How Exporters Grow*

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

March 2022[¶]

Abstract

We show that in successful episodes of export market entry, there are statistically and economically significant post-entry dynamics of quantities, but no post-entry dynamics of markups. We structurally estimate a model of customer base accumulation to match these moments. In our model, firms can increase customer base through spending on marketing and advertising, or by temporarily charging low markups. Our estimates suggest that firms use marketing and advertising, but not dynamic pricing to acquire new customers, and that selling expenses associated with building and maintaining customer base are considerable.

1 Introduction

Recent research makes customer base central to the analysis of firm dynamics, business cycles, and international trade.[1] Although much has been learned, there is as yet no consensus on

*Appendix available at www.doireann.com. This work makes use of data from the Central Statistics Office, Ireland, which is CSO copyright. The possibility for controlled access to confidential micro data sets on the premises of the CSO is provided for in the Statistics Act 1993. The use of CSO data in this work does not imply the endorsement of the CSO in relation to the interpretation or analysis of the data. This work uses research data sets that may not exactly reproduce statistical aggregates published by the CSO. We thank the staff of the CSO for making this project possible. Expert research assistance was provided by Adrian Corcoran, Matt Shapiro and Anthony Priolo. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. Yaniv Yedid-Levi is grateful for financial support from the Social Sciences and Humanities Research Council of Canada. We thank Manuel Amador, Costas Arkolakis, Tim Kehoe, Kim Ruhl, James Tybout, Daniel Xu, and participants in the 2015 and 2016 NBER Summer Institute 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, doireann.fitzgerald@gmail.com

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

§Tiomkin School of Economics, Reichman University (IDC Herzliya), yaniv.yl1@gmail.com

¶First draft: July 2015.

[1]See, among others, Arkolakis (2016), Drozd and Nosal (2012a,b) Eaton et al. (2011), Eaton et al. (2014), Foster et al. (2008, 2016), Gilchrist et al. (2017), Gourio and Rudanko (2014), Hottman et al. (2016), Ravn et al. (2006), Ruhl and Willis (2017).

what firms do to expand customer base after entry into a market. There are two competing theories of customer base accumulation in the literature: (1) firms acquire customer base through engaging in non-price activities such as marketing and advertising (we refer to this as the "marketing and advertising model"), and (2) customer base is a function of past sales, so firms expand in a market by first charging low markups to shift out demand, then gradually increasing markups as customer base rises (we refer to this as the "customer markets model").

We use customs data for Ireland for the period 1996-2009, and a combination of reduced form and structural estimation to confirm the role of customer base in export growth, and to distinguish between these two competing theories. Customs data are well suited to this exercise, because we observe quantities and prices for the same firm selling the same product in multiple segmented markets. We show that in the years following successful entry by a firm into a new market, export quantities grow steadily in the new market relative to old markets, while prices in the new relative to old markets are flat. Assuming that marginal cost is the same for all markets the firm serves, this is consistent with gradual accumulation of customer base in new markets through marketing and advertising, but not through intertemporal distortion of markups.

To test this hypothesis more formally, we estimate a model of the exporting firm's problem, where we nest the two theories of customer base accumulation. In our model, customer base shifts demand conditional on prices, but does not affect the price elasticity of demand. Customer base does not fully depreciate between one period and the next. The firm decides whether to participate in a given export market, knowing that it can invest in future customer base through spending on marketing and advertising (which is subject to an adjustment cost), or by depressing markups below their statically optimal level in order to increase future sales. The relative effectiveness of these two types of investment is governed by a parameter, which we estimate.

We estimate the model using simulated method of moments. The moments we match are the post-entry behavior of normalized quantities and markups, as well the exit hazard and the export entry rate. The estimated model is constrained to match a trade elasticity of 3, consistent with the literature, and with evidence provided by Fitzgerald and Haller (2018) based on the Irish customs data we use. Our estimates imply that markup distortion is ineffective in adding to customer base, so firms use only marketing and advertising to invest.

An advantage to structurally estimating the model is that we can quantify selling expenses as a fraction of revenue. They are highest on entry, and decline as the firm approaches its steady state customer base in a market. We estimate that selling expenses are 18.5% of revenue in the latter years of export spells which last 7 or more years. This is considerable, though within the range of what the empirical literature on this issue finds.[2]

One reason we care about how firms acquire customer base is because it matters for measurement, especially of productivity. If selling expenses are substantial (as we find), and if they are treated as a production cost rather than as investment, both physical productivity and revenue productivity will be systematically mismeasured. The extent of mismeasurement will differ for growing versus stagnant firms, and for periods when firms expand across export markets versus periods where they do not. In contrast, if firms build customer base by distorting markups, physical productivity is not subject to mismeasurement, although revenue productivity is.[3]

In an application of our estimated structural model, we highlight the impact of selling expenses on the potential for productivity mismeasurement. We simulate a trade liberalization episode assuming that true firm level productivity is constant, and show that measured productivity may appear to be depressed immediately after export entry due to high selling expenses in the entry year. But as new exporters converge to their steady state export intensity, measured productivity may actually increase above the pre-liberalization baseline. This is because the steady state customer base investment-to-sales ratio is lower in export markets than in the home market due to the lower probability of survival in the export market. It is possible that this effect could contribute at the margin to the appearance of within-firm productivity gains from exporting. This suggests caution in interpreting findings in the literature (such as De Loecker (2007) and Lileeva and Trefler (2010)) that there are within-firm productivity gains associated with export entry due to trade liberalizations.[4]

Our contribution is related to a number of other papers that investigate quantity and price dynamics using manufacturing census and customs data. This includes Foster et al. (2008) who use plant census data for the US and focus on implications for productivity measurement, and Bastos et al. (2018), Berman et al. (2019) and Piveteau (2021) who use customs data for Portugal and France and focus on export dynamics. In contrast to other work in this area, we use a comprehensive set of quantity, price and exit moments to structurally estimate competing models of customer base accumulation, to ask which one

[2]Using Compustat data from the U.S. for 1971-2006, Gourio and Rudanko (2014) find that firms for which "Selling, General, and Administrative Expenses" (SG&A) are above the industry median have an average share of SG&A in sales of 27%, while firms for which SG&A are below the industry median have an average share of SG&A in sales of 17%.

[3]This latter point is made by Foster et al. (2008, 2016).

[4]De Loecker (2007) and Lileeva and Trefler (2010) cannot distinguish between selling expenses and production cost in their data.

the data favors, and to obtain estimates of selling costs.

Our work is also related to a literature in macroeconomics that posits that current sales affect future demand. Models with this feature are used extensively in work that focuses on flexible-price markup-based explanations for the countercyclicality of the labor wedge in business cycles. The idea is that booms are times when there are many potential new customers, and to attract these customers, firms choose low markups (e.g. e.g. Bils (1989), Ravn et al (2006) and Gilchrist (2017)). While we cannot rule out that firms use markups to attract customers at a business cycle frequency, our results do not provide support for this mechanism over a longer horizon.

Beyond the specific areas of productivity measurement and markup cyclicality, our findings have potential implications for the many areas of macroeconomics and international economics where models with demand and customer base are used. Gourio and Rudanko show that sluggish adjustment of customer base matters for firm responses to shocks, and for the relationship between investment and Tobin's Q. Drozd and Nosal (2012b) investigate the performance of the customer markets model in matching international business cycle comovements, and find it does poorly. In a direct application of our findings, Fitzgerald et al. (2019) show that the marketing and advertising model we estimate here has the potential to rationalize the very different responses of exports to movements in exchange rates and changes in tariffs that we see in the data.

