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

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 and Matt Schapiro. 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, Daniel Xu, and seminar participants at the NBER Summer Institute (Price Dynamics), the NBER IFM meetings, the NBER Macroeconomics Within and Across Borders meetings, LETC, DEGIT, SED, CREI, the Minneapolis Fed, Ohio State, Rochester, and UBC 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.

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

ABSTRACT

We show that after firms enter new export markets, there are striking dynamics of quantities, but no dynamics of prices, controlling for both costs and selection. This points to an important role for demand in the growth of successful exporters, and to a nonprice mechanism through which quantity demanded grows. A model where firms engage in costly investment in customer base through marketing and advertising, and learn about their idiosyncratic demand, can qualitatively match these facts, along with a declining exit hazard. We structurally estimate the model and find that costs of adjusting customer base are key to explaining how exporters grow.

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

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

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

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

How Exporters Grow*

Doireann Fitzgerald[†] Stefanie Haller[‡] Yaniv Yedid-Levi[§]

January 2016

Abstract

We show that after firms enter new export markets, there are striking dynamics of quantities, but no dynamics of prices, controlling for both costs and selection. This points to an important role for *demand* in the growth of successful exporters, and to a *nonprice* mechanism through which quantity demanded grows. A model where firms engage in costly investment in customer base through marketing and advertising, and learn about their idiosyncratic demand, can qualitatively match these facts, along with a declining exit hazard. We structurally estimate the model and find that costs of adjusting customer base are key to explaining how exporters grow.

1 Introduction

Recent years have seen increased interest in the role of demand in explaining firm dynamics. Exporting firms provide a unique laboratory in which to examine this issue, because we have good data on sales broken down by market, allowing us to separate out market-specific

*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 and Matt Schapiro. 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, Daniel Xu, and seminar participants at the NBER Summer Institute (Price Dynamics), the NBER IFM meetings, the NBER Macroeconomics Within and Across Borders meetings, LETC, DEGIT, SED, CREI, the Minneapolis Fed, Ohio State, Rochester, and UBC 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

demand from factors such as productivity that affect a firm's sales in all markets. We use rich data on firms and exports for Ireland to show that demand plays a quantitatively important role in explaining how exporters grow in a market and to provide evidence that nonprice actions such as marketing and advertising play a key role in expanding demand.

More concretely, in this paper we do two things. We first document a novel set of facts about the evolution of quantities and prices post export entry. Controlling for both costs and selection, we show that (1) quantities grow dramatically in the first five years of successful export spells;[1] (2) within successful export spells, there are no statistically or economically significant dynamics of prices; (3) higher initial quantities predict longer export spells; (4) initial prices do not help to predict export spell length. In sum, there are very striking dynamics of quantities with respect to market tenure (defined as the age of an export spell), but no statistically or economically significant dynamics in prices.

The post-entry dynamics of quantities cannot be generated by any factor that operates at the level of the firm or the firm-product, as our empirical strategy controls for this dimension of heterogeneity. This rules out explanations based on dynamics in productivity, capital, or financial constraints. Instead, these dynamics must arise from something that varies at the level of the firm-market pair. This clearly points to an important role for *demand* in explaining exporter growth.

At the same time, the *joint* behavior of quantities and prices rules out a number of possible demand-based explanations for post-entry exporter dynamics and, by extension, demand-based explanations for firm dynamics in general. The absence of price dynamics implies that quantity growth cannot be purely a by-product of learning about idiosyncratic demand (à la Jovanovic (1982)), as this would require prices to move in the opposite direction to quantities. It also rules out the possibility that high current sales lead to higher future sales, as this would give firms an incentive to set initial markups below their long-run optimum in order to acquire customer base. Instead, our quantity and price findings are jointly suggestive of a role for customer base that is accumulated through *nonprice* actions of the firm, for example, advertising and marketing.

The second thing we do in the paper is to propose a model of post-entry export dynamics that can match qualitatively the post-entry behavior of quantities and prices, as well as the behavior of the exit hazard conditional on costs. We then estimate the model.

In our model, firms can invest in customer base through expenditure on marketing and advertising (i.e., nonprice actions), thus shifting demand conditional on prices. They also

[1] An export spell is defined as a consecutive series of years in which a firm exports a product to a destination market. We call an export spell that lasts seven years or more a *successful* spell.

learn about their idiosyncratic demand. Firms set prices (in the face of uncertainty about demand, it matters whether they set prices or quantities) and face constant elasticity of demand, so markups are constant, generating a flat path of prices conditional on costs. Slow convergence to steady state customer base, and hence post-entry growth of quantities, arises because of costly and irreversible investment, or gradual learning at the firm level about idiosyncratic demand, or a combination of both. Finally, learning about demand leads to a declining exit hazard, as firms exit a market on learning that they have low idiosyncratic demand.

The estimated model matches the relationship between initial quantities and export spell length, the evolution of quantities in export spells of different lengths, and the evolution of the exit hazard. By construction, it matches the behavior of prices. Key parameter estimates are in line with the literature on advertising, and the implied share of revenue devoted to advertising and marketing is consistent with data for US manufacturing firms.[2] Parameter estimates imply an important role for *both* costs of adjusting customer base *and* gradual learning about demand in explaining post-entry export dynamics. In particular, while gradual learning is necessary to match the exit hazard, costs of adjusting customer base are necessary to match the rate of growth of quantities in successful export spells.

We illustrate the quantitative importance of our estimates of the costs of adjusting customer base by simulating the response of aggregate exports to permanent unexpected changes in tariffs in our baseline model. We compare this with the response in a restricted model that shuts down costs of adjusting customer base and relies only on learning about demand to match steady state export dynamics. Our baseline model generates long-run elasticities of aggregate exports to tariffs that are significantly higher than short run elasticities. In contrast, when costs of adjusting customer base are shut down, short and long run responses are very similar.

Our work is related to several literatures. We build on the work of Roberts and Tybout (1997), Das, Roberts, and Tybout (2007), Eaton et al. (2008), Eaton et al. (2014), and Ruhl and Willis (2015), who document facts about export entry and post-entry export dynamics. This literature documents rapid growth of revenues in successful export spells. In contrast to this literature, we document quantity and price facts separately. As we show, this is key to identifying the precise mechanisms through which demand-based factors generate revenue dynamics. Crucially, we extend the approach of Ruhl and Willis (2015) to identify how quantities and prices evolve with market tenure controlling nonparametrically for both

[2]We do not have data on these expenditures for our firms.

firm-product heterogeneity (e.g., in marginal costs) and selection on idiosyncratic demand.

Contemporaneous with our work, papers by Bastos et al. (2015), Piveteau (2015), and Berman et al. (2015) also document quantity and price facts using customs data. But unlike us, Bastos et al. and Piveteau do not control for selection on firm- or firm-product-level unobserved heterogeneity, while Berman et al. do not control for selection on idiosyncratic demand. Moreover, unlike these papers, we provide very extensive evidence on the robustness of our quantity and price findings across firms, products, and markets.

A recent and growing literature in the tradition of Jovanovic (1982) explores the role of learning about demand in explaining facts about post-entry export dynamics. See, for example, Albornoz et al. (2012), Fernandes and Tang (2012), Berman et al. (2015), Timoshenko (2015a, 2015b) and Bastos et al. (2015). Arkolakis et al. (2015) have a model of firm dynamics with learning about demand that also has implications for post-entry export dynamics. We show that the behavior of prices is inconsistent with dynamics driven *purely* by learning about demand. But given a role for accumulation of customer base, we show that learning about demand is important to match post-entry export dynamics, in particular the behavior of exit.

Our work is closely related to Arkolakis (2010, 2015), who develops a static model with convex costs of acquiring customers in export markets. Eaton et al. (2014) extend this model to a dynamic setting. Our model shares many of the features of Eaton et al. (2014) but is more parsimonious and does not require transaction-level data to estimate.

By showing that export quantities are systematically related to market-specific tenure, whereas prices are not, we provide evidence that what is often labeled "quality" in trade models is likely to have a substantial customer base component. In this, our findings are similar to those of Hottman, Redding and Weinstein (2015), though they use data on consumer products and a very different empirical methodology from ours.

Finally, our work is also related to a macro literature on customer base, both empirical (Foster, Haltiwanger, and Syverson (2008, 2013)) and theoretical (Gourio and Rudanko (2014) and Drozd and Nosal (2012)). In contrast to Foster, Haltiwanger, and Syverson (2008), our empirical strategy controls for selection on idiosyncratic demand, and we do not restrict attention to homogeneous goods industries. By showing there are no post-entry dynamics of prices once selection is controlled for, we point to nonprice mechanisms for accumulating customer base. In this we are consistent with Drozd and Nosal (2012), and in contrast to Foster, Haltiwanger, and Syverson (2013) and Gourio and Rudanko (2014), who posit that firms initially distort markups downward in order to acquire customer base.

