}, A_t^{ik}\right) = \begin{cases} A_t^{ik} + \phi\left(\frac{A_t^{ik}}{D_t^{ik}} - \delta\right)^2 D_t^{ik} & \text{if } A_t^{ik} > 0 \\ \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

This portion of the model resembles Arkolakis (2010), with the modification that customer base can be accumulated. The irreversibility assumption is very natural in the context of investment in an intangible such as customer base.

Idiosyncratic demand has two components, permanent and transitory: $\varepsilon_t^{ik} = \nu^{ik} + \eta_t^{ik}$, with $\nu^{ik} \sim N\left(0, \sigma_\nu^2\right)$, $\eta_t^{ik} = \rho\eta_{t-1}^{ik} + \zeta_t^{ik}$, and $\zeta_t^{ik} \sim N\left(0, \sigma_\eta^2\right)$. The structure of information is as follows. When making choices about participation, investment, and prices, both potential entrants and incumbents observe C^i, F_t^{ik}, and S_t^{ik}.[37] Potential entrants do not observe ε_t^{ik} or its components. As a result, they use the unconditional distribution of $\nu^{ik} + \eta_t^{ik}$ to form expectations about current and future ε_t^{ik}. We model learning as follows. Incumbents may be "uninformed" or "informed." Uninformed incumbents use the unconditional distribution to form expectations, just like potential entrants. With probability γ per period, uninformed participants become informed. This happens at the end of a period. Incumbents that are informed entering period t observe ν^{ik} and η_{t-1}^{ik} and hence use the distribution of $\nu^{ik} + \eta_t^{ik}$ *conditional on* ν^{ik} and η_{t-1}^{ik} to form expectations about current and future demand. As long as an informed incumbent remains in the market, it remains informed. As soon as an incumbent exits, it loses its current draws of ν^{ik} and η_{t-1}^{ik} and becomes an uninformed potential entrant.[38] Let N_{t-1}^{ik} be an indicator variable that takes the value 0 if the firm is uninformed in market k entering period t, and 1 if it is informed. The firm's information set I_t^{ik} is therefore a state variable of its problem, which evolves as follows:

$$I_t^{ik} = \begin{cases} \left\{\nu^{ik}, \eta_{t-1}^{ik}\right\} & \text{if } \left\{X_{t-1}^{ik} = 1, N_{t-1}^{ik} = 1\right\} \\ \\ \emptyset & \text{if } \left\{X_{t-1}^{ik} = 0\right\} \text{ or } \left\{X_{t-1}^{ik} = 1, N_{t-1}^{ik} = 0\right\}. \end{cases} \tag{9}$$

[36] For computational tractability, we assume full depreciation of customer base on exit.
[37] S_t^{ik} is irrelevant to the choices of incumbents.
[38] This assumption is made for computational tractability.

This way of modeling information acquisition is both tractable and flexible, in that it allows the average speed of learning to be fast (γ close to 1) or slow (γ close to 0), irrespective of the parameters of the idiosyncratic demand process.[39]

Since there is uncertainty about demand, it matters whether firms set prices or quantities. We assume firms set prices. Because demand is CES and there are no strategic interactions with other firms, the optimal price is equal to the statically optimal markup over marginal cost ($\frac{\theta}{\theta-1}$), irrespective of the firm's participation history, information set, or customer base. In this sense, the model hardwires in a flat path of prices with respect to market tenure.

Assuming that it discounts the future at rate β, we can then write the firm's intertemporal optimization problem as follows:

$$V\left(D_{t-1}^{ik}, X_{t-1}^{ik}, I_t^{ik}, F_t^{ik}, S_t^{ik}, C^i\right) =$$

$$\max_{\substack{X_t^{ik} \in \{0,1\} \\ A_t^{ik}}} \left\{ \begin{array}{c} X_t^{ik} \frac{(\theta-1)^{\theta-1}}{\theta^\theta} \left(C^i\right)^{1-\theta} \left(D_t^{ik}\right)^\alpha \mathbb{E}\left(\exp\left(\nu^{ik} + \eta_t^{ik}\right) | I_t^{ik}\right) \\ \\ -X_t^{ik}\left(F_t^{ik} + \left(1 - X_{t-1}^{ik}\right) S_t^{ik}\right) - c\left(D_t^{ik}, A_t^{ik}\right) \\ \\ +\beta \mathbb{E}\left(V\left(D_t^{ik}, X_t^{ik}, I_{t+1}^{ik}, F_{t+1}^{ik}, S_{t+1}^{ik}, C^i\right) | I_t^{ik}\right) \end{array} \right\} \tag{10}$$

subject to (7), the accumulation equation for D, (8), the cost of investment, and (9), the updating of information, which includes the process for N_t^{ik} as a function of lagged participation.

6 Model estimation and results

6.1 Estimation

We use simulated method of moments to estimate the model. Given values for parameters β, α, δ, ϕ, θ, σ_ν^2, ρ, σ_η^2, F, ω, S, λ, and γ and assumptions about C^i, we discretize both exogenous and endogenous states[40] and use value function iteration to solve for the optimal policies for participation and investment. Using the model parameters and the corresponding optimal policies, we then simulate post-entry trajectories for 50,000 "firm-markets." To account for

[39] In a robustness check, we examine the more conventional case of Bayesian learning where firms learn about idiosyncratic demand from observing ε_t^{ik}, and where the speed of learning is linked to the parameters of the idiosyncratic demand process.

[40] We use three states each for the permanent and transitory idiosyncratic demand shocks (ν and η). The number of endogenous states depends on parameter values.

the fact that there are part-year effects in the data, the length of a period in our model is 6 months, but we aggregate up to an annual frequency to construct the equivalents of the moments we estimate in Section 4. The goal of our estimation is to choose the vector of parameters that best matches these moments.