The paper is organized as follows. In Section 2, we lay out our model. In Section 3 we describe our data. In Section 4 we describe our reduced form empirical strategy. In Section 5, we describe our reduced form results. In Section 6 we structurally estimate the model to match these results. In Section 7 we illustrate the implications of our findings for measured productivity around a trade liberalization. The final section concludes.

2 A model of customer base accumulation

Our model generalizes the exporter problem in Arkolakis (2010). Arkolakis presents a static model in the tradition of informative advertising, where a firm can shift current demand in a market by engaging in marketing and advertising in that market. Marketing and advertising effort has decreasing returns, and does not affect the price elasticity of demand.

We generalize in two ways. First, we make the model dynamic through less than full depreciation of customer base between one period and the next. Second, we assume that future customer base may be increasing in current sales as well as in current marketing and advertising effort. In this, we follow a long tradition of "customer markets" models in macroeconomics (e.g. Bils (1989), Ravn et al (2006)), as well as more recent work by Foster, Haltiwanger and Syverson (2016) on firm dynamics, whose specification we adopt.[5]

When future customer base is increasing in current sales, and customer base is below its steady state level, the firm has an incentive to set a below-steady-state markup. In doing so it trades off lower profit today in return for higher customer base and higher profits in the future. The price elasticity of demand is a key parameter governing this tradeoff: the more price elastic is demand, the smaller is the reduction in today's markup necessary to generate a given shift in future demand. In this model the transition to steady state customer base is slow by assumption.

We also allow for adjustment costs of advertising and marketing investment, so transitions may be slow even in the limiting case where future customer base is unaffected by lagged sales.[6]

We nest this model of demand within a parsimonious model of export entry and exit in mutiple markets due to stochastic fixed and sunk costs of exporting.

More formally, firms are indexed by i, and markets are indexed by k. Markets are segmented, so the firm is able to price discriminate. Only within-market marketing and advertising effort and sales affect future within-market customer base. The only channel through which a firm's decisions across different markets are linked is through its common exogenous marginal cost of production, C_t^i.[7] In each export market k, the firm faces iid sunk (S_t^{ik}) and fixed (F_t^{ik}) costs of participation. Sunk costs lead to selection on entry, while fixed costs lead to selection on both entry and exit.

Conditional on participation, demand is isoelastic:

$$Q_t^{ik} = Q_t^k \left(\frac{P_t^{ik}}{P_t^k} \right)^{-\theta} \left(D_t^{ik} \right)^{\alpha} \exp\left(\nu_t^{ik} \right). \tag{1}$$

Here, Q_t^k is aggregate demand in market k and P_t^k is an index of competitors' prices. Going forward, we write $Y_t^k = Q_t^k \left(P_t^k \right)^{\theta}$, and refer to Y_t^k as market size.[8] Under monopolistic competition, market size is exogenous to the firm. Demand also depends on the exogenous

[5]Our preferred interpretation of this specification is that new customers may learn about a firm by observing the purchases of current customers. In this case, there are no strategic interactions between a firm and its current customers.

[6]This aspect of our setup is related to Drozd and Nosal (2012a) who build a two-country dynamic general equilibrium model with firms which accumulate customer base through marketing expenditures.

[7]In Appendix A, we lay out a model augmented with firm-level heterogeneity in quality as well as cost.

[8]In Appendix A we show that iceberg trade costs and tariffs enter the firm's problem in the same way as market size.

idiosyncratic demand shifter ν_t^{ik}.

The firm affects the quantity sold in market k by choosing the price, P_t^{ik}, and by taking actions which affect customer base, D_t^{ik}. If $\alpha \in (0,1)$, demand is increasing in customer base, but at a diminishing rate. If $\alpha = 0$, demand does not depend on customer base. Note that customer base shifts demand, but does not affect the price elasticity of demand.

Let $X_t^{ik} = \{0,1\}$ be an indicator for participation. Firms entering market k start with customer base $\underline{D}^k$. Customer base in market k accumulates according to:

$$D_{t+1}^{ik} = \left(1 - X_t^{ik}\right) \underline{D}^k + X_t^{ik} \left(\left(1-\delta\right) D_t^{ik} + \psi A_t^{ik} + \left(1-\psi\right) P_t^{ik} Q_t^{ik}\right) \tag{2}$$

Customer base depreciates at rate δ conditional on continued participation. There is full depreciation on exit. The parameter ψ governs the relative effectiveness of marketing and advertising and current sales in adding to future customer base. ψA_t^{ik} is the increment to customer base at date $t+1$ due to marketing and advertising at date t, while $\left(1-\psi\right) P_t^{ik} Q_t^{ik}$ is the increment to customer base at date $t+1$ due to sales at date t. The cost of marketing and advertising is given by $c\left(D_t^{ik}, A_t^{ik}\right)$, where $c\left(\cdot, 0\right) = 0$, and $c\left(\cdot, \cdot\right)$ is differentiable in both arguments, with $c_A > 0$, $c_{AA} \geq 0$ and $c_D \leq 0$. Since customer base is intangible, it is natural to assume irreversibility, i.e. $A_t^{ik} \geq 0$.

There are three sources of potentially persistent heterogeneity in the model: C_t^i, Y_t^k, and ν_t^{ik}. The firm knows the processes for these variables, as well as the processes for fixed and sunk costs, and observes current realizations of all shocks before making decisions.[9]

Let $Z_t^{ik} = \left\{S_t^{ik}, F_t^{ik}, C_t^i, Y_t^k, \nu_t^{ik}\right\}$. Flow profit conditional on participation at t is:

$$\pi\left(Z_t^{ik}, X_{t-1}^{ik}, D_t^{ik}, A_t^{ik}, P_t^{ik}\right) = \left(P_t^{ik} - C_t^i\right) Y_t^k \left(P_t^{ik}\right)^{-\theta} \left(D_t^{ik}\right)^{\alpha} \exp\left(\nu_t^{ik}\right) \tag{3}$$

$$-c\left(D_t^{ik}, A_t^{ik}\right) - F_t^{ik} - \left(1 - X_{t-1}^{ik}\right) S_t^{ik}$$

Current participation (X_t^{ik}), marketing and advertising effort (A_t^{ik}), and current prices (P_t^{ik})

[9] Berman et al (2019) propose that learning about idiosyncratic demand may induce post-entry export dynamics. In Appendix C we describe a model where learning about ν_t^{ik} rather than customer base accumulation induces dynamics in quantities and prices. We structurally estimate this model (results in Appendix I), and it provides a very poor fit to the data.

affect future as well as current profits. The Bellman equation for the firm's problem is:

$$V\left(Z_t^{ik}, X_{t-1}^{ik}, D_t^{ik}\right) = \max_{\substack{X_t^{ik} \in \{0,1\} \\ A_t^{ik} > 0 \\ P_t^{ik} > 0}} \left\{ \begin{array}{c} X_t^{ik}\pi\left(Z_t^{ik}, X_{t-1}^{ik}, D_t^{ik}, A_t^{ik}, P_t^{ik}\right) + \\ \\ \beta\mathbb{E}\left\{V\left(Z_{t+1}^{ik}, X_t^{ik}, D_{t+1}^{ik}\right)|Z_t^{ik}\right\} \end{array} \right\}$$

subject to the accumulation of customer base.