The paper is organized as follows. In the second section, we describe our data and summary statistics. In the third section, we describe our empirical strategy. In the fourth section, we describe 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 show that several popular models of post-entry dynamics are inconsistent with these facts and describe a model that is capable of matching the facts qualitatively. In the sixth section, we describe how we estimate the model, report our estimation results, and discuss their interpretation. In the seventh section, we perform a simulation exercise that illustrates the model's predictions for aggregate export responses to tariff changes. 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 appendix to Fitzgerald and Haller (2015); this appendix also describes a third data set (the PRODCOM survey), which we use to obtain the number of products produced at the firm level.

2.1 Census of Industrial Production

The CIP, which covers manufacturing, mining, and utilities, takes place annually. Firms with three or more persons engaged are required to file returns.[3] 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 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.

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

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)[4] 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 CSO informs us that their ability to match customs records to firm identifiers is best for the period after 1999. For the period 2000-2009, customs records that match to a firm identifier account for 76% of published merchandise exports. The appendix to Fitzgerald and Haller (2015) provides additional summary statistics 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 reporting threshold for intra-EU exports (635,000 euro per year in total shipments within the EU) is different from that for extra-EU exports (254 euro per transaction).[5] The high threshold for intra-EU exports likely leads to censoring of exports by small exporters to the EU. However, since the threshold is not applied at the market level but to exports to the EU as a whole, 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).[6] 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 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

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

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

[6]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.

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 in Irish manufacturing is high. At least half of exporters participate in multiple markets, and we observe exports to 140 distinct markets in the period 1996-2009. 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 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.[7] This is illustrated in Table 2, which reports summary statistics on churn in the number of export markets from year to year. In any given year, on average 49% of exporters change the *number* of markets they participate in. This 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 that are systematically related to market tenure while controlling for time-varying heterogeneity at the firm-product level. This can be heterogeneity in marginal cost[8] (and for brevity we often say "controlling for marginal cost") but it can also be heterogeneity in demand that affects all markets equally (e.g., because of differences in product quality).

At the same time, in order to identify true dynamics of quantities and prices related to market tenure, we must also control for selection on unobserved idiosyncratic demand: firms are likely to stay longer in markets where idiosyncratic demand is high than in markets

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

[8] 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.

where idiosyncratic demand is low. The structure of our data allows us to control for both dimensions of heterogeneity in a straightforward way.

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

First, it is instructive to show how we construct two key variables, *market tenure* and *spell length*. The top panel of Table 3 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.

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. An export spell is defined as a period of continuous export participation. 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).[9] If the spell is neither left- nor right-censored, we observe completed spell length (markets B, C, D, E). 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).

Market A (a spell that is both left- and right-censored) is a common occurrence in our data. Neither market tenure nor spell length is observed for these spells.

As Table 3 illustrates, there is cross-sectional variation in both market tenure and spell length within a firm-product-year. By focusing on variation within a firm-product-year, we control for marginal cost. By further focusing only on variation in market tenure within a given spell length, we control for selection on idiosyncratic demand, and by focusing only on variation in spell length given a market tenure of 1, we can put some bounds on the variation in idiosyncratic demand.

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 total 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

[9] 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.

multiple export spells of different length for firm i, product j, and market t 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.[10] 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 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 both 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.

[10] 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. The choice to top-code at 7 years is a compromise. 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.

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 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.[11] 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. We organize our results into initial values conditional on spell length and spell trajectories, normalizing the start of

[11] 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.

each spell to 0. The results are illustrated in Figures 1, 2, and 3, which graph the trajectories of revenues, quantities and prices implied by taking the exponential of the relevant sums of coefficients from Table 4.[12] The online appendix compares summary statistics on the firm-years included in this analysis with those for all firm-years in our data. Exporters included in the analysis of product revenue, quantity, and price are on average bigger and more export-intensive, and export to more destinations than the average exporter.

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 by part-year effects (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 help predict export spell length.

Additional findings on quantities and prices are as follows. 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 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.

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.

Finally, as measured by the coefficients on the market dummy variables, quantities (and prices) are systematically higher in some markets than in others. The correlation between the estimated coefficients in the quantity equation and the estimated coefficients in the price equation is weakly negative (-0.07), that is, higher quantities are weakly associated with lower prices.

The first column of Table 5 reports results for the baseline estimation of equation (2).

[12] 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.

The probability of exit is initially decreasing in market tenure and then flattens out after four years in a market. 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.[13] 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 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.[14] 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 the full set of results (including some not described here) are reported in an online appendix to the paper.

[13] 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.

[14] 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 or product-market-year fixed effects (if 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. 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 sample, and in the range 7 to 10 in our 19-year sample. The key results are qualitatively unchanged. 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%).[15] 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.

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. 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. Quantitatively, the one statistically

[15] This does not affect our conclusion that prices are "flat" very much because if we use the point estimates of quantity deviations from year 2 quantities (thus avoiding contamination from part-year effects) and corresponding price deviations from year 2 prices to calculate implied price elasticities of demand, we obtain numbers in the range 18-149.

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. 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. 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 most spells are first 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. 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 then estimate different trajectories based on firm size (as measured by employment) at the time of firm-product-market or firm-market entry. 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. This categorization is based on the NACE Revision 1.1 three-digit sector of the firm.[16] Estimates for consumer nonfood nondurables and consumer durables are noisy, as

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

there are relatively few firms in these industries.[17] 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. The key stylized facts are qualitatively replicated for all product types.

4.4 Relation to the literature

As mentioned in the introduction, 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 (2015), Timoshenko (2015a)), though we extend the analysis of these papers to cleanly characterize the relationship between initial revenue and completed spell length, and the evolution of revenues in short spells as well as successful spells of long duration.

Contemporaneous with our analysis, Berman et al. (2015) and Piveteau (2015) document quantity and price facts based on customs data for France, while Bastos et al. (2015) document quantity and price facts based on data for Portugal. None of these papers estimate our exact empirical specification. Berman et al. do not allow the evolution of quantity and price with market tenure to depend on spell length, while Piveteau does not control for firm-product level heterogeneity. Meanwhile, Bastos et al. do not control for either selection or firm-product heterogeneity. As a result, the findings of these papers are not directly comparable to ours.[18] However, in the specifications that are closest to ours (i.e., go closest to controlling for both firm-product heterogeneity and selection on demand), these papers find that prices are flat with respect to market tenure.[19] This confirms that our findings on the behavior of prices are not driven by some special feature of the Irish data.

and Haller (2013).

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

[18]Piveteau's baseline specification does not control for firm-level unobserved heterogeneity and finds prices increasing with market tenure. Bastos et al. (2015) have similar findings for Portuguese firms. When we estimate Piveteau's specification on our data, our results are similar to his, consistent with firms with high average market tenure charging systematically higher prices than firms with low average market tenure.

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

Foster, Haltiwanger, and Syverson (2008) use a very different empirical strategy and data for a select set of US manufacturing industries to show that, controlling for productivity, older firms charge higher prices. As we note in the introduction, our empirical strategy differs from theirs in that we control for selection on idiosyncratic demand, and we do not restrict attention to homogeneous goods industries.

Our results on the importance of the product extensive margin are quantitatively very consistent with those of Hottman, Redding, and Weinstein (2015) on the contribution of the extensive margin of products to explaining variation in firm size.

5 A model of post-entry export dynamics

5.1 Motivation

We now describe a model of the firm's problem that can match qualitatively all the facts we document. Before going into the details, we motivate our modeling choices and give some intuition.

First, and most importantly, we focus on dynamics that are driven by demand, not supply. 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 we document, since all our estimates are conditional on firm-year or firm-product-year fixed effects as appropriate. Since the product extensive margin accounts for at most 20-30% of dynamics of revenue at the market level, we ignore this margin and build a model of a single-product firm. We focus on a model of a monopolistically competitive firm that faces downward-sloping demand but takes competitors' actions as given.[20] Finally, we do not build a 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. To capture this, we make entry at the market level stochastic.