We match moments of four types. The first three sets of moments are based on our estimates conditional on marginal cost and market effects: the ratios of initial quantities across spells of different length; the evolution of quantities with market tenure within spells of different length; and the evolution of exit probabilities with market tenure, all of these conditional on costs. The last moment is the average exit rate in the first year in a market. We match the average exit rate across all export spells in our data for which entry is not censored. We do not target moments related to prices, as they are matched automatically in our baseline model. The full set of 32 targeted moments is reported in Table 6. We also match exactly a target entry rate of 1% at the firm-market level. This is the rate of entry by nonparticipating firms at the market level, averaged across the 56 export markets which account for 99% of exports in our data.

We first preset some parameters not identified by our target moments. Since we calibrate to a 6-month period, we set $\beta = 1.05^{-0.5}$. In our baseline model, θ is not well identified by our target moments. We pick $\theta = 2$ as a baseline value for θ. This is consistent with a markup over *marginal* cost (which does not include fixed costs of production or costs related to marketing and advertising) of 100%.[41] The key data moments are estimated conditional on costs, so we normalize $C^i = 1$. We set $S = 0$, as this parameter is redundant given that there is no cost heterogeneity. The export entry rate in our model is equal to $\lambda(1-\omega)$. We set $\lambda = 0.01/(1-\omega)$.[42]

This leaves us with nine parameters, $\left\{\alpha, \delta, \phi, \sigma_\nu^2, \rho, \sigma_\eta^2, F, \omega, \gamma\right\}$. We choose these parameters to minimize the criterion function $m'Vm$, where m is the difference between the data moments and the equivalent population moments in the model, and V is a diagonal matrix, with the inverse of the standard deviation of the estimates of the data moments on the diagonal (we do not include the entry rate in this matrix, as we hit this target by construction). We use a combination of a particle swarm algorithm and the simplex method to optimize.

6.2 Baseline results

[41] The online appendix describes results based on alternative values for θ in Tables A.70 and A.71.

[42] The average entry rate for non-participants across the 56 export markets that account for 99% of exports is 1%.

Figure 7: Model fit: Quantities

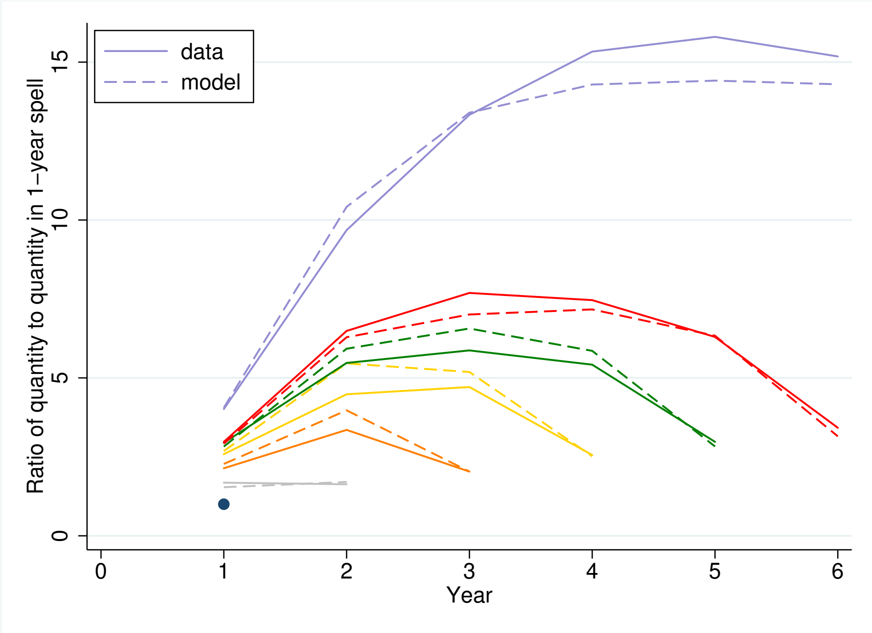

Notes: Figure shows data on evolution of quantities at the firm-product-market level with tenure by spell length from 2, and corresponding quantity trajectories for the structural model. All quantities are expressed relative to the quantity in a 1-year spell. Source: CSO and authors' calculations.

Figure 8: Model fit: Exit

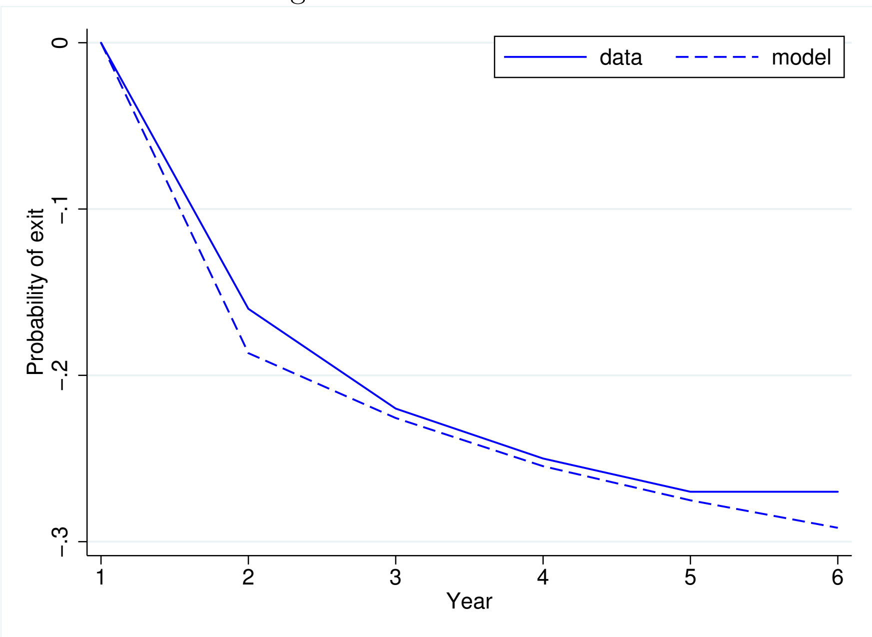

Notes: Figure shows data on reduction in probability of exit at the firm-market level relative to probability of exit in the first year in a market, and corresponding evolution for the structural model. Figure does not illustrate exit rate in year 1. Source: CSO and authors' calculations.