2.1 Model predictions

If we make assumptions about functional forms and the statistical processes for the exogenous variables, we can prove some propositions for the polar cases where $\psi = 1$ (marketing and advertising model) and $\psi = 0$ (customer markets model). We prove three propositions showing that markups behave differently in these two polar cases. This suggests that markup behavior can be used to distinguish between them empirically. We prove a fourth proposition to provide support for our approach to dealing with unobserved heterogeneity in idiosyncratic demand in our empirical strategy.

Assume that marginal cost is constant for each firm ($C_t^i = C^i$), market size is constant for each market ($Y_t^k = Y^k$), and ν_t^{ik} remains fixed within an export spell (i.e. an episode of continuous participation by a firm in a particular market). On exit, the firm loses its draw of ν_t^{ik}, and receives a new draw of ν_t^{ik} in every subsequent period of non-participation. Assume that F_t^{ik} can take only two values, zero and infinity (so exit is exogenous), and the adjustment cost function for marketing and advertising takes the form $c(D, A) = A + \phi(A^2/D)$.[10] Assuming the resulting value functions are differentiable and concave, we prove the following:

Proposition 1. When $\psi = 1$ (marketing and advertising model), the markup is *independent* of C^i, Y^k and ν_t^{ik}.

Proof See Appendix B.

The intuition for Proposition 1 is very straightforward. When $\psi = 1$, current sales do not affect future customer base. The optimal price is a static decision, and is given by the standard CES markup over marginal cost, $\theta/(\theta - 1)$.[11]

[10] This is the functional form we use in the structural estimation.

[11] This is obviously true under much more general assumptions than those we make here.

Proposition 2. When $\psi = 0$ (customer markets model), the markup on entry is *increasing* in C^i and *decreasing* in Y^k and ν_t^{ik}.

Proof See Appendix B.

Proposition 3. When $\psi = 0$ (customer markets model), if customer base on entry is below steady state customer base, then (a) the markup converges to the steady state markup from below, and (b) growth in the markup on entry (i.e. growth between the first and second periods of participation) is *decreasing* in C^i, and *increasing* in Y^k and ν_t^{ik}.

Proof See Appendix B.

The intuition for Propositions 2 and 3 is that low markups are an investment in customer base. Steady state customer base is decreasing in marginal cost and increasing in market size and idiosyncratic demand. Holding fixed customer base on entry, the greater is steady state customer base, the greater the incentive to invest.

We verify in model simulations that the behavior of markups in Propositions 1-3 carries over to the case of endogenous exit. In Appendix B, we also prove a series of propositions characterizing the behavior of quantities, and show that in both polar cases ($\psi = 1$ and $\psi = 0$), quantities on entry and quantity growth depend on costs, market size, and idiosyncratic demand.

Our last proposition allows us to use duration i.e. how long a firm survives in a particular market, as a proxy for unobserved idiosyncratic demand. The following holds under endogenous exit, both when $\psi = 1$ and when $\psi = 0$:

Proposition 4. Fix C^i, Y^k, and D_t^{ik}. Let $\bar{\nu}^{ik}$ be the current draw of ν_t^{ik}. Then the probability of survival is increasing in $\bar{\nu}^{ik}$.

Proof See Appendix B.

The intuition is straightforward. On exit, the firm loses its current draw of the permanent component of idiosyncratic demand. The value of exit is therefore independent of this draw, while the value of continued participation is increasing in this draw. Hence, the probability that the firm will continue to participate is increasing in its current draw of the permanent component of idiosyncratic demand. This implies that duration, which is an increasing function of the survival probability, is also increasing in $\bar{\nu}^{ik}$. We conjecture that the probability of survival is also decreasing in costs and increasing in market size, as long as customer base is above customer base on entry (i.e. $D > \underline{D}$). In Appendix B.3, we verify that this is the case numerically.[12]

3 Data

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 customs records for merchandise exports. The data are described in detail in Appendix D.

3.1 Census of Industrial Production

The CIP is an annual census of manufacturing, mining, and utilities. Firms with three or more employees are required to file returns.[13] 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 total revenue, employment, and country of ownership.

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.

3.2 Customs records

Our main source of data is customs records of Irish merchandise exports for the years 1996-2014. The f.o.b. value (euros) and quantity (tonnes)[14] 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. Appendix D provides additional information on this match and on other details of the data. We make use of all matched records.

[12] This is tricky to prove formally because the value of participation and the value of staying out of the market are both decreasing in cost and increasing in market size.

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

[14] The value is always available, but quantity is missing for about 10% of export records. For a limited set of observations, an additional quantity measure (besides tonnes) is available. We make use of this in a robustness check.

An important feature of the customs data is that the eight-digit CN classification system changes every year. In order to make time-series comparisions, 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 et al. (2012).[15] The breakdown of exports by HS 2-digit classification over the sample period is reported in Appendix D (Table 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 for this analysis 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. When we do so, we do not condition on 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 panel data on the value and quantity of exports at the firm-product-market level, where the product is defined at the eight-digit (concorded) level, and the market refers to the destination country. We use this to construct a price (unit value) by dividing value by quantity, where available. In aggregate trade statistics, unit value data at the product level are notoriously noisy. However, conditioning on the exporting firm as well as the product considerably reduces this noise.[16]

3.3 Summary statistics

Table 1 shows summary statistics on exporting behavior in our data. Export participation is high, export intensity conditional on participation is high, and more than half of exporters export to more than one market (we observe 141 distinct export markets over the course of the panel). These facts are typical of small open European economies (see ISGEP (2008)).

There is a good deal of churn in export participation at the firm-market level, and entry and exit are not synchronized across different export markets within a given firm.[17] This is illustrated in Table 2, which reports the distribution of exporter-level changes 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 is a lower bound on churn, as firms may keep the total number of export markets constant, while switching between markets. This churn is consistent with stochastic fixed and sunk costs of the type we include in our model. Fitzgerald and Haller (2018) document additional facts about entry and exit in these data.

[15] We examine robustness to conditioning on products for which the concordance is 1-1 for all pairs of years in the sample.

[16] We check that our results are robust to dropping unit value outliers.

[17] This is consistent with Lawless (2009), who uses a different data set for a selected sample of Irish firms.

4 Reduced form empirical strategy

There are three goals of our reduced form analysis. First, we want to confirm that customer base accumulation is an important source of post-entry dynamics in our data. Second, we want to see whether there are post-entry dynamics in markups. Third, the moments we construct will be used as inputs into our simulated method of moments estimation in Section 6. Before describing exactly what we do, we define some useful variables, and give an overview of how our approach achieves these three goals.

4.1 Definitions and measurement

As explained in Section 3.2, an observation in our data is a firm-product-market-year.

An *export spell* is a continuous episode of market participation, i.e. an episode in which there are positive exports for that firm-product-market in every consecutive year. Spell entry is censored if we see exports in the first year of the sample.[18] Spell exit is censored if we see exports in the last year of the sample. 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.[19]

Export *age* is set equal to 1 in the first year of exporting after not exporting in the previous year. Age is incremented by 1 in each subsequent year of continuous participation. If we see no exports for the firm-product-market for some year, age is reset to 1 in the first subsequent year of participation. Age is censored if spell entry is censored.

Duration is defined as age in the year prior to exit. Duration is observed only if an export spell is neither left- nor right-censored. However for right-censored spells, we know that duration is at least as great as age at the end of the sample. By top-coding duration, we can make use of this information. As a baseline, we topcode both age and duration at 7 years.[20]

4.2 Overview

To confirm that customer base accumulation is an important source of post-entry quantity dynamics, we need to isolate variation in quantities that is not driven by changes in marginal cost. To examine the dynamics of markups, we must similarly purge changes in marginal

[18] Our sample starts in 1996, but because of issues with the match between customs data and firms in that year, we also consider spell entry to be censored if it takes place in 1997.