Within this class of models, the literature has proposed several ways to generate post-entry export dynamics. These can be illustrated as follows. Let Q_t^{ik} be the quantity firm i sells in market k at date t. Let P_t^{ik} be the price the firm charges to buyers from k. Let ε_t^{ik} be a shock idiosyncratic to the firm and the market, which shifts demand conditional on price

[20]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, but the price-taking assumption is unlikely to apply to all the firms and markets for which we document similar patterns in the data.

and is exogenous to the firm. Let D_t^{ik} be 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.[21] Firm i's demand in market k at time t is given by[22]

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

Post-entry dynamics in export revenue ($P_t^{ik}Q_t^{ik}$) conditional on firm-year and market-year effects can arise because of dynamics in ε_t^{ik},P_t^{ik} , D_t^{ik}, or a combination of all three. It turns out that the behavior of quantities and prices is very helpful in distinguishing between competing explanations for dynamics in revenue.

Exogenous dynamics: Post-entry export dynamics in quantities can be generated purely through the process for ε_t^{ik}. If the price elasticity of demand does not depend on the idiosyncratic demand shock, there are no corresponding dynamics in prices. We assume the price elasticity of demand does not depend on ε_t^{ik}, and we allow for a class of processes for ε_t^{ik} that is consistent with bounded market shares.

Learning: A growing trade literature models post-entry export dynamics as arising from *learning* about the process for idiosyncratic demand, following Jovanovic (1982).[23] This can induce post-entry dynamics in quantities through inducing dynamics in prices P_t^{ik}. Informally, this works as follows.[24] A potential entrant has a prior belief about the distribution of ε_t^{ik}. Since it does not observe ε_t^{ik}, it matters whether it chooses prices or quantities. Suppose it chooses *quantities* in order to maximize expected profit given its belief. Ex post, by observing the realized price of the good in the market, it infers the true value of ε_t^{ik}. A high price implies high idiosyncratic demand, whereas a low price implies low idiosyncratic demand. The incumbent firm uses this realization to update its belief about the distribution of ε_t^{ik}. Given this updated belief, next period it chooses whether or not to participate, and conditional on participation, adjusts its quantities to maximize

[21] The trade literature has a tradition of referring to anything that shifts demand conditional on price as "quality." To explain the facts we document, "quality" would have to be market-specific, systematically related to tenure in a market, and unrelated to marginal cost (as this would generate price dynamics). This sounds remarkably like customer base to us. This does not rule out a firm-product-specific quality that *is* related to marginal cost, but is the same across all markets served by the firm. This is absorbed by our firm-product-year fixed effects.

[22] 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.

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

[24] We show this formally in the online appendix.

expected profit given its new belief. Seeing high realizations induces the firm to increase its quantities (and hence realized prices fall), whereas seeing low realizations induces it to reduce its quantities (and hence realized prices rise) or, alternatively, to exit. The converse is true when the firm sets prices instead of quantities.

The behavior of *revenue* that we see in the data could be driven by learning of this type, but the joint behavior of quantities and prices could not. Learning can generate dynamics in quantities *only* if there are corresponding (and opposite) dynamics in realized prices.[25] In addition, the fact that initial quantities are positively correlated with spell length, whereas initial prices do not help forecast spell length, is consistent with firms setting prices rather than quantities in an environment with uncertainty about idiosyncratic demand. But in a pure learning model, if firms set prices rather than quantities, prices should (weakly) rise and quantities should (weakly) fall with market tenure in successful spells.[26]

However, learning about idiosyncratic demand remains an appealing ingredient in a model of post-entry export dynamics, because it can rationalize the large number of export spells that start small and last only one period, as well as the fact that the exit hazard initially declines rapidly and eventually flattens out. We incorporate learning about demand into our model in such a way that it does not generate dynamics in prices. We do this by assuming that firms set prices (rather than quantities) and face constant elasticity of substitution (CES) demand, so the optimal markup over marginal cost is constant. We also simplify the process of learning by making the evolution of beliefs stochastically dependent on past participation rather than dependent on past realizations of prices through Bayes' law.

Customer base: Dynamics in customer base D_t^{ik} can generate dynamics in quantities.[27] Foster et al. (2013) and Piveteau (2015) assume that customer base is increasing in past sales. If this is how customer base is accumulated, even if customer base does not affect the price elasticity of demand, 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.

However, if D_t^{ik} is accumulated purely through *nonprice* actions of the firm (such as expenditures on marketing and advertising), dynamics in D_t^{ik} can generate quantity dynamics in the absence of price dynamics (as long as the price elasticity of demand is unaffected by

[25] Under its belief, the firm does not anticipate price dynamics; but we observe realized prices, not firms' expectations of prices.

[26] Both prices and quantities are flat under CES demand and price setting.

[27] Ruhl and Willis (2015) generate dynamics by assuming that demand depends deterministically on the history of participation.

D_t^{ik}). We assume that this is how customer base is accumulated, in a process analogous to what is usually assumed for accumulation of physical capital.

In order to generate the gradual convergence of quantities to steady state that we observe in successful export spells, there must be some force that slows down accumulation of D_t^{ik}. One possibility is that there are costs of adjusting customer base (as in, e.g., Arkolakis (2010, 2015) and Eaton et al. (2014)). Just like costs of adjusting physical capital, this can lead to a slow transition to steady state customer base. Another possibility is that learning about the process for ε_t^{ik} can slow down accumulation (as in, e.g., Eaton et al. (2014)). This is because the firm's optimal customer base depends on its belief about idiosyncratic demand. As beliefs evolve, so too does investment and hence customer base. This mechanism has the added attraction that, as just noted, it can explain the evolution of the probability of exit with market tenure. We allow for both costs of adjustment and learning about demand in our model.

5.2 Model description

Firm i has marginal cost C_t^i, in terms of some numeraire. It faces a random sunk cost S_t^{ik} of participating in market k. With probability λ (independent across firms, markets, and over time), the sunk cost is equal to $S < \infty$. With probability $1-\lambda$, the sunk cost is infinity. It also faces a random fixed cost of participating in market k given by F_t^{ik}. With probability $1-\omega$ (independent across firms, markets, and over time), the fixed cost is equal to $F < \infty$. With probability ω, the fixed cost is equal to infinity. These random 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 idiosyncratic to the firm and the market.

Demand for firm i in market k depends on its own price, on customer base, and on idiosyncratic demand. It takes the following form:

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

Idiosyncratic demand has two components, permanent and transitory: $\varepsilon_t^{ik} = \nu^{ik} + \eta_t^{ik}$, with $\nu^{ik} \sim N\left(0, \sigma_\nu^2\right)$ and $\eta_t^{ik} = \rho\eta_{t-1}^{ik} + \zeta_t^{ik}$, with $\zeta_t^{ik} \sim N\left(0, \sigma_\eta^2\right)$. If $\alpha \in (0,1)$ there is a finite positive steady state for customer base D_t^{ik}. Customer base accumulates as follows:

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

where A_t^{ik} is investment in customer base, which we assume to be subject to both convex costs of adjustment governed by ϕ and irreversibility, since it is likely that expenditures on advertising and marketing are sunk:

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

The structure of information is as follows. When making choices about participation, investment, and prices, both potential entrants ($X_{t-1}^{ik} = 0$) and incumbents ($X_{t-1}^{ik} = 1$) observe C_t^i, F_t^{ik}, and S_t^{ik} and know the processes from which they are drawn.[28] However, they do not observe their current idiosyncratic demand ε_t^{ik}. Potential entrants use the unconditional distribution of $\nu^{ik} + \eta_t^{ik}$ to form expectations about current and future ε_t^{ik}. Incumbents may be "uninformed" (like potential entrants) or "informed." With probability γ per period, previously uninformed incumbents become informed. This implies that they observe ν^{ik} and η_{t-1}^{ik} and 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.[29] This way of modeling information acquisition is flexible, in that it allows the average speed of learning to be fast (γ close to 1) or slow.

Let N_t^{ik} be an indicator variable that takes the value 0 if the firm is uninformed in market k 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:

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

Since there is uncertainty about demand, it matters whether the firm sets prices or quantities. We assume it sets prices. Because demand is CES and we assume monopolistic competition, 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, our model hardwires in a flat path of prices with respect to market tenure.

[28] S_t^{ik} is irrelevant to the choices of incumbents.

[29] This assumption is made for computational tractability.

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

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

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

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.

This model captures two key features of Eaton et al. (2014), that is, learning about demand through participation and investment in future customer base, in a relatively tractable reduced form model.

6 Model estimation and results

6.1 Estimation

We use a 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 a process for C_t^i, we discretize both exogenous and endogenous states[30] 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 construct the population equivalents of the moments we estimate in Section 4. The goal of our estimation is to choose the vector of parameters that best matches a subset of the moments estimated in the previous section.