Table 6 reports the data moments in the first column and the corresponding fitted values of the moments from the baseline model in the second column. Figures 7 and 8 illustrate the fit of the model by graphing target and fitted moments for quantities and exit. The estimated model matches all of the key facts in the data. It generates dispersion in initial quantities that is positively correlated with spell length and of the right order of magnitude, quantities that increase with market tenure in successful spells as in the data, and an exit hazard that declines substantially between the first and second year in a market and continues to fall until about five years in the market, closely matching the data. The 1-year exit rate is somewhat lower than in the data, at 42% compared to 46%.

Table 7 reports the estimated parameters and an overall measure of fit based on the criterion function.[43] Our estimate of α, the elasticity of sales with respect to customer base, is 0.50. This indicates that there is a significant role for customer base in explaining post-entry export dynamics. This customer base depreciates at a relatively rapid rate: our estimate of δ, 0.46, implies an annual rate of depreciation of 70%, which is in line with the findings reported in Bagwell (2007) of an empirical literature on the annual depreciation rates of brand advertising.[44] Our estimated value of ϕ is nonzero, consistent with adjustment costs à la Arkolakis (2010). In order to provide some intuition for the magnitude of ϕ, we calculate the weighted average expenditure on marketing and advertising (i.e., $A_t^{ik} + c\left(D_t^{ik}, A_t^{ik}\right)$) as a share of revenue. For export spells lasting at least seven years, 33% of revenues are devoted to marketing and advertising in the initial year in a market, while this declines to just under 23% of revenues by year 6 (see Table A.68 in the online appendix). We do not have data on marketing and advertising expenditures for firms in our data. Nevertheless, these shares seem somewhat large.[45] However this is a function of our baseline choice of θ. Higher values of θ yield lower shares of marketing and advertising investment in revenue, without affecting the fit of other moments.

We find that there is a nontrivial variance of both permanent and mean-reverting idiosyncratic demand shocks. Idiosyncratic demand shocks are persistent, but not strikingly so. Meanwhile, our estimate of γ, the probability of an uniformed incumbent learning its true ν^{ik} and η_{t-1}^{ik}, is equal to 0.57 on a 6-month basis. This implies that by the beginning

[43] Standard errors are constructed using the method suggested by Gourieroux, Monfort and Renault (1993), as described in the online appendix. They are sensitive to the calculation of the numerical derivatives of the moments with respect to individual parameters. We believe that the exercises where we estimate restricted versions of the model are more informative about the role played by individual parameters.

[44] Eaton et al. (2014) report very high rates of attrition of exporter-importer matches, which is consistent with our finding of a high depreciation rate for customer base.

[45] Gourio and Rudanko (2014) provide evidence from Compustat for manufacturing and nonmanufacturing firms, suggesting higher shares, in the range 17-27%.