[19] In our baseline analysis we treat these reentry spells the same as first spells.

[20] Using our full panel of customs data, which lasts for 19 years, we show that our key results are robust to top-coding at levels up to 10 years.

cost from prices. Under the assumption that marginal cost is the same for all markets served by a particular firm-product pair, we achieve both of these aims by differencing quantities and prices across markets within a firm-product i.e. by regressing logs of these variables on firm-product-year fixed effects.

We are not interested in dynamics that are due purely to exogenous changes in market size. So for both quantities and prices, we control for time-varying market size by differencing quantities and prices across firms selling the same product in the same market, i.e. by regressing logs of these variables on product-market-year fixed effects. This is valid under the assumption of monopolistic competition, where all market participants face the same market size.

We are concerned that selection on persistent unobserved idiosyncratic demand may give the appearance of a relationship between relative quantities and age, and relative markups and age, even when there no true underlying dynamics. The underlying problem is that firm-product-market spells which survive for a long time are likely to have higher idiosyncratic demand than the average entrant. Regressing, say, quantities on age may therefore lead us to infer dynamics even when there are none. So motivated by Proposition 4, we control for persistent idiosyncratic demand by conditioning on duration, i.e. how long a firm-product pair survives in a particular market.[21]

We then identify dynamics in relative quantities and markups by examining how the residual variation in quantities and prices after these controls evolves with a firm-product pair's age in a market.

If dynamics are driven by investment in customer base, *growth* in relative quantities and markups will vary with the factors that drive the return to investment, i.e. marginal cost, market size, and idiosyncratic demand. In order to allow dynamics to vary with these factors, we interact age with proxies motivated by Proposition 4 and related conjectures. We can then get an estimate of the dynamic behavior of relative quantities and markups for fixed marginal cost, market size and idiosyncratic demand by evaluating at fixed values of these proxies.

We proxy marginal cost using the number of distinct markets a firm exports to over the entire sample period. We proxy market size using the number of distinct firms which export to a market over the entire sample period. As already noted, we proxy idiosyncratic demand using duration. We verify in model simulations that these proxies are correlated

[21]This follows Abraham and Farber (1987), who use duration to proxy for unobserved heterogeneity in the quality of worker-firm matches in order to obtain an unbiased estimate of returns to job tenure.

with marginal cost, market size, and idiosyncratic demand as hypothesized.[22]

In addition to examining dynamics in relative quantities and markups, we also characterize the behavior of the exit hazard, verifying that it is downward-sloping conditional on firm- and market-specific factors, consistent with the hypothesis that there is selection on idiosyncratic demand. This supports our use of duration as a proxy for idiosyncratic demand, and will be useful for quantifying heterogeneity in idiosyncratic demand in the structural estimation.

4.3 Implementation

Let w_t^{ijk} denote log quantity, or log price at the firm-product-market level, where i indexes firms, j indexes products, and k indexes markets. Let c_t^{ij} be a a set of firm-product-year fixed effects. Let d_t^{jk} be a set of product-market-year fixed effects. Let $\mathbf{a}_t^{ijk}$ be a vector of indicator variables for firm i's age in market k with product j. Let $\mathbf{l}_t^{ijk}$ be a vector of indicators for the duration of the relevant spell. This indicator does not vary within a spell. Let $\mathbf{cens}^{ijk}$ be a vector of indicators for spells that are left-censored ($cens_l^{ijk}$), and right-censored but not left-censored ($cens_r^{ijk}$). Then let:

$$\mathbf{s}_t^{ijk} = \begin{bmatrix} \mathbf{l}_t^{ijk} \otimes \mathbf{a}_t^{ijk} \\ \mathbf{cens}^{ijk} \end{bmatrix}$$

where the symbol $\otimes$ indicates the Kronecker product. We do not observe age of greater than l for a spell that lasts l years, so the redundant interactions are dropped. We code $\mathbf{s}_t^{ijk}$ such that the dependent variable (log quantity or price) is normalized to 0 for 1-year export spells.

Let m^i and f^k be the number of markets per firm and the number of firms per market as described above. Our baseline estimating equation is then:

$$w_t^{ijk} = c_t^{ij} + d_t^{jk} + \boldsymbol{\beta}' \left(\mathbf{s}_t^{ijk} \otimes \begin{pmatrix} 1 \\ m^i \\ f^k \end{pmatrix} \right) + \epsilon_t^{ijk}. \tag{4}$$

We include all observations with positive exports in estimating this equation.

Because $\mathbf{s}_t^{ijk}$ includes all possible combinations of age and duration, our specification is fully non-parametric with respect to age and duration, and semi-parametric with respect to

[22] See Appendix B.3 for details.

m^i and f^k. By controlling nonparametrically for duration and age, we allow for nonlinearities induced by entry or exit part-way through a year. In interpreting our results, we take these effects into account, and we explicitly incorporate them into our structural model.[23]

We can then use the estimated coefficients $\boldsymbol{\beta}$ to evaluate

$$E\left(w_t^{ijk} - c_t^{ij} - d_t^{jk} | a, l, m^i, f^k\right) - E\left(w_t^{ijk} - c_t^{ij} - d_t^{jk} | 1, 1, m^i, f^k\right)$$

for values of $\left\{a, l, m^i, f^k\right\}$. Assuming our proxies span the underlying unobserved heterogeneity, this gives us estimates of the dynamics of relative quantities and relative markups conditional on marginal cost, market size, and idiosyncratic demand.

To characterize the conditional exit hazard at the firm-market level, we estimate the following linear probability model:

$$\Pr\left[X_{t+1}^{ik} = 0 | X_t^{ik} = 1\right] = c_t^i + d_t^k + \boldsymbol{\beta}' \left(\begin{pmatrix} \mathbf{a}_t^{ik} \\ cens_l^{ik} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ m^i \\ f^k \end{pmatrix} \right) + \epsilon_t^{ik} \quad (5)$$

We include all observations where exit is not censored by the end of the sample. The terms c_t^i, d_t^k, $\mathbf{a}_t^{ik}$ and $cens_l^{ik}$ are defined as above. We code $\mathbf{a}_t^{ik}$ and $cens_l^{ik}$ such that the exit hazard is normalized to zero for for the first year of an export spell. Based on our estimates of this expression, we can trace out

$$E\left(\Pr\left[X_{t+1}^{ik} = 0 | X_t^{ik} = 1\right] - c_t^i - d_t^k | a, m^i, f^k\right) - E\left(\Pr\left[X_{t+1}^{ik} = 0 | X_t^{ik} = 1\right] - c_t^i - d_t^k | 1, m^i, f^k\right)$$

for values of $\left\{a, m^i, f^k\right\}$. This gives us estimates of the exit hazard conditional on marginal cost and market size, providing insight into selection on idiosyncratic demand. We also estimate this equation at the firm-product-market level.

Because they are useful as target moments in the structural estimation, we also regress entry rates and one-year exit rates at the firm-market level on m^i and f^k. This allows us to calculate entry and 1-year exit rates conditional on values of m^i and f^k.

[23] Some authors who have access to high-frequency data correct for part-year effects by dating the beginning of an export spell e.g. from its first month, and aggregating to an annual frequency starting each "year" in the entry month. We do not have the data to do this, and it would in any case invalidate the use of calendar year fixed effects to control for marginal cost and market size.

5 Reduced form results

We first confirm that our proxies for marginal cost, market size and idiosyncratic demand have desirable properties.