We match moments of four types: the ratios of initial quantities across spells of different length; the evolution of quantities with market tenure within spells of different length; the average exit rate in the first year in a market (we pick the US market); and the evolution of exit probabilities with market tenure, all of these conditional on costs. We do not target the initial size of longer spells relative to one-year spells, or moments based on quantities in

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

the first and last year of a spell relative to other years of the spell, as asymmetric part-year effects may be at work, and these are not present in our model. We do not target moments related to prices, as they are matched automatically given the structure of the model. The full set of moments is reported in Table 6. We also match exactly the average rate of entry into exporting in our data.

We first preset some parameters not identified by our target moments. Since our data are annual, we set $\beta = 1.05^{-1}$. 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%. The key data moments we try to match are estimated conditional on costs, so we normalize $C_t^i = 1$. Results are not sensitive to this choice.[31] 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 $\lambda(1-\omega)$, while the exit rate converges to ω as market tenure gets very large. Based on an annual rate of entry into exporting of 6%,[32] we set $\lambda = 0.06/(1-\omega)$.

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 first do a global grid search of the parameter space and then use a derivative-free algorithm to optimize starting from local minima of the global grid search.[33]

6.2 Baseline results

Table 6 reports the data moments in the first column and the corresponding moments from the baseline model in the second column. Figures 7, 8, and 9 illustrate the fit of the model in terms of initial quantities, growth of quantities, and the evolution of the probability of exit with market tenure. 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.

[31] The online appendix describes results based on alternative values for θ.
[32] This is at the level of exporting as a whole, not at the level of an individual market.
[33] We use a combination of a particle swarm algorithm and the simplex method.

Table 7 reports our estimated parameters.[34] Our estimate of α, the elasticity of sales with respect to customer base, is 0.54. 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, 0.52, is in line with the findings reported in Bagwell (2007) of an empirical literature on the annual depreciation rates of brand advertising.[35] Our estimated value of ϕ is large, consistent with substantial adjustment costs à la Arkolakis (2010). In order to provide some intuition for the magnitude of ϕ, we calculate average expenditure on marketing and advertising (i.e., $A_t^{ik} + c\left(D_t^{ik}, A_t^{ik}\right)$) as a share of expected revenue. Figure 10 shows how the average share evolves with market tenure for "successful" export spells, (i.e., those lasting at least seven years). These firms spend 14% of expected revenues on marketing and advertising in the initial years in a market, but this declines to just under 6% of revenues by year 6. Unfortunately, we do not have data on marketing and advertising expenditures for firms in our data. However, our estimates of the share of revenue devoted to these expenditures in the long run are in the ballpark of the 5% average for manufacturing reported by the Duke University Chief Marketing Officer Survey (2015).

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 0.87. This points to a relatively rapid rate of learning. We calculate the share of participants who are informed at each market tenure and find that by the fourth year, 99.6% of incumbents are informed.[36]

We find that there are fixed costs of export participation, but they are tiny, amounting to 0.03% of expected revenue in the initial year in a market. This is because expenditures on advertising and marketing act to some extent like a fixed cost as in Arkolakis (2010). Finally, we estimate a modest probability of the firm-market experiencing a "death shock," since $\omega = 0.03$. Most of the exit out to seven years in a market is endogenous.

6.3 Identification

In order to illustrate which moments of the data are key to identifying different parameters of the baseline model, we reestimate the baseline model, targeting in turn different subsets of

[34]In the online appendix we report standard errors constructed using the method suggested by Gourieroux, Monfort and Renault (1993). Standard errors based on this method are small because the numerical derivatives of moments with respect to individual parameters are very large.

[35]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.

[36]The shares of informed participants are reported in the online appendix.

the moments we target in our baseline estimation. The results are reported in Table 8. From this exercise we learn the following. Within-spell quantity trajectories are key to identifying jointly a modest value of α, a high rate of depreciation of customer base (high δ), a modest degree of costs of adjustment (low ϕ), and relatively fast learning (high γ). Within-spell quantity trajectories, together with the dispersion of initial quantities, are jointly key to identifying a relatively modest dispersion of the innovation of the transitory component of idiosyncratic demand (relatively low σ_η^2). The exit trajectory is key to identifying relatively slow learning ($\gamma < 1$), a modest degree of costs of adjustment (low ϕ), and low fixed costs F. Finally, within-spell quantity trajectories for short spells (spells where we observe in-sample exit) are key to identifying a modest degree of costs of adjustment (low ϕ).

6.4 Inspecting the mechanism

In order to examine which features of the model are key to matching the moments, we first eliminate these features one by one, and reestimate the model, targeting the full set of moments.

First, we shut down customer capital by setting $\alpha = 0$ (the parameters δ and ϕ become redundant). Now the only intertemporal choice faced by the firm is whether or not to participate.[37] Table 6 reports the fit and Table 7 reports the estimated parameters of this model, while we provide results in the form of figures in the online appendix. This model can generate an increasing relationship between initial quantities and spell length, and a reduction in exit hazard with market tenure. But it is incapable of generating within-spell quantity growth. This illustrates the fact that if a model with Jovanovic-style learning about demand is set up to match the facts about within-spell price dynamics, it cannot simultaneously match the facts about quantity dynamics. Accordingly, the fit of this model is poor.

Second, 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).[38] Now the only source of dynamics in quantities after period 2 arises from learning. Table 6 reports the fit and Table 7 reports the estimated parameters of this model (figures are provided in the online appendix). Qualitatively, the model is a success. It can generate an increasing relationship between initial quantities and spell length, within-spell quantity growth in the

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

[38] In the online appendix, we report results based on estimating the model shutting down in turn convex adjustment costs and irreversibility of investment in customer base.

longest export spells, and a reduction in exit hazard with market tenure. Relative to the baseline model, it has difficulty matching the quantity trajectories in shorter spells, and it does a somewhat worse job in capturing the fact that in the data, the exit hazard flattens out after about five years. The degree of slow learning required to fit the moments ($\gamma = 0.58$) is such that even after seven years in the market, 2% of participants remain uninformed.

Third, we assume that learning takes only one period by setting $\gamma = 1$. 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. Table 6 reports the fit and Table 7 reports the estimated parameters of this model (figures are provided in the online appendix). The fit is worse than the baseline model, though better than that of the model with slow learning and no adjustment costs. Qualitatively, the model is a success. It can generate an increasing relationship between initial quantities and spell length, within-spell quantity growth in the longest export spells, and a reduction in exit hazard with market tenure. Relative to the baseline model, it has difficulty in matching the quantity trajectories in shorter export spells, overpredicts the decline in exit hazard between period 1 and period 2, and underpredicts the decline in subsequent periods.

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 fit and Table 7 reports the estimated parameters (figures are provided in the online appendix). Qualitatively, this model can reproduce the key facts, but it is unable to generate different quantity dynamics in export spells of different lengths. In this sense, although our point estimate of ρ may seem low, we can see that it is crucial to allow for autocorrelation of transitory idiosyncratic demand in order to match the behavior of short spells.[39]

To sum up, these experiments confirm the view that accumulable customer base plays a key role in simultaneously matching post-entry export dynamics of quantities and prices. We find that qualitatively, either costs of adjusting customer base or gradual learning could explain the dynamic evolution of quantities and exit. However quantitatively, both of these mechanisms play a role. This is important because costs of adjustment can generate dynamics in the responses of large incumbent exporters to shocks, in a way that slow learning about idiosyncratic demand cannot. We illustrate this in the next section.

[39] 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.

7 Simulation

Our estimated model of steady state post-entry export dynamics has important implications for how firms respond to shocks. We now perform a simple simulation exercise to illustrate this. In particular, we assess the implications of our estimated model for the elasticity of aggregate exports with respect to a permanent unexpected change in tariffs in the destination market. We focus on whether this elasticity varies with the time horizon (i.e., whether long-run responses differ from short-run responses), and whether there are significant asymmetries in responses to tariff reductions and tariff increases.

Because we have a purely partial equilibrium model of the firm's problem, and moreover, the probability of entry depends only on the estimated parameters λ and ω, our exercise has a restricted interpretation. Implicitly, we hold domestic costs and foreign aggregate demand constant. Moreover, our exercise cannot match the response of the rate of entry into an export market to a change in tariffs. Instead, we are able to capture that part of the response of exports that is due to the responses of incumbents and of firms who would have entered anyway. We note that Fitzgerald and Haller (2015) estimate modest responses of export participation to tariff changes using the same firm and customs data as in this paper.