Table 6: Data and model moments: Baseline and alternative structural models

	Data		Model							
		s.e.	baseline	no PY	$\alpha = 0$	no AC	$\gamma = 1$	$\rho = 0$	fullinf	bayes
$\ln\left(Q_1^2/Q_1^1\right)$	0.52	(0.02)	0.43	0.95	0.35	0.63	0.73	0.42	0.73	0.73
$\ln\left(Q_1^3/Q_1^1\right)$	0.76	(0.04)	0.82	1.32	0.65	0.83	0.80	0.82	0.83	0.92
$\ln\left(Q_1^4/Q_1^1\right)$	0.95	(0.05)	0.99	1.44	1.04	0.89	0.86	1.08	0.91	0.93
$\ln\left(Q_1^5/Q_1^1\right)$	1.08	(0.07)	1.04	1.48	1.30	0.93	0.90	1.23	1.04	1.01
$\ln\left(Q_1^6/Q_1^1\right)$	1.09	(0.08)	1.07	1.52	1.14	0.96	0.96	1.14	1.06	1.03
$\ln\left(Q_1^7/Q_1^1\right)$	1.39	(0.05)	1.40	1.86	1.25	1.39	1.30	1.20	1.10	1.57
$\ln\left(Q_2^2/Q_1^2\right)$	-0.03	(0.03)	0.10	-0.44	-0.01	-0.07	-0.09	0.04	-0.23	-0.09
$\ln\left(Q_2^3/Q_1^3\right)$	0.45	(0.05)	0.56	0.09	1.00	0.54	0.53	0.79	0.56	0.48
$\ln\left(Q_3^3/Q_1^3\right)$	-0.05	(0.05)	-0.11	-0.60	0.05	-0.05	-0.02	0.19	-0.10	-0.09
$\ln\left(Q_2^4/Q_1^4\right)$	0.55	(0.06)	0.71	0.27	0.87	0.62	0.60	0.75	0.67	0.53
$\ln\left(Q_3^4/Q_1^4\right)$	0.60	(0.06)	0.66	0.21	0.91	0.62	0.63	0.86	0.78	0.50
$\ln\left(Q_4^4/Q_1^4\right)$	-0.01	(0.07)	-0.06	-0.53	0.05	0.04	0.08	0.24	0.08	-0.07
$\ln\left(Q_2^5/Q_1^5\right)$	0.62	(0.09)	0.74	0.30	0.73	0.68	0.67	0.70	0.63	0.59
$\ln\left(Q_3^5/Q_1^5\right)$	0.69	(0.09)	0.84	0.37	0.82	0.70	0.71	0.80	0.82	0.66
$\ln\left(Q_4^5/Q_1^5\right)$	0.61	(0.09)	0.73	0.25	0.72	0.68	0.71	0.88	0.82	0.63
$\ln\left(Q_5^5/Q_1^5\right)$	0.01	(0.09)	0.00	-0.46	-0.04	0.11	0.18	0.24	0.12	0.09
$\ln\left(Q_2^6/Q_2^6\right)$	0.78	(0.11)	0.77	0.33	0.94	0.75	0.76	0.79	0.83	0.62
$\ln\left(Q_3^6/Q_1^6\right)$	0.95	(0.11)	0.88	0.41	0.95	0.76	0.80	0.89	0.89	0.76
$\ln\left(Q_4^6/Q_1^6\right)$	0.92	(0.11)	0.90	0.41	0.98	0.75	0.81	1.01	0.95	0.78
$\ln\left(Q_5^6/Q_1^6\right)$	0.75	(0.11)	0.78	0.30	0.87	0.74	0.81	1.03	0.97	0.72
$\ln\left(Q_6^6/Q_1^6\right)$	0.14	(0.11)	0.08	-0.39	0.15	0.18	0.28	0.33	0.24	0.14
$\ln\left(Q_2^7/Q_1^7\right)$	0.88	(0.06)	0.94	0.55	0.84	1.19	1.12	0.75	0.83	0.80
$\ln\left(Q_3^7/Q_1^7\right)$	1.20	(0.06)	1.19	0.73	0.87	1.23	1.23	0.86	1.08	1.12
$\ln\left(Q_4^7/Q_1^7\right)$	1.34	(0.06)	1.26	0.77	0.91	1.23	1.25	0.95	1.15	1.33
$\ln\left(Q_5^7/Q_1^7\right)$	1.37	(0.06)	1.27	0.78	0.90	1.24	1.25	0.99	1.19	1.45
$\ln\left(Q_6^7/Q_1^7\right)$	1.33	(0.07)	1.26	0.77	0.85	1.23	1.25	1.03	1.20	1.54
$exit_1$	0.46	(0.004)	0.42	0.46	0.54	0.41	0.45	0.43	0.30	0.28
$exit_2 - exit_1$	-0.16	(0.005)	-0.19	-0.23	-0.01	-0.16	-0.20	-0.16	-0.16	-0.27
$exit_3 - exit_1$	-0.22	(0.005)	-0.23	-0.28	-0.19	-0.21	-0.23	-0.23	-0.18	-0.27
$exit_4 - exit_1$	-0.25	(0.006)	-0.25	-0.30	-0.29	-0.23	-0.26	-0.25	-0.19	-0.28
$exit_5 - exit_1$	-0.27	(0.006)	-0.28	-0.32	-0.31	-0.25	-0.28	-0.26	-0.19	-0.28
$exit_6 - exit_1$	-0.27	(0.007)	-0.29	-0.33	-0.31	-0.27	-0.31	-0.26	-0.20	-0.28

Notes: Data quantity moments are based on second column of Table 4. Exit moments are based on second column of Table 5. $exit_1$ is the average 1-year exit rate across all spells in the data for which entry is not censored. Standard error for $exit_1$ is based on assuming the indicator for 1-year exit is binomially distributed. The full information model additionally targets the entry rate, which set equal to 1% by construction in all the other models. The target is 1% (s.e. 0.0001) and the fitted value for the entry rate for this model is 0.25%. Parameter estimates are reported in Table 7.

Table 7: Parameters and fit: Baseline and alternative structural models

	Parameter											
Model	α	δ	ϕ	γ	ρ	σ_ν	σ_η	$\frac{F}{E(R_1)}$	ω	λ	θ	Fit
baseline	0.50	0.46	3.03	0.58	0.40	0.52	0.34	0.31	0.03	n.a.	n.a.	3.71
s.e.	(0.23)	(0.15)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	†	(0.00)			
$\alpha = 0$	n.a	n.a.	n.a.	0.49	0.00	0.45	2.56	0.52	0.11	n.a.	n.a.	32.44
no AC	0.76	0.80	n.a.	0.87	0.44	0.37	0.19	0.10	0.03	n.a.	n.a.	5.46
$\gamma = 1$	0.70	0.86	17.63	n.a.	0.31	0.34	0.17	0.12	0.04	n.a.	n.a.	6.62
$\rho = 0$	0.29	0.10	10.23	0.42	n.a.	0.68	1.89	0.40	0.08	n.a.	n.a.	19.86
fullinf	0.29	0.47	13.27	n.a.	0.07	1.95	2.05	0.00	0.05	0.01	n.a.	18.45‡
setq	n.a.	n.a.	n.a.	0.96	0.93	0.98	2.00	4.74	0.01	n.a.	36	24.33‡
bayes	0.51	0.29	35.82	n.a.	0.83	0.66	0.16	1.72	0.00	n.a	n.a.	16.99

Notes: Standard errors on the baseline parameter estimates are calculated following Gourieroux, Monfort and Renault (1993) as described in the online appendix. † The baseline estimate of F is 0.0131 and the standard error is 0.0105. "Fit" is the value of the criterion function, $m'Vm$, where m is the difference between data moments and moments of the model conditional on the parameter vector, and V is a diagonal matrix with the vector of inverses of the standard errors of the data moments on the diagonal. ‡ For the full information model, the criterion function minimized in the estimation includes the entry rate, but for comparability the fit value reported here includes only the baseline moments. For the quantity setting model, the criterion function minimized in the estimation includes price moments, but for comparability the fit value reported here includes only the baseline moments.

of their fourth year, over 99% of incumbents are informed (see Table A.69 in the online appendix).