Table 3 reports the correlation of the number of markets per firm (m^i) with log average employment, average sales per worker and average revenue-based TFP (averages are taken over all the periods the firm is present in the CIP). It also reports the correlations between number of firms per market (f^k), market k's average share in world GDP over the sample period, and the bilateral distance between Ireland and market k.[24] We find that m^i is positively correlated with log employment, sales per worker and TFP, giving us confidence that it captures an important dimension of the firm's underlying cost advantage. Meanwhile, f^k is strongly positively correlated with market k's share in world GDP and negatively correlated with market k's distance from Ireland, giving us confidence that it summarizes the attractiveness of market k to Irish firms.

Table 3 also reports the correlations of the firm- and market-level proxies with each other. At the export spell level, m^i and f^k are negatively correlated, consistent with "bad" firms exporting only to "good" markets, and conversely, only "good" firms exporting to "bad" markets. This is in line with what we expect based on our model.

In Table 4, we report the distribution of duration across export spells and export observations. Short-duration spells account for a large fraction of spells, but a substantially smaller fraction of export observations. In Table 5, we regress duration on m^i and f^k. As we expect, the coefficients on m^i and f^k are positive and strongly significant. However the R-squared of the regression is less than 1%. Conditional on entry, firm- and market-specific heterogeneity does not account for much of the variation in duration. Unlike m^i and f^k, we do not have alternative proxies for idiosyncratic demand with which to compare duration. But the fact that there is a good deal of residual variation in duration conditional on m^i and f^k is consistent with an important role for selection on idiosyncratic demand.

5.1 Baseline results

We now present the results from estimating our baseline specification, equation (4). We report the results in the form of figures which show the (exponential of the) fitted values of the dependent variables, evaluated at the mean values of our proxies for costs and market size, i.e. $\left\{\bar{m}^i, \bar{f}^k\right\}$, for all possible combinations of age and duration. These fitted values are

[24] See Appendix D for more information on data sources and construction.

graphed against age. The omitted category in all regressions is export spells which last one year. The log of the dependent variable for these spells is therefore normalized to 0, and the exponential is normalized to 1. Figure 1 shows the results for quantities, and Figure 2 shows the results for prices. Full results from estimating our baseline specification are reported in Appendix Tables 9-11.[25]

Figures 1 and 2 therefore illustrate the average trajectories of relative quantities and relative prices (i.e. quantities and prices net of the relevant the firm-product-year and market-product-year fixed effects) for export spells of different duration, for a firm of type $\bar{m}^i$ in a market of type $\bar{f}^k$. Note that these means are taken across all firm-product-market-year observations. So $\bar{m}^i$ is at the 96th percentile in terms of exporting firms, at the 57th percentile in terms of export spells, and at the 55th percentile in terms of export observations. Meanwhile $\bar{f}^k$ is at the 95th percentile of export markets, at the 63rd percentile of export spells, and at the 59th percentile of export observations. So these are average trajectories for low-cost firms in large markets.[26]

Four key findings emerge from Figures 1 and 2. First, relative quantity on entry is increasing in duration. Second, there is no statistically significant relationship between relative markups on entry and duration. Third, relative quantity grows fourfold between years one and five of export spells that last at least seven years. This growth is statistically significant up to a horizon of four years and is not driven purely by part-year effects in the first year i.e. there is economically and statistically significant growth between years two and four. Fourth, within export spells that last at least seven years, there are no systematic dynamics in (relative) markups up to a horizon of six years.

Next we confirm that there is a decreasing hazard of exit conditional on firm- and market-specific heterogeneity. Figure 3 graphs the fitted values of the firm-product-market and firm-market export hazards, evaluated at $\left\{\bar{m}^i, \bar{f}^k\right\}$. The omitted category is observations in their first year of export participation, so in both cases the exit probability is normalized to 0 for age equal to 1. Again, full results are reported in the Appendix (Tables 15, 16 and 17).[27] Conditional on our proxies for cost and market size, the probability of exit at both firm-product-market and firm-market levels is initially steeply decreasing in market tenure.

[25]We include exponentials of standard errors in the quantity figure. To make the price figure easier to read, we include only standard errors for the longest spell.

[26]Relative quantity on entry is increasing in f^k. Relative markups on entry are unrelated to either m^i or f^k. We do not see evidence of a systematic relationship between dynamics in either relative quantity or markups and m^i or f^k.

[27]At the firm-product-market level, the exit hazard is less steeply decreasing for high m^i firms than for low m^i firms. At the firm-market level, the exit hazard is less steeply decreasing for high m^i firms and high f^k markets than for low m^i firms and low f^k markets.

This is consistent with selection on the permanent component of idiosyncratic demand. This both implies that it is important to control for this dimension of heterogeneity to correctly identify dynamics, and justifies our use of duration a proxy for the idiosyncratic demand.

In Appendix Table 18, we report entry rates and 1-year exit rates conditional on t $\{\bar{m}^i, \bar{f}^k\}$. These are moments we will match in the structural estimation.[28]

5.2 Interpretation

Figure 1 makes it clear that demand plays an important role in post-entry export behavior. When a firm simultaneously enters multiple equally-sized markets with the same product, initial quantities are on average more than two times bigger in the market where the firm ultimately survives at least seven years than in the market where the firm exits two years after entry. Note that the impact of part-year effects is similar for these two markets; this is not true for the market where the firm survives only one year. Since this is the same firm, and since we control for market size, and for the relationship between duration, marginal cost, and market size, this can only be attributed to a combination of differences in idiosyncratic demand and customer base accumulated in the first year.

In addition, in the market where the firm ultimately survives at least seven years, quantities (relative to average quantities) grow on average 75% between the second and sixth years of participation. This is not driven by part-year effects, since neither of these years are entry or exit years. It is not driven by reductions in marginal cost, since we difference across markets within a given year. Taken together with the behavior of markups, this is consistent with shifts in demand due to accumulation of customer base rather than movement along a demand curve.

As regards prices and markups, Figure 2 shows that when a firm simultaneously enters multiple equally-sized markets, the price it charges in the market where it ultimately survives for at least seven years is statistically indistinguishable from the price it charges in the market where it exits two years after entry. In fact this is true for all durations. Since by definition, price is a markup over marginal cost, under the assumption that marginal cost is the same for all the markets the firm serves, this implies that markups on entry are statistically indistinguishable across markets. This is despite the fact that, according to the behavior of quantities, there are differences in idiosyncratic demand between the two markets which should induce differential incentives to accumulate future customer base.

In addition, in the market where the firm ultimately survives at least seven years, despite

[28]At both levels, entry is increasing in m^i and f^k, while exit is decreasing in m^i and f^k.

quantities growing 75% between years two and six, over this horizon there is no statistically significant evolution of markups relative to average markups. This is contrary to what our model would predict if firms were initially depressing markups to shift out future demand. In this case, we would expect to see relative markup growth accompanying increases in relative quantities sold due to customer base accumulation.

In the next section of the paper, we provide a more formal test of the hypothesis that firms do not use markup distortion to accumulate customer base. We do this by structurally estimating the model. But first, we explain the role of the different components of our baseline estimating strategy, and examine the the robustness of our results along a number of dimensions.

5.3 Building up our specification

Our empirical strategy has many elements. To illustrate the role of each element, we now build it up step-by-step. As a benchmark, we regress log quantities and prices on firm-product-year and product-market-year fixed effects, without any other controls. These fixed effects explain a substantial fraction of the variation in the data: 78% for quantities, and 87% for prices.