The exercise is as follows. 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 solve for the policy functions using in turn a zero current and expected tariff ($\tau_L = 1$), and a 10% current and expected tariff ($\tau_H = 1.1$). We then simulate exports for a panel of 5000 firms, first under the assumption that tariffs in the destination market are initially equal to τ_H and then fall unexpectedly and permanently to τ_L, and second under the assumption that tariffs are initially equal to τ_L and then rise unexpectedly and permanently to τ_H. In each case, we start all firms as non-exporters and simulate the model for 80 periods before shocking the tariff.[40] We then simulate the panel of firms for 20 years after the tariff change and calculate the elasticity of aggregate exports with respect to the tariff change at all horizons from 1 to 20 years as follows, letting R_0 denote aggregate exports in the period before the change in the tariff, and R_t denote aggregate exports t periods after the change:

$$\sigma_t = \frac{\ln R_t - \ln R_0}{\ln \tau_t - \ln \tau_0}$$

[40]In the absence of general equilibrium feedback effects from resource demand to costs, it takes a long time for aggregate exports to converge to steady state.

We repeat this exercise 50 times and calculate the average of the resulting elasticities across the 50 repetitions as well as the 95% confidence interval. We do this first using our baseline model and parameter estimates, second, using the restricted model where we shut down costs of adjusting customer base (learning only), and third, using the model with costs of adjustment but one-period learning (one-period learning).

Figure 11 illustrates for the baseline model the mean elasticity at all horizons between 1 year and 20 years based on the tariff reduction (in blue) and the tariff increase (in red). The elasticity is increasing and convex in the time horizon in both cases, and there do not appear to be strong asymmetries in responses.

In Figure 12 we compare the implications of the baseline model with those of the two alternative models, by reporting σ_t/σ_1 for both the tariff decrease and the tariff increase. The elasticities for the one-period learning model are qualitatively similar to those in our baseline model. However, the elasticities in the learning-only model behave very differently. Short-run and long-run elasticities do not differ much, and moreover, there is an asymmetry in responses to tariff reductions and tariff increases.

We conclude from this exercise that our baseline model not only fits the pattern of steady state dynamics better than a model without costs of adjusting customer base, but also has a better chance of matching differences between short- and long-run responses to trade liberalizations.

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 in explaining how exporters grow.

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

We then 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 from pure 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 than through lagged sales.

We estimate a model with this feature in addition to both costly adjustment of customer base and Jovanovic-style learning about the components of idiosyncratic demand. Our model can match the facts we document and implies reasonable estimates of the share of revenue devoted to expenditures on advertising and marketing. Both costs of adjusting customer base and slow learning about idiosyncratic demand are necessary in order to match quantitatively the dynamic evolution of quantities and exit.

Finally, we show that our baseline model can generate quantitatively important differences in short- and long-run elasticities of aggregate trade to permanent unanticipated tariff changes, without relying on extensive margin adjustment. This is due to costs of adjusting customer base, as a model that generates dynamics purely through the interaction of customer base accumulation with learning about demand cannot generate this difference.

Our results are important for several reasons. First, we provide clean evidence of the quantitative importance of demand-based factors in explaining firm dynamics. Second, we provide evidence on the precise mechanisms through which demand-based factors generate firm dynamics. This is important, because dynamics driven by learning have different implications for how firms respond to shocks from dynamics due to costs of adjustment. As we illustrate in our simulation, if demand-driven dynamics are due partly to costs of adjusting customer base, then large incumbent firms will respond sluggishly to shocks. However if dynamics are due to learning about idiosyncratic demand alone, these firms will respond quickly, and dynamics will be due only to the behavior of marginal participants. This has important implications both for the responses of aggregate exports to, for example, trade liberalizations and movements in real exchange rates, and more generally for how firms respond to shocks.

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

[3] Arkolakis, C. (2015), "A Unified Theory of Firm Selection and Growth," *Quarterly Journal of Economics*, forthcoming.

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

[5] 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.

[6] Bastos, P., D. A. Dias and O. Timoshenko (2015), "Learning, Prices, and Firm Dynamics," mimeo.

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

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

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

[10] 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.

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

[12] Fernandes, A., and H. Tang (2012), "Learning from Neighbors' Export Activities: Evidence from Exporters' Survival," mimeo.

[13] Fitzgerald, D., and S. Haller (2013), "Pricing-to-Market: Evidence From Plant-Level Prices," *Review of Economic Studies* 81 (2), 761-786.

[14] Fitzgerald, D., and S. Haller (2015), "Exporters and Shocks: Dissecting the International Elasticity Puzzle," mimeo.

[15] 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.

[16] Foster, L., J. Haltiwanger, and C. Syverson (2013), "The Slow Growth of New Plants: Learning About Demand," mimeo.

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

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

[19] Hottman, C., S. Redding, and D. Weinstein (2015), "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, forthcoming.

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

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

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

[23] Piveteau, P. (2015), "An Empirical Dynamic Model of Trade with Customer Accumulation," mimeo.

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

[25] 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.

[26] Ruhl, K., and J. Willis (2015), "New Exporter Dynamics," mimeo.

[27] The CMO Survey (2015), *CMO Survey Report*, August 2015, www.cmosurvey.org

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

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

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

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
Probability of entry into exporting	0.06
Probability of exit from exporting	0.12
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 revenue is concorded product export revenue 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.

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.

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

Year	1	2	3	4	5	6
Market	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 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					
	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					

Notes: Fictitious illustration of sources of identifying variation.

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

Obs. level	Firm-product-market						Firm-market			
Dep. var.	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.

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.

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

	Data		Model				
	moment	s.e.	Baseline	$\alpha = 0$	$\rho = 0$	no AC	$\gamma = 1$
$\ln(Q_1^3/Q_1^2)$	0.25	(0.04)	0.31	0.33	0.24	0.23	0.27
$\ln(Q_1^4/Q_1^2)$	0.44	(0.05)	0.42	0.45	0.43	0.40	0.43
$\ln(Q_1^5/Q_1^2)$	0.57	(0.07)	0.53	0.50	0.57	0.53	0.54
$\ln(Q_1^6/Q_1^2)$	0.58	(0.08)	0.60	0.54	0.67	0.61	0.65
$\ln(Q_1^7/Q_1^2)$	0.88	(0.05)	0.87	0.88	0.84	0.89	0.84
$\ln(Q_3^4/Q_2^4)$	0.05	(0.06)	-0.01	-0.00	0.14	-0.09	0.15
$\ln(Q_3^5/Q_2^5)$	0.07	(0.08)	0.13	-0.00	0.14	0.19	0.06
$\ln(Q_4^5/Q_2^5)$	-0.01	(0.09)	-0.01	-0.00	0.25	0.00	0.09
$\ln(Q_3^6/Q_2^6)$	0.17	(0.11)	0.21	-0.00	0.14	0.20	0.12
$\ln(Q_4^6/Q_2^6)$	0.14	(0.11)	0.19	-0.00	0.25	0.29	0.01
$\ln(Q_5^6/Q_2^6)$	-0.03	(0.11)	-0.00	-0.00	0.33	0.07	-0.01
$\ln(Q_3^7/Q_2^7)$	0.32	(0.06)	0.31	0.00	0.14	0.29	0.33
$\ln(Q_4^7/Q_2^7)$	0.46	(0.06)	0.42	0.00	0.25	0.39	0.43
$\ln(Q_5^7/Q_2^7)$	0.49	(0.06)	0.46	0.00	0.33	0.43	0.45
$\ln(Q_6^7/Q_2^7)$	0.45	(0.06)	0.48	-0.00	0.42	0.44	0.45
$exit_1$	0.36	(0.007)	0.35	0.35	0.35	0.37	0.33
$exit_2 - exit_1$	-0.16	(0.005)	-0.17	-0.18	-0.20	-0.12	-0.20
$exit_3 - exit_1$	-0.22	(0.005)	-0.22	-0.23	-0.22	-0.20	-0.23
$exit_4 - exit_1$	-0.25	(0.006)	-0.25	-0.24	-0.24	-0.25	-0.25
$exit_5 - exit_1$	-0.27	(0.006)	-0.26	-0.25	-0.25	-0.28	-0.26
$exit_6 - exit_1$	-0.27	(0.007)	-0.27	-0.25	-0.26	-0.30	-0.27

Notes: Quantity moments based on Table 4. Exit moments based on second column of Table 5. $exit_1$ refers to the intercept for the US market in this regression. 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^2	σ_η^2	$\frac{F}{\mathbb{E}(R_1)}$	ω	Fit
Baseline	0.54	0.52	6.39	0.87	0.32	0.65	1.64	0.00	0.03	0.39
$\alpha = 0$	0	0	0	0.82	0.71	1.17	0.00	0.03	0.06	13.68
$\rho = 0$	0.29	0.31	4.69	0.40	0	1.10	3.77	0.02	0.07	4.44
no AC	0.73	0.63	0	0.58	0.75	0.46	0.25	0.00	0.01	1.60
$\gamma = 1$	0.52	0.63	20.67	1	0.51	0.56	4.68	0.01	0.02	1.12

Notes: Fit indicates 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 error of the data moments on the diagonal.