We find that there are nontrivial fixed costs of export participation. For the ultimately successful spells (i.e those lasting at least 7 years), these amount to to 15% of revenue in the initial year in a market and decline to 6% of revenue by year 6.

Finally, we estimate a modest probability of the firm-market experiencing a "death shock," since $\omega = 0.03$ is the probability of exogenous exit on a 6-month basis. Most of the exit out to seven years in a market is endogenous.

6.3 Restricted versions of the baseline model

We now examine the key features of the model. First we address the role of part-year effects. Then we examine the role of other features by eliminating them one by one, and re-estimating the model each time. This is helpful in understanding the mapping between moments and parameters.

To examine the role of part-year effects in matching the dynamics in the data we keep all parameters fixed. We calculate the targeted moments assuming that all entrants enter at the beginning of a year, rather than half entering in the first 6-month period, and half entering in the second 6-month period. The implied moments are reported in Table 6. Figures A.54 and A55 in the online appendix illustrate the results. Part-year effects are key to matching the relationship between initial quantities in 1-year export spells and initial quantities in longer export spells, but not to matching the relationships between initial quantities in spells longer than 1 year. They also play a role in matching within-spell growth rates, not just in the first and last years in a spell, but also to some extent in the second year of a spell, though they are not responsible for all of within-spell quantity growth. Finally, they play a role in matching the decline in the exit hazard between year 1 and year 2 in a market, but they are not key to the subsequent behavior of exit.

We shut down customer capital by setting $\alpha = 0$. When we do this, the parameters δ and ϕ become redundant. Now the only intertemporal choice faced by the firm is whether or not to participate.[46] Table 6 reports the fitted moments and Table 7 reports the estimated parameters of this model, while figures A.56 and A.57 in the online appendix illustrate the results. This model can generate an increasing relationship between initial quantities and spell length, but it is incapable of generating within-spell quantity growth beyond that due to part-year effects and exogenous idioysncratic demand. In order to try to match within-spell quantity dynamics, the estimated standard deviation of the mean-reverting component of idiosyncratic demand is very large, which implies that the fit of the exit trajectory is very poor.

We shut down all adjustment costs of investment by setting $\phi = 0$ and allowing for full reversibility of investment expenditures (the firm can "consume" inherited D). This means that learning and part-year effects are the only sources of dynamics in the model. Table 6 reports the fitted moments and Table 7 reports the estimated parameters of this model, while Figures A.58 and A.59 in the online appendix illustrate the results. Somewhat surprisingly, learning is faster in this model than in the baseline. This model can generate an increasing relationship between initial quantities and spell length, and a reduction in exit hazard with market tenure. Relative to the baseline model, it has difficulty matching within-spell quantity growth: almost all of the dynamics of within-spell quantities happen within the first year. This is not a foregone conclusion - with slower learning, this model could generate dynamics over a longer period, but it appears that slower learning distorts other moments more than it improves the fit of within-spell quantity growth.

We partially shut down learning by setting $\gamma = 1$, which implies that at the end of the first period, all participants learn their permanent and (lagged) transitory idiosyncratic demand. The only source of dynamics in quantities after period 2 is due to adjustment costs and residual dynamics from part-year effects. Table 6 reports the fitted moments and Table 7 reports the estimated parameters of this model, while Figures A.60 and A.61 in the

[46] This model resembles that of Das, Roberts, and Tybout (2007), but with the addition of learning about idiosyncratic demand.

online appendix illustrate the results. Unsurprisingly, the value of ϕ increases, as costs of adjustment substitute for learning in generating post-entry dynamics. The fit of this model is in fact remarkably similar to the learning-only model, though the fit for quantity growth is marginally better, while fit of initial quantities and the exit is somewhat worse.

We also estimate the model imposing $\rho = 0$ (i.e., zero autocorrelation of the transitory component of the exogenous idiosyncratic demand shock). Table 6 reports the fitted moments and Table 7 reports the estimated parameters, while Figures A.62 and A.63 in the online appendix illustrate the results. This model is unable to generate different quantity dynamics in export spells of different lengths (apart from part-year effects in the year of exit). Although our point estimate of ρ may seem low, it is crucial to allow for autocorrelation of transitory idiosyncratic demand in order to match the different behavior of quantities in short and long export spells.[47]

6.4 Non-nested alternative models

We also examine what happens when we modify the model in ways that are not nested in our baseline.

We estimate a model where there is no customer base, firms set quantities rather than prices, and this alone combined with learning about idiosyncratic demand generates dynamics in quantities.[48] In this case, we target price moments in addition to quantity moments (since prices are not flat by construction) and we also estimate θ, the price elasticity of demand, since in this model it is well identified by the combination of quantity and price moments. Table 7 reports the estimated parameters of this model, while Table A.67 and Figures A.64, A.65 and A.66 in the online appendix report the fit of the various moments. In order to match the fact that there are substantial dynamics of quantities, but no dynamics of prices, the model requires $\theta = 36$. Notwithstanding this very high value, the estimated model implies that initial prices are more than 10% higher in spells that last at least 3 years than in spells that last only 1 year, in direct contradiction to the data. It also implies that falling prices signal exit, in contrast to the data. Finally, the model provides a very poor fit to the exit trajectory, as learning is estimated to be very fast.