Next, we add our controls in steps. The results of this exercise for quantities are reported in Table 6, while results for prices are reported in Table 7. Column (1) of each table reports results from regressing log quantities and prices on the two sets of fixed effects, and a set of indicator variables for duration. Column (2) regresses log quantities and prices on the fixed effects, a set of indicator variables for age, and a dummy for next period exit. The exit dummy allows for nonlinear dynamics prior to exit. In column (3), we include indicator variables for duration, age, and the exit dummy simultaneously. In column (4), we include the full set of interactions between duration and age and report the resulting coefficients on duration when age is equal to 1, and the coefficients on age for spells of duration 7+ years. Full results for this specification are reported in Appendix Table 23.

Finally, in column (5), we reproduce the corresponding coefficients from estimating our baseline specification, which interacts each age-duration indicator with m^i and f^k, evaluating at $\bar{m}^i$ and $\bar{f}^k$.

From Table 6 we learn that for quantities, inference about the nature of dynamics and selection on unobserved heterogeneity is quite sensitive to the specification used. Both selection and dynamics appear to be present, so failure to control for one leads to misleading inference about the other. Based on specification (1), we might infer a bigger role for selection than is actually present, because dynamics are ignored. Based on specification (2), we might infer more dramatic dynamics in quantities than are actually present, because we fail to take account of the fact that initial quantities are systematically lower in short spells than in long spells. Comparing specification (3) with specifications (4) and (5) illustrates that controlling simultaneously for duration and age is not sufficient. If we do not allow dynamics to vary with duration, we might not do too badly in capturing the dynamics in long spells, but we would draw very misleading inference about selection. Finally, allowing dynamics to vary with costs and market size, as well as duration, as we do in our baseline specification (5), affects our understanding of selection, but does not much affect our understanding of dynamics.

In the case of prices, none of the controls add much in terms of R^2. Irrespective of specification, we find no evidence of markup dynamics or of a systematic impact of selection on idiosyncratic demand on markups.

5.4 Robustness

5.4.1 Specification

We estimate a specification which includes firm-product-market-spell fixed effects to control for the first order effect of idiosyncratic demand. This specification has the advantage that it uses only within-spell variation in age to identify dynamics, but it cannot identify how initial quantities and markups are related to idiosyncratic demand. If anything, the dynamics in quantities identified using this approach are a little steeper than those based on our baseline specification, while the behavior of markups is very similar to the baseline. The relevant tables are reported in Appendix F and the relevant figures are reported in Appendix G, as are the results of all robustness checks in this subsection.

We experiment with alternatives to our baseline proxies for marginal cost and market size, i.e. number of markets per firm, m^i, and number of firms per market, f^k. Results are very similar when we use logs instead of levels of m^i and f^k. Results are also very similar when we use the number of markets per firm-product and the number of firms per product-market (m^{ij} and f^{jk}) instead of m^i and f^k.

We estimate a specification where we drop unit value outliers. Our criterion for an outlier is an observation where the absolute value of the the log change in the unit value between the current and previous period exceeds 2. Results are very similar to the baseline. We also estimate a specification where we use only observations which have a measure of quantity which is not tonnes (the default measure of quantity). This reduces the sample size by a factor of 8, and results are very noisy as a result.

We estimate a specification where we topcode age and duration at 10 years rather than at 7 years. To do so, we make use of the longer sample of customs data (1996-2014) which is not matched to CIP firms. The behavior of relative quantities and markups is qualitatively very similar to the baseline using this specification and data.

5.4.2 Firm, product, and market characteristics

Foreign multinationals have a substantial presence in Ireland in the manufacturing sector. They are export-intensive (mainly platform FDI) and due to this, account for 55% of the firms in our baseline analysis of quantities and prices. We confirm that our results are very similar for domestic-owned and foreign-owned firms by splitting the sample, and re-estimating on the two subsamples.

It is possible that the importance of marketing and advertising relative to markup distortions in accumulating customer base could differ across sectors. We split the sample into four broad sectors, two consumer-facing (consumer food, consumer non-food non-durables), and two business-facing (intermediate goods, capital goods), and estimate our baseline equation separately for each sector. Splitting the sample in this way greatly reduces sample size, and results are correspondingly noisy. However they are qualitatively consistent with the baseline in all sectors.

We also check whether results are similar for markets which are close, and for which trade barriers are low,[29] and for markets which are distant, and barriers are greater. We do this by splitting the sample into EU markets and non-EU markets. Results are very similar for the two different samples.

5.4.3 Are our results driven by special features of the Irish data?

One possible concern is that our results may be due to some special features of the Irish data. To alleviate such concerns, we note that we can replicate the findings of a large body of literature working with firm and customs micro data for other countries. As mentioned above, summary statistics on the cross-sectional dimension of exporting in our data are in line with those for other small open European economies. Our findings on the post-entry dynamics of revenues and exit are consistent with those in the previous literature, (e.g.,

[29] Note that as long as trade barriers are the same for all firms selling a given product to a given market, they are captured by the product-market-year fixed effects.

Eaton et al. (2014), Ruhl and Willis (2017)). Using our data, we can also replicate a number of facts about the behavior of unit values in customs data for other countries. In the cross-section, export prices vary with destination market characteristics just as in the literature surveyed in Harrigan et al. (2015). Meanwhile, in Fitzgerald et al. (2019), we show that the degree of pricing-to-market in our data is very similar to that for other countries (e.g. France, as shown by Berman et al. (2012)). The consistency of these findings with those based on other data sets suggests to us that our results cannot be attributed to special features of the Irish data.

5.4.4 Relation to the empirical literature on price dynamics

A series of papers prior to and contemporaneous with ours investigate price dynamics in a variety of empirical settings. Some of these papers arrive at results which differ from ours. We hypothesize that some of these contrasting results are due to the use of different empirical specifications, while others are due to the use of different data.

Foster et al. (2008) use the quinquennial U.S. manufacturing census data on plants in a narrow set of commodity-like sectors, and find that older plants have *higher* prices. Their specification does not control for firm- or firm-product-level unobserved heterogeneity or condition on duration, and may therefore confound selection on this dimension of heterogeneity with dynamics. We estimate a specification which resembles theirs using our data and find an increasing (though not always statistically significant) relationship between prices and age. We hypothesize that this may be due to selection on quality, or to quality upgrading by successful firms. Both of these are controlled for in our baseline specification by the inclusion of firm-product-year fixed effects.

Using customs data for France, Berman et al. (2019) show a *decreasing* relationship between prices and age in an export market. In contrast to us, they do not control for product-market-year effects or duration in their empirical specification. When we implement their specification, we also obtain negative point estimates, but the magnitudes are quantitatively small, and we do not find statistical significance. We hypothesize that they may find statistical significance because they work with a much larger data set.

Piveteau (2021) uses the same French customs data as Berman et al. (2019), and concludes that prices are *increasing* with tenure in a market. He regresses log price on age, including firm-product-year as well as product-market-year fixed effects. With this specification, he finds that prices in observations with age of 7 years are 4% higher than prices in observations which have just entered. We cannot replicate this finding in our data.

Finally, Argente et al (2021) use Nielsen data on consumer food sales in 206 distinct geographical markets in the U.S. to estimate a restricted version of our baseline specification (they do not interact the vector of duration-age indicators with proxies for firm- and market-level heterogeneity). They find that markups are invariant to tenure, with the exception of markup declines immediately before exit from a market which are associated with fire-sales at the store level.