Table 8: Parameters and fit: Baseline model fitted to restricted sets of moments

	Parameter									
Moments	α	δ	ϕ	γ	ρ	σ_ν^2	σ_η^2	$\frac{F}{\mathbb{E}(R_1)}$	ω	Fit
Baseline	0.54	0.52	6.39	0.87	0.32	0.65	1.64	0.00	0.03	0.39
no growth	0.71	0.35	16.19	0.81	0.68	0.48	3.59	0.00	0.07	52.99
no initial Q	0.51	0.67	12.73	0.84	0.46	0.50	4.04	0.05	0.05	15.85
no exit	0.41	0.54	23.53	1.00	0.81	0.54	0.96	0.09	0.01	11.26
no short spells	0.45	0.61	37.22	0.87	0.39	0.91	1.26	0.01	0.07	6.47

Notes: Estimates are based on matching a criterion function that sets the weight on some subsets of moments equal to zero. Fit indicates 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 error of the data moments on the diagonal. Fit takes account of all data moments and is comparable to the fit of the baseline model.

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

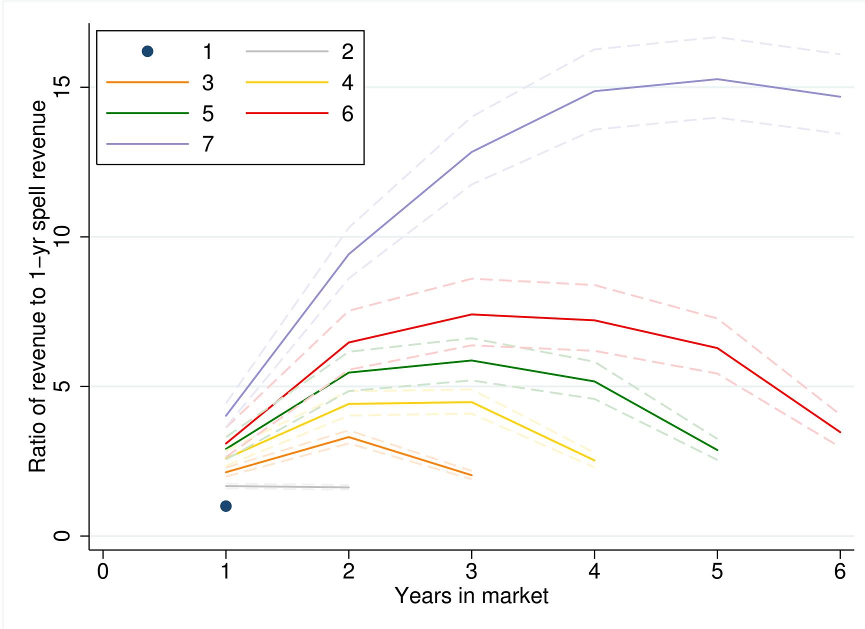

Notes: Figure shows evolution of revenue at the firm-product-market level with tenure in the market, allowing trajectories to differ with completed spell length. Trajectories are conditional on firm-product-year and market effects. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

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

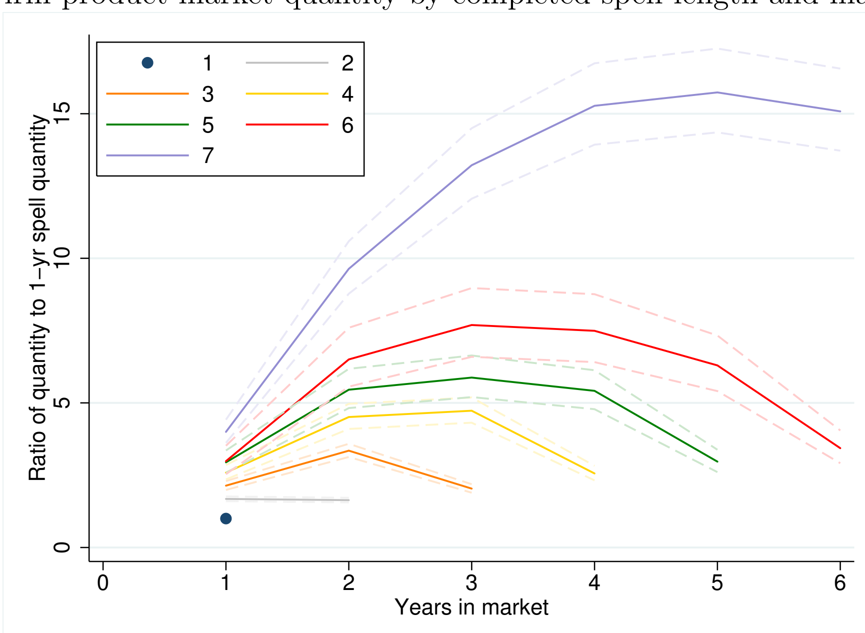

Notes: Figure shows evolution of quantities at the firm-product-market level with tenure in the market, allowing trajectories to differ with completed spell length. Trajectories are conditional on firm-product-year and market effects. 95% confidence intervals are plotted. Source: CSO and authors' calculations.

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

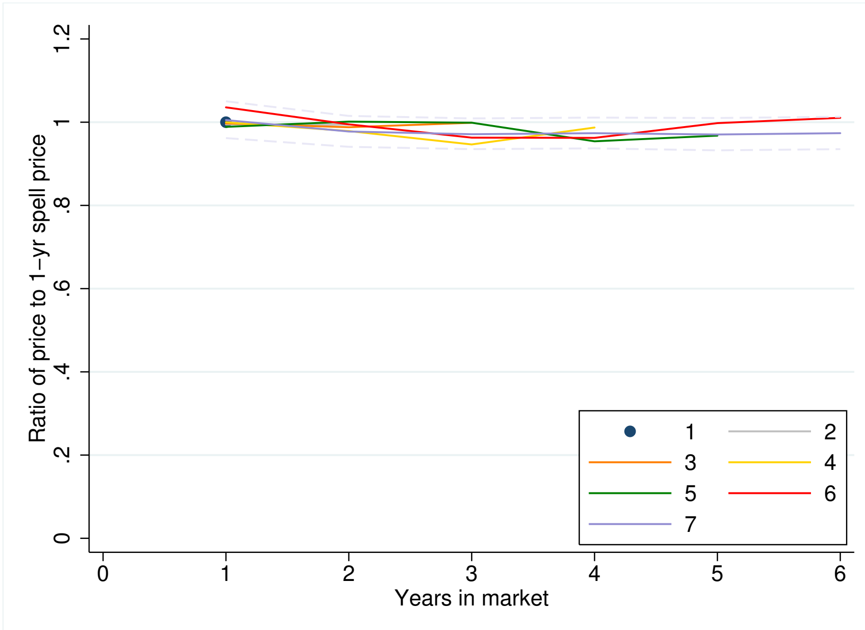

Notes: Figure shows evolution of prices at the firm-product-market level with tenure in the market, allowing trajectories to differ with completed spell length. Trajectories are conditional on firm-product-year and market effects. The 95% confidence interval for spells of 7+ years is plotted. Source: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