Additionally, we estimate a model without any learning, where firms observe both components of their idiosyncratic demand (ν and η_t) before taking any decisions. Table 6 reports

[47] Papers that generate dynamics through learning about demand, such as Berman et al. (2015) and Arkolakis et al. (2015), typically assume that the transitory component of idiosyncratic demand is i.i.d.

[48] This model is similar to that in Berman et al. (2015) and Arkolakis et al. (2015), except that we allow idiosyncratic demand to be autocorrelated ($\rho > 0$) and model learning as stochastic rather than Bayesian.

the fitted moments and Table 7 reports the estimated parameters, while Figures A.67 and A.68 in the online appendix illustrate the results. The full information model has difficulty matching 1-year exit rates, and the exit trajectory more generally. It requires very high standard deviations of both components of idiosyncratic shocks in order to generate a positive relationship between initial quantities and survival.

Finally, we modify our baseline model by assuming that firms are Bayesian, and use the Kalman filter to update their beliefs about the two components of idiosyncratic demand based on realizations of quantities. This model is considerably more computationally demanding to solve and estimate than the other models we consider.[49] Table 6 reports the fitted moments and Table 7 reports the estimated parameters, while Figures A.69 and A.70 in the online appendix illustrate the results. This model overpredicts initial quantities and growth in the longest spells, and it fails strikingly to match the behavior of exit, generating too little exit in the first year, and being unable to generate dynamics in exit beyond the second year. This is because the parameters of the idiosyncratic demand process necessary to match the behavior of quantities imply very rapid learning: after participating one period in the market, firms assign 99% probability to their true ν. Separating the speed of learning from the parameters of the idiosyncratic demand process is key to the success of the baseline model in matching the data.

6.5 Summary

To sum up, our baseline model provides a good fit to all the moments of the data. Accumulation of customer base appears to play a key role in explaining the post-entry dynamics of export quantities and export prices. Slow learning about idiosyncratic demand is key to matching the behavior of exit. While qualitatively, either costs of adjusting customer base or slow learning could explain the evolution of quantities and exit, quantitatively both of these mechanisms are necessary.

7 Implications for adjustment to shocks

An empirical literature finds that long-run elasticities of exports with respect to prices are bigger than short run elasticities.[50] A quantitative literature shows that sunk costs of export

[49] Restricting the SMM algorithm for our estimates of the Bayesian model to 25,000 rather than 50,000 independent draws of "firms," this model still takes an order of magnitude longer to estimate than the baseline model.

[50] See e.g. Hooper, Johnson and Marquez (2000) and Gallaway, McDaniel and Rivera (2003).

participation can help deliver this difference.[51] However as argued by Ruhl and Willis (2016), the existence of many export spells which start small and swiftly exit is inconsistent with substantial sunk costs of export entry. In this section, we simulate our model, showing that it generates sluggish adjustment in response to shocks even in the absence of entry responses and general equilibrium feedback through marginal cost, the channels through which sunk cost models generate sluggish adjustment. In addition, we show that it can generate responses to news shocks. For illustrative purposes, we focus on shocks to ad valorem tariffs, but the implications are not limited to this type of shock.

All of our simulations take the same form. We fix all parameters at our baseline estimates. This includes setting $C_t^i = 1$ (constant costs, the same for all firms), setting λ such that the rate of entry is fixed at 1%, and setting θ, the price elasticity of demand, equal to 2. We solve for the policy functions of the firm under two different tariff environments. We simulate the behavior of 1,000,000 "firms" for 100 years (i.e. 200 periods) using the initial policy function in the relevant experiment. This gives us an initial distribution of market participants in terms of customer base, information and idiosyncratic demand in year 0. In the first half of year 1, the particular experiment is initiated, and we simulate responses for 20 years (i.e. 40 periods). We also use the same draws for idiosyncratic shocks to simulate a counterfactual world where tariffs do not change. We then sum exports across all firms in order to illustrate the implied dynamics of aggregates in both experiment and counterfactual.

7.1 Short vs long run responses to shocks

We first examine the timing of responses to shocks. Demand for each firm is given by:

$$Q_t^{ik} = \left(\tau_t^k P_t^{ik}\right)^{-\theta} \left(D_t^{ik}\right)^{\alpha} \exp\left(\varepsilon_t^{ik}\right).$$

We assume first that tariffs are equal to 5% ($\tau_H = 1.05$) through year 0. At the begining of year 1, it is announced, unexpectedly, that in that period and from then on, tariffs will fall to 0 ($\tau_L = 1$). We simulate responses of aggregate exports under this experiment, as well as under the counterfactual where tariffs remain (and are expected to remain) equal to 5% throughout. It is convenient to report the results in terms of the elasticity of aggregate exports with respect to the shock at different time horizons. We calculate this elasticity as follows. Let X_t denote aggregate exports at year t in our experiment, and let X_t' denote aggregate exports at year t in the counterfactual. Similarly, let τ_t denote the tariff at year

[51] See e.g. Alessandria and Choi (2014).