6 Structural estimation

We now turn to structural estimation of the model from Section 2 by simulated method of moments. There are two possible approaches to estimating the model. One is to estimate a model with fixed firm and market types, targeting moments that condition on firm and market heterogeneity. An alternative is to estimate processes for firm and market heterogeneity, and to include target moments that are informative about this heterogeneity.

We take the first approach. This is one of the reasons we are careful to evaluate the post-entry trajectories of relative quantities and markups in Figures 1 and 2 at fixed values of our proxies for marginal cost and market size, i.e. $\left\{\bar{m}^i, \bar{f}^k\right\}$. This also implies that some of our parameter estimates (we note below which ones) are tied to these fixed firm and market types.

6.1 Assumptions about distributions and functional forms

As described above, we abstract from variation in marginal cost and market size in our estimated model, so we normalize both costs (C_t^i) and market size ($Y_t^k = Q_t^k \left(P_t^k\right)^\theta$) to 1.

We assume the following process for idiosyncratic demand: $\nu_t^{ik} = \bar{\nu}^{ik} + \tilde{\nu}_t^{ik}$, where $\bar{\nu}^{ik} \sim N\left(0, \sigma_\nu^2\right)$ and $\tilde{\nu}_t^{ik} = \rho \tilde{\nu}_{t-1}^{ik} + \eta_t^{ik}$, with $\eta_t^{ik} \sim N\left(0, \sigma_\eta^2\right)$.

We assume the following process for the sunk cost:

$$S_t^{ik} = \begin{cases} 0 & \text{with probability } \lambda \\ \infty & \text{with probability } 1-\lambda \end{cases}$$

With probability $\lambda \in [0,1]$, entry is possible, while with probability $1-\lambda$, entry is not possible. In the absence of cost and market size heterogeneity, this assumption is without loss of generality, though the value of λ is tied to the mean values of the cost and market size proxies at which our target moments are evaluated. Note that just because entry is possible, it does not mean that a firm will choose to enter a market. The decision to enter depends also on the realizations of fixed cost, F_t^{ik} and idiosyncratic demand $\left\{\bar{\nu}^{ik}, \tilde{\nu}_t^{ik}\right\}$.

We assume the following process for the fixed cost:

$$F_t^{ik} = \begin{cases} 0 & \text{with probability } (1-\omega)\,\gamma \\ F & \text{with probability } (1-\omega)\,(1-\gamma) \\ \infty & \text{with probability } \omega \end{cases}$$

where $\omega \in [0,1]$ and $\gamma \in [0,1]$. If $\gamma < 1$ and $0 < \omega < 1$, this process generates both exogenous and endogenous exit. If $0 < \gamma < 1$, it contributes to selection on idiosyncratic demand and a downward-sloping exit hazard. This is because a firm may participate in a market where idiosyncratic demand is weak as long as $F_t^{ik} = 0$, but as soon as it draws a realization of $F_t^{ik} = F$, exit is triggered. In contrast, in markets where idiosyncratic demand is strong, the firm will continue to participate as long as $F < \infty$. Again, the parsimony of this specification is justified by the absence of cost and market size heterogeneity, but the estimated parameters are tied to the mean values of the cost and market size proxies.

We assume that the cost of investment takes the form:

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

This builds in quadratic adjustment costs, as is standard in the literature on investment in physical capital, and irreversibility, which makes sense in the context of intangible investment.

Because all firms and markets are identical, we estimate a single value $\underline{D}$ for the exogenous inital customer base. This is also tied to the mean values of the cost and market size proxies at which our target moments are evaluated.

6.2 Estimation procedure

Given values for the parameters, we discretize both exogenous and endogenous states and use value function iteration to solve for the optimal policies (i.e. participation X_t^{ik}, advertising and marketing effort A_t^{ik}, and prices P_t^{ik}). Using the model parameters, the value functions, and the corresponding optimal policies, we simulate post-entry trajectories for 10,000 "firm-markets" for 14 periods. To account for the fact that there are part-year effects in the data, the length of a period in our model is 6 months. We stagger entry across 6-month periods, and aggregate up to an annual frequency to construct the equivalents of the moments we report in Section 5 out to 7 years post-entry. The goal of our estimation is to choose the vector of parameters that best matches these moments.

We match the following 59 moments: the ratios of initial relative quantities and markups across spells of different duration to relative quantities and markups in 1-year spells; the evolution of relative quantities and markups with age for export spells of different duration; the evolution of exit probabilities at the firm-market level with age, normalized by the exit probability in the first year; the rate of entry at the firm-market level, and the exit rate in the first year in a market, also at the firm-market level. See Tables 62 and 63 in the Appendix for the complete list of moments.[30]

We preset the rate at which firms discount the future. Since the period length is 6 months, we set $\beta = 1.05^{-0.5}$. We also preset the trade elasticity (i.e. the long-run elasticity of exports with respect to e.g. tariffs), which in our model is given by $\theta/(1-\alpha)$. Based on the work of Fitzgerald and Haller (2018), we set it equal to 3. This value is within the range of estimates in the literature.[31]

There are then 12 parameters to be estimated: $\left\{\sigma_\nu^2, \sigma_\eta^2, \rho, \lambda, F, \omega, \gamma, \underline{D}, \alpha, \delta, \phi, \psi\right\}$. Note that given our assumption about the trade elasticity the implied value of θ will be given by $3(1-\hat{\alpha})$. We choose the parameters to minimize the criterion function $m'Vm$, where m is the difference between the data moments and the equivalent 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 use a combination of a particle swarm algorithm and the simplex method to optimize over the parameter vector.

6.3 Baseline results

Table 8 reports the estimated parameters and the optimized value of the criterion function $m'Vm$. Strikingly, our preferred parameter vector is a corner solution, with $\psi = 1$, so firms use only marketing and advertising to accumulate customer base, and implicitly, current sales do not affect future customer base. We estimate decreasing returns to customer base ($\alpha = 0.41$), as well as a high depreciation rate. Since a period in our model is six months,

[30]We match entry and exit moments at the firm-market level rather than at the firm-product-market level because we do not wish to address the role of the extensive margin of products in our analysis.

[31]See Head and Mayer (2014) for a survey.

the annual depreciation rate implied by our estimate of δ is 85%.[32]

Figures 4, 5 and 6 illustrate the fit of the estimated model to the target quantity, price and exit moments. The corresponding tables are reported in Appendix H (Tables 62 and 63). The model provides a good fit to all moments. It can generate dispersion in initial quantities that is positively correlated with spell duration, and of the right order of magnitude. Quantities increase with tenure in successful spells as in the data. Since $\psi = 1$, markups are constant. Initial markups are therefore uncorrelated with spell duration, and markups are flat with respect to age. The exit hazard closely matches the data, and the rate of entry matches that in the data.

Our structural estimates thus confirm the findings of our reduced form analysis, and suggest that firms use non-price activities such as marketing and advertising rather than markups to accumulate customer base and increase market share.

6.4 Robustness

We first check that our results are robust to the choice of trade elasticity. To understand why this matters, it is useful to note that when $\psi = 1$ (i.e. firms use only marketing and advertising to invest in customer base as in our baseline estimates), the fit of the model is invariant to the trade elasticity.[33] However when future customer base depends on current sales (i.e. $\psi < 1$), the trade elasticity matters, because given $\hat{\alpha}$, it governs the price elasticity of demand θ. When θ is low, big markup discounts today are required to generate a given shift in demand tomorrow. When θ is high, small markup discounts today can generate the same shift in demand tomorrow.

We re-estimate the model, setting the trade elasticity equal to 5 and 8, which are within the range of values found in the literature. We also set the trade elasticity equal to 100, which is well outside the plausible range. In each case, the criterion is minimized at our baseline parameter estimates (with F and $\underline{D}$ appropriately scaled).