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

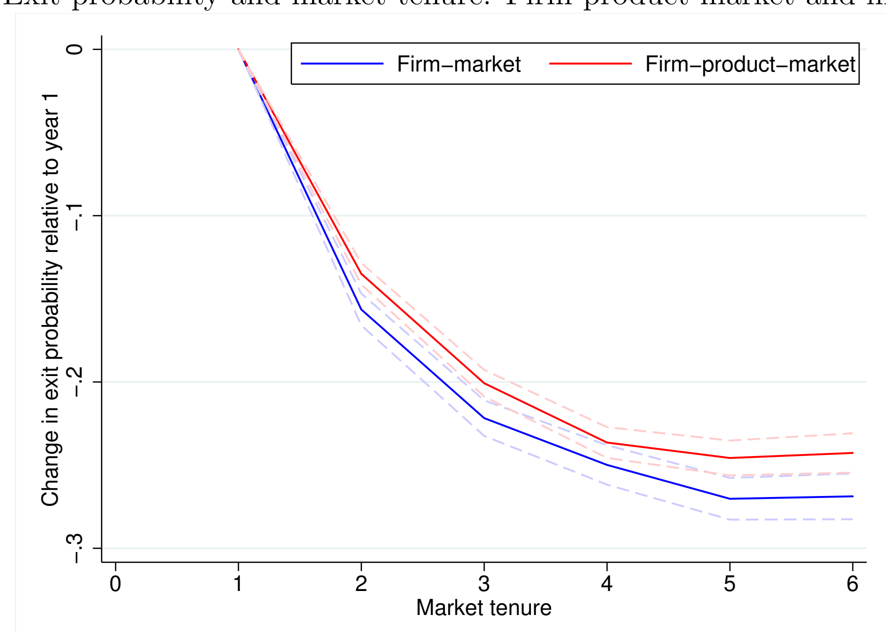

Notes: Figure shows reduction in probability of exit at the firm-market and firm-product-market levels with respect to probability of exit in the first year in a market. Trajectories are conditional on firm-year and market and firm-product-year and market effects, respectively. 95% confidence intervals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

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

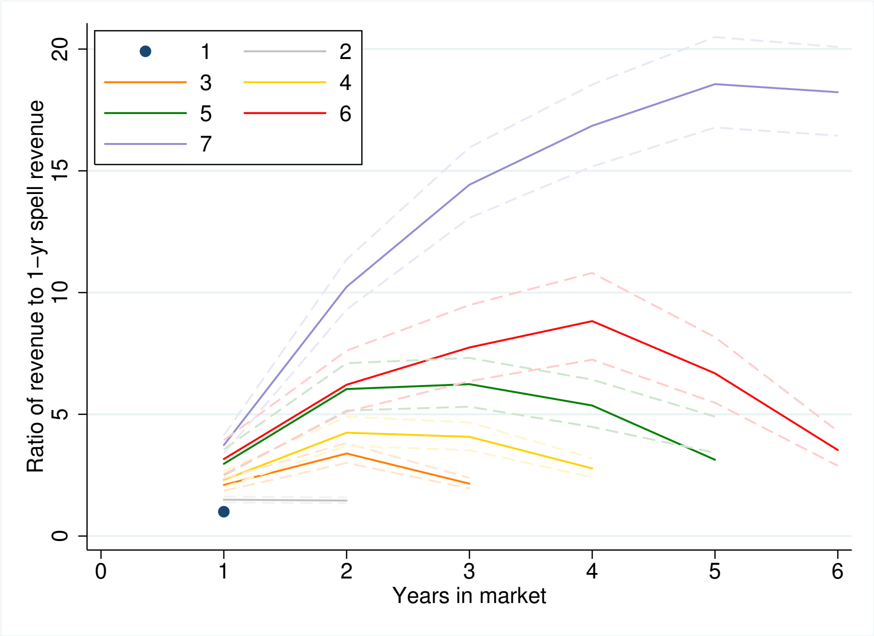

Notes: Figure shows evolution of revenue at the firm-market level with tenure in the market, allowing trajectories to differ with completed spell length. Trajectories are conditional on firm-product-year and market effects. Source: CSO and authors' calculations.

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

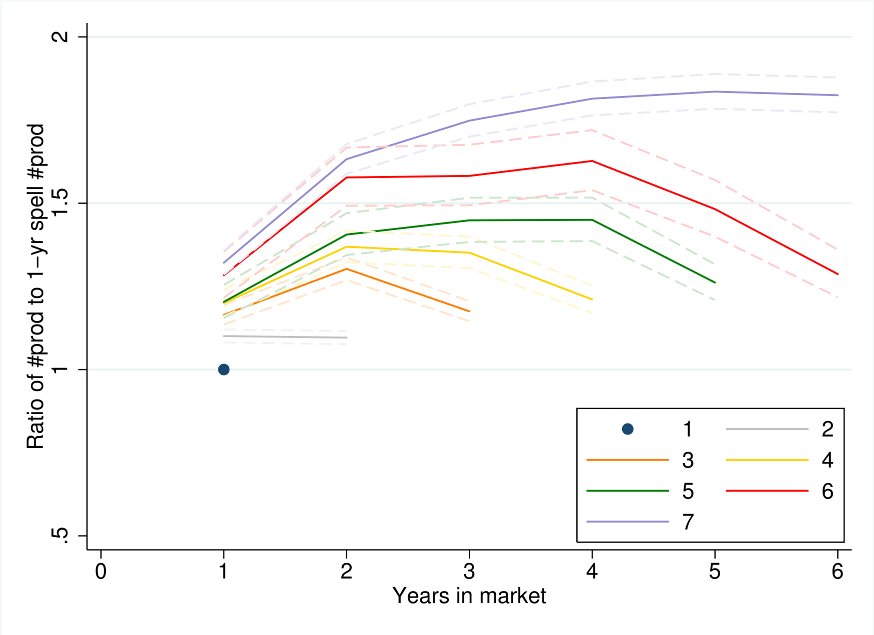

Notes: Figure shows evolution of revenue at the firm-market level with tenure in the market, allowing trajectories to differ with completed spell length. Trajectories are conditional on firm-product-year and market effects. Source: CSO and authors' calculations.

Figure 7: Model fit: Initial quantities

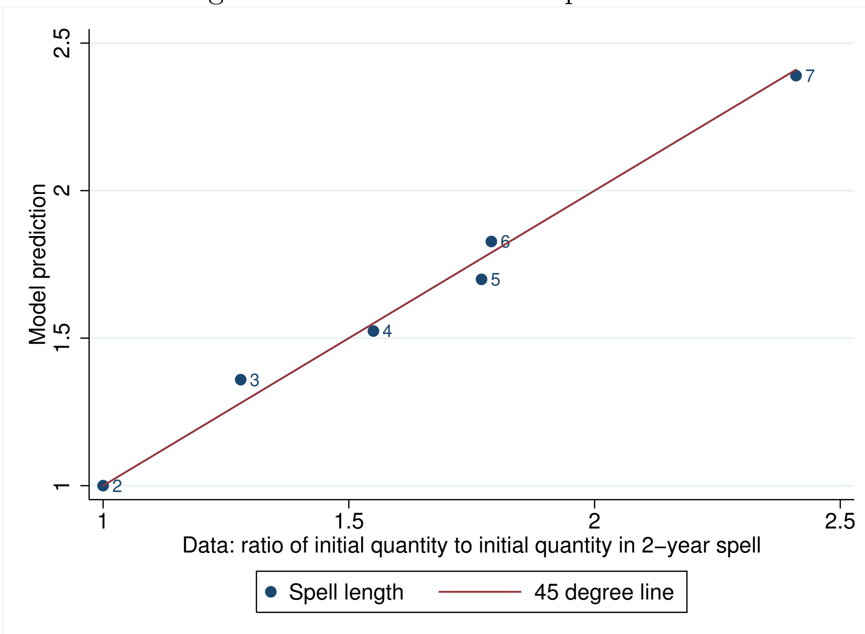

Notes: Figure shows scatter plot of data on initial quantity relative to initial quantity in two-year spell against corresponding ratio from the structural model. Source: CSO and authors' calculations.

Figure 8: Model fit: Growth of quantities

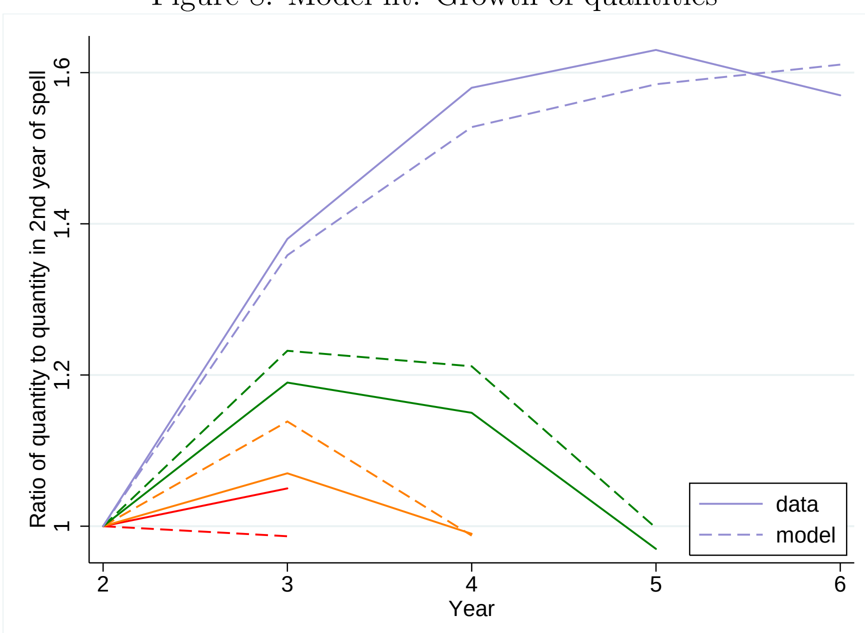

Notes: Figure shows data on evolution of quantities at the firm-product-market level with tenure by spell length, and corresponding evolution for the structural model. All quantities are expressed relative to quantity in year 2 of relevant spell. Source: CSO and authors' calculations.