Figure 9: Simulation results: Short vs long run elasticities with respect to tariff shocks

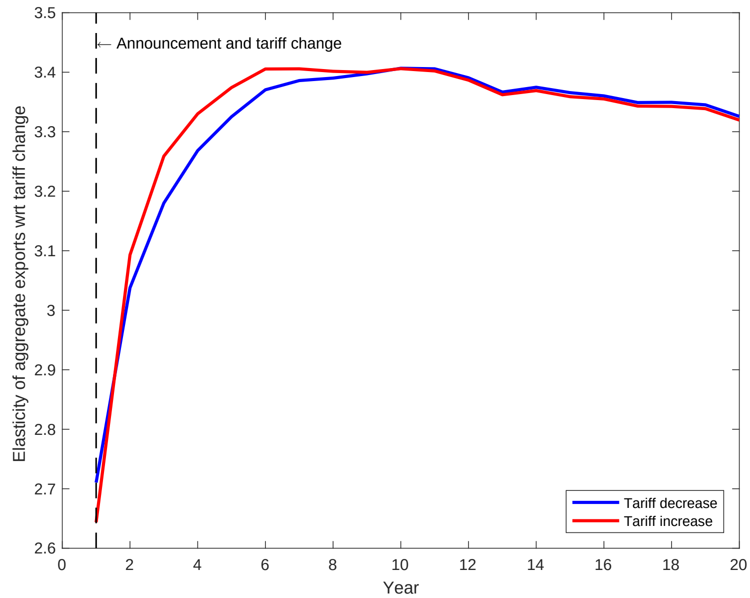

Notes: Figure shows simulated elasticity of aggregate exports with respect to an unexpected 5% change in tariffs at the beginning of year 1, based on 1 million "firms" and baseline parameter estimates. Costs, foreign demand and entry are held fixed. Source: Authors' calculations.

t in our experiment, and let τ_t' denote the tariff at year t in the counterfactual. Then we calculate the elasticity using:

$$\sigma_t = \frac{\ln X_t - \ln X_t'}{\ln \tau_t - \ln \tau_t'}$$

The time series of this elasticity is reported in Figure 9. On impact, the elasticity is 2.7 (note that this is greater than the price elasticity of demand which is given by $\theta = 2$), and rises to 3.4 after 10 years.

In a second version of this exercise, we assume that tariffs are equal to 0 through year 0, and that the unexpected announcement at the beginning of year 1 is that tariffs have risen to 5%. The elasticity based on this second exercise and the corresponding counterfactual where tariffs remain constant at 0 throughout is also reported in Figure 9. Adjustment to this negative shock also takes time. There is a mild degree of asymmetry in responses to positive and negative shocks. On impact, the response to the negative shock is smaller (elasticity is 2.65 rather than 2.7), but after that, adjustment is faster (elasticity is 3.4 by 6 years after the announcement rather than 10 years after). Across both exercises, on average the long-run elasticity (after 10 years) is 1.3 times larger than the short-run elasticity (after 1 year). This compares to a ratio of long to short run elasticities in Galloway et al. of 2.

Figure 10: Simulation results: Elasticities with respect to news shocks about tariffs

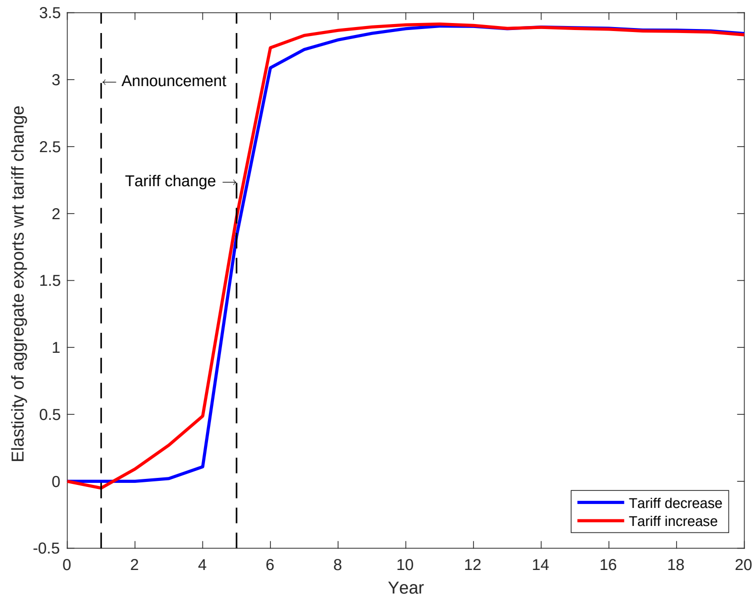

Notes: Figure shows simulated elasticity of aggregate exports with respect to an unexpected announcement at the beginning of year 1 that tariffs will change by 5% at the beginning of year 5, based on 1 million "firms" and baseline parameter estimates.. Costs, foreign demand and entry are held fixed. Source: Authors' calculations.

7.2 Responses to news shocks

In our second exercise, we examine how exports respond to news about future tariffs. We first look at good news. We assume that tariffs are equal to 5% through year 0. At the beginning of year 1 it is announced unexpectedly that from the start of year 5, tariffs will fall to 0. We simulate responses of aggregate exports under this experiment, as well as under the counterfactual where tariffs remain and are expected to remain at 5% throughout. We then calculate the elasticity of aggregate exports with respect to the tariff change at different time horizons, using the future tariff change for years 1 through 4, as actual tariffs remain unchanged during this period:

$$\sigma_t = \frac{\ln X_t - \ln X_t'}{\ln (1.05) - \ln (1.00)}$$

The time series of this elasticity is reported in Figure 10. Aggregate exports do respond in advance of the change in the actual tariff, though the magnitude of the response to positive news is small.

In a second version of this exercise, we assume that tariffs are equal to 0 through year 0. At the beginning of year 1 it is announced unexpectedly that from the start of year 5, tariffs will rise to 5%. The elasticity based on this experiment (and the corresponding counterfactual) is also reported in Figure 10. In the case of bad news, there is a non-trivial response in advance of the change in the actual tariff.

8 Conclusion

We use the joint dynamics of quantities, prices, and exit post-export entry to provide reduced form and structural evidence of an important role for costly investment in customer base and slow learning about demand in explaining how exporters grow.

We show that conditional on costs and completed spell length, the initial years of a successful export spell are characterized by growth in revenues and quantities, but there are no dynamics in prices. At the same time, higher quantities on entry predict survival in an export market, but there is no relationship between initial prices and survival.

We argue that the fact that quantities grow while prices do not change points strongly to a nonprice mechanism through which firms expand their demand. This rules out quantity dynamics arising purely from learning about the components of idiosyncratic demand, pointing instead to a role for customer base. Moreover, it suggests that firms accumulate customer base through advertising and marketing, rather directly through lagged sales. At the same time, the downward slope of the exit hazard combined with the relationship between initial quantities and survival is suggestive of a role for learning about idiosyncratic demand.

We estimate a model with accumulation of customer base through non-price actions, costs of adjusting customer base, and slow learning about idiosyncratic demand. To match the behavior of prices, we assume firms set prices as a constant markup over marginal cost. Our model can match quantitatively the behavior of quantities, prices and exit. Our parameter estimates imply that there are costs of adjusting customer base, and learning about idiosyncratic demand is slow. Both of these features constitute market-specific demand-side frictions, which slow down firms' adjustment to shocks. These findings have important implications for the responses of aggregate exports to changes in trade policy and movements in real exchange rates, both of which we plan to explore in future work.

References

[1] Albornoz, F., H. F. Calvo Pardo, G. Corcos, and E. Ornelas (2012), "Sequential Exporting," *Journal of International Economics* 88 (1), 17-31.

[2] Alessandria, G. and H. Choi (2014), "Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization," *Journal of International Economics* 94 (2), 207-223.

[3] Alessandria, G., H. Choi and K. Ruhl (2015), "Trade Adjustment Dynamics and the Welfare Gains from Trade," mimeo.

[4] Altonji, J. and R. Shakotko (1987), "Do Wages Rise with Job Seniority?" *Review of Economic Studies* 54, 437-459.

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

[6] Arkolakis, C., T. Papageorgiou, and O. Timoshenko (2015), "Firm Learning and Growth," mimeo.

[7] Bagwell, K. (2007), "The Economic Analysis of Advertising," in M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization*, vol. 3, 1701-1844, Amsterdam: North-Holland.

[8] Bastos, P., D. A. Dias and O. Timoshenko (2016), "Learning, Prices, and Firm Dynamics," mimeo.

[9] Berman, N., P. Martin and T. Mayer (2012), "How Do Different Exporters React to Exchange Rate Changes," *Quarterly Journal of Economics* 127, 437-492.

[10] Berman, N., V. Rebeyrol, and V. Vicard (2015), "Demand Learning and Firm Dynamics: Evidence from Exporters," mimeo.

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

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

[13] Eaton, J., M. Eslava, M. Kugler, and J. R. Tybout (2008), "Export Dynamics in Colombia: Firm-Level Evidence," in E. Helpman, D. Marin, and T. Verdier, eds., *The Organization of Firms in a Global Economy,* 231-272, Cambridge, MA: Harvard University Press.

[14] Eaton, J., M. Eslava, D. Jinkins, C. Krizan, and J. R. Tybout (2014), "A Search and Learning Model of Export Dynamics," mimeo.

[15] Fitzgerald, D., Y. Yedid-Levi and S. Haller (2016), "Costs of Adjustment and Exporter Behavior," mimeo.

[16] 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.

[17] Foster, L., J. Haltiwanger, and C. Syverson (2016), "The Slow Growth of New Plants: Learning About Demand," *Economica* 83, 91-129.

[18] Gallaway, M., C. McDaniel and S. Rivera (2003), "Short-run and long-run industry-level estimates of U.S. Armington elasticities," *North American Journal of Economics and Finance* 14, 49-68.

[19] Gourieroux, C., A. Monfort and E. Renault (1993), "Indirect Inference," *Journal of Applied Econometrics* 8 (SI), 85-118.

[20] Gourio, F., and L. Rudanko (2014), "Customer Capital," *Review of Economic Studies* 81 (3), 1102-1136.

[21] Harrigan, J., X. Ma and V. Shlychkov (2015), "Export Prices of U.S. Firms," *Journal of International Economics* 97, 100-111.

[22] Hooper, P., K. Johnson, and J. Marquez (2000), "Trade Elasticities for the G-7 Countries," *Princeton Studies in International Economics*, 87.

[23] Hottman, C., S. Redding, and D. Weinstein (2016), "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, forthcoming.

[24] ISGEP (2008), "Understanding Cross-Country Differences in Exporter Premia: Comparable Evidence for 14 Countries," *Review of World Economics* 144, 596-635.

[25] Jovanovic, B. (1982), "Selection and the Evolution of Industry," *Econometrica* 50 (3), 649-670.

[26] Lawless, M. (2009), "Firm Export Dynamics and the Geography of Trade," *Journal of International Economics* 77 (2), 245-254.

[27] 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.

[28] Piveteau, P. (2016), "An Empirical Dynamic Model of Trade with Customer Accumulation," mimeo.

[29] Rauch, J. (1999), "Networks versus Markets in International Trade," *Journal of International Economics* 48 (1), 7-35.

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

[31] Ruhl, K., and J. Willis (2016), "New Exporter Dynamics," *International Economic Review*, forthcoming.

[32] Timoshenko, O. A., (2015a), "Learning versus Sunk Costs Explanations of Export Persistence," *European Economic Review* 79, 113-128.

[33] Timoshenko, O. A., (2015b), "Product Switching in a Model of Learning," *Journal of International Economics* 95 (2), 233-249.

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