Figure 9: Model fit: Exit

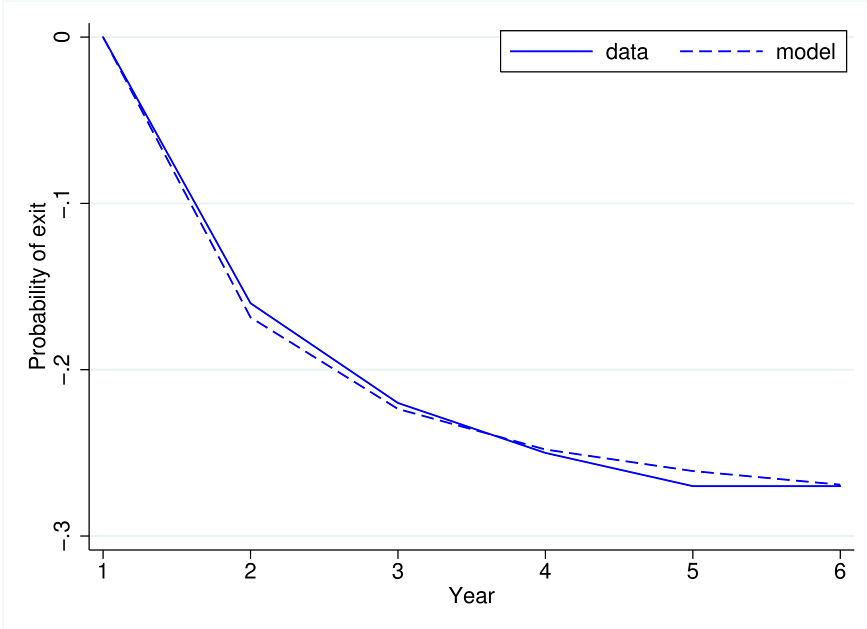

Notes: Figure shows data on reduction in probability of exit at the firm-market level relative to probability of exit in the first year in a market, and corresponding evolution for the structural model. Figure does not illustrate exit rate in year 1. Source: CSO and authors' calculations.

Figure 10: Model prediction: Advertising and marketing share of expected revenue

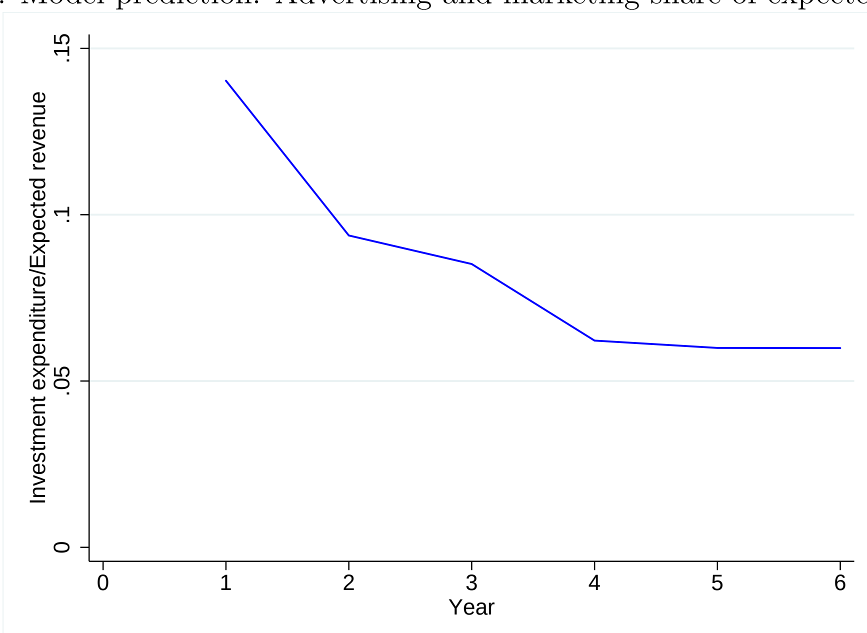

Notes: Population average of $((A + c(A, D)) / \mathbb{E}(REV))$ for export spells lasting at least seven periods, based on structural parameter estimates in Table 7. Source: Authors' calculations.

Figure 11: Simulated elasticity of aggregate exports with respect to tariffs by time horizon: Baseline model

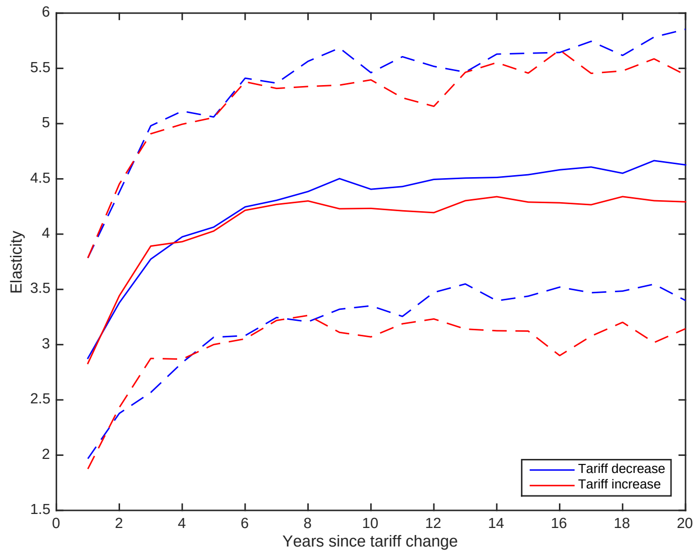

Notes: Figure shows simulated elasticity of aggregate exports with respect to tariff decreases and increases based on baseline structural model. Solid blue line is mean of $\sigma_t = \frac{\ln R_t - \ln R_0}{\ln \tau_t - \ln \tau_0}$ across 50 simulations of 5000 firms based on a tariff reduction. Dashed lines indicate the 95% confidence interval. Red lines indicate elasticity based on a tariff increase. Source: Authors' calculations.

Figure 12: Long-run versus short-run elasticities of aggregate exports with respect to tariffs: Comparing models

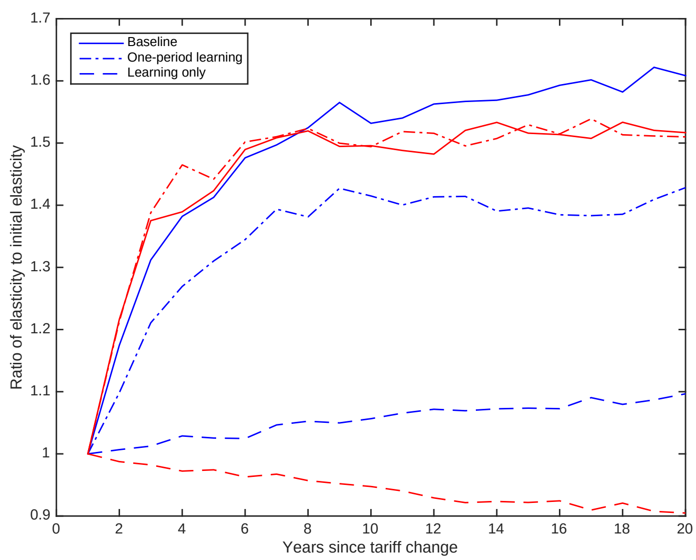

Notes: Figure shows simulated elasticity of aggregate exports with respect to tariffs relative to on-impact elasticity, (i.e., σ_t/σ_1), for three structural models: the baseline model with both costs of adjusting customer base and learning about demand; the model with costs of adjustment and only one-period learning about demand; and the model with no costs of adjustment and only learning about demand. Elasticities are calculated as $\sigma_t = \frac{\ln R_t - \ln R_0}{\ln \tau_t - \ln \tau_0}$ across 50 simulations of 5000 firm. Blue lines indicate elasticities based on tariff reductions. Red lines indicate elasticities based on tariff increases. Source: Authors' calculations.