In addition, we estimate a version of the model where we set $\psi = 0$ (i.e. assume no role for marketing and advertising in accumulating customer base), and instead of using the trade elasticity to pin down θ, we estimate θ along with the other parameters. The fit is worse than that of the baseline, model and most strikingly, the trade elasticity implied by the resulting estimates of θ and α is just over 100, an implausibly large value. These results are reported in Table 64 in Appendix H, and Figures 51-53 in Appendix I.

[32]We report standard errors constructed using the method suggested by Gourieroux et al. (1993). Details of this method are reported in Appendix E. We do not report a one-sided standard error for ψ.

[33]F and $\underline{D}$ must be scaled by a function of θ and α, but the remaining parameters are invariant to θ.

We conclude that our baseline result that the data prefers a model where firms invest in customer base only through marketing and advertising does not depend on our choice of trade elasticity.

6.5 Selling expenses

An advantage of structural estimation is that we can use our estimates to back out the ratio of selling expenses in a market to revenue in that market:[34]

$$sell_t^{ik} = \frac{c\left(D_t^{ik}, A_t^{ik}\right)}{P_t^{ik} Q_t^{ik}}$$

In Figure 7, we show the evolution of this ratio for export spells of different duration. As a share of revenue, selling expenses are highest at the beginning of an export spell, and decline with age thereafter.[35] Initial selling expenses are also higher in spells that are ultimately successful than in spells that are ultimately unsuccessful, as firms invest more when idiosyncratic demand, and therefore the probability of survival, is higher. Our estimates suggest that selling expenses on average account for 18.5% of revenue in the 6th year of export spells that last 7+ years.

We do not have any data on selling expenses for our firms, but Arkolakis (2010) calculates that marketing and advertising expenditures may account for 7-8% of U.S. GDP. The CMO Survey of chief marketing officers in the U.S. finds that over the period 2008-2018, firms in goods-producing sectors report spending between 7% and 11% of revenue on marketing.[36] As reported in the Introduction, Gourio and Rudanko (2014) note that the share of Selling, General & Administrative expenses in total revenue for Compustat firms is even higher, on the order of 17-27%. Traina (2018) notes that this share has increased from 12% on average in 1950 to 22% today. Our estimates for the share of selling expenses in revenue are within this range.

[34] We do not treat the fixed cost of export participation as part of selling expenses, and in this sense, the numbers we report are a lower bound.

[35] The steady state ratio of selling expenses to revenue in the model is governed by $\{\alpha, \delta, \phi, \theta\}$ and the probability of exit. We report the formula in Appendix B.

[36] See answers to the question "Marketing expenses account for what percent of your firm's revenues?" for goods-producing firms selling business-to-business and business-to-consumer.

7 Measured productivity around a trade liberalization

Because of the potential impact on gains from trade, economists have long been interested in the possibility that exporting has a positive causal impact on firm-level productivity.[37] But as we note in the Introduction, especially around episodes of market entry, measured productivity may be affected by changes in selling costs. This may systematically bias what researchers infer about the relationship between exporting and productivity. We now use our structural estimates to get a sense of how big this mismeasurement might be by simulating a trade liberalization episode which induces some firms which were previously non-exporters to enter the export market, and which induces incumbents to export more.

The exercise is set up as follows.[38] We simulate a panel of 50,000 firms whose marginal costs (i.e. unit labor requirement, or the inverse of productivity) are drawn from a lognormal distribution. There is a home market and a foreign market. In each market, firms face the process for idiosyncratic demand we estimate in the previous section. In the home market, there are no sunk or fixed costs of participation; nor are there any shocks which trigger exit. In the foreign market, firms face both sunk and fixed costs of participation. We choose the dispersion of productivity to generate a 90-10 ratio of productivity of just over 3.[39] We choose the relative size of the home and export markets, and the process for sunk and fixed costs of export participation to generate an export participation rate of 29%, an export intensity conditional on exporting of 32% and an exporter size premium in terms of sales of 1.75.[40]

We then assume that the ad valorem tariff in the export market is unexpectedly reduced from 25% to 0%. This implies that the export market size increases by a factor of 1.25^{θ}. We solve for the optimal responses to this change, and simulate the evolution of exports over 10 years, assuming that costs for each firm (and therefore, implicitly also wages) are held fixed.[41] Figure 8 shows how export participation, export intensity, and the exporter premium evolve over these 10 years. Export participation increases gradually from 29% to 34%, while export intensity increases much more rapidly, by more than 10 percentage points. The exporter size premium overshoots in initially, rising to close to 2.2, but falls back to a value just over 2 as more high cost firms gradually enter the export market.

Based on the simulated data, we calculate measured revenue-based TFP (i.e. attributing

[37] See e.g. De Loecker (2007) and Lileeva and Trefler (2010).

[38] We provide more details in Appendix J.

[39] This is similar to the ratio reported by Hsieh and Klenow (2009) for U.S. manufacturing in 1997.

[40] The respective values in the Irish data are 44%, 32%, and 1.85.

[41] Implicitly, we assume there is an outside sector which pins down wages.

selling costs to production cost), using the formula:

$$TFPR_t^i = \frac{Q_t^{ih} P_t^{ih} + Q_t^{if} P_t^{if}}{\left(Q_t^{ih} + Q_t^{if}\right) C^i + c\left(D_t^{ih}, A_t^{ih}\right) + c\left(D_t^{if}, A_t^{if}\right)} \tag{6}$$

Figure 9 shows the average behavior of measured TFPR, for firms which start exporting in response to the change in market size, for firms which exported both before and after the change in market size, and for all firms. In each case, measured TFPR is expressed relative to measured TFPR under the no trade liberalization baseline. Note that true TFPR is constant, equal to $\theta/(\theta-1)$.

There are two distinct biases in measured TFP. Note that in our simulation, the behavior of quantity-based TFP (TFPQ) is almost identical to that of TFPR. In the immediate aftermath of the trade liberalization, measured TFP is depressed, as firms devote substantial resources to acquiring customers in the foreign market. This effect is bigger for incumbents than for entrants, as they have higher export intensity. However as firms reach their new steady state customer base, measured TFP eventually rises above pre-trade liberalization TFP. This is because the investment-sales ratio for the export market is lower than that for the home market due to the lower survival probability firms face in the foreign market. Though modest, it is possible that this effect could contribute at the margin to the appearance of within-firm productivity gains from exporting in the literature cited above. This suggests that some caution is merited in interpreting the findings of the literature which finds within-firm productivity gains in response to export entry.

8 Conclusion

We use customs data for Ireland to show that successful entry into an export market is associated with substantial growth in quantities conditional on costs, but no change in markups. This is compelling evidence that customer base and demand play an important role in post-entry export dynamics. These facts also suggest that firms invest in customer base primarily through marketing and advertising rather than by distorting markups to shift future demand. We show this formally by structurally estimating a model which nests these two possibilities. We find that the data prefer a model where firms use only marketing and advertising to attract customers. Our parameter estimates allow us to back out the selling expenses associated with building and maintaining customer base. We find that they are substantial. In a simulation exercise, we show that misattibution of selling expenses to production cost could potentially contribute to productivity mismeasurement, and in particular, could lead to the appearance of within-firm productivity dynamics due to export entry, even when true underlying productivity does not change.

References

[1] Abraham, K. and S. Farber (1987), "Job Duration, Seniority, and Earnings" *American Economic Review* 77, 278-297.

[2] Argente, D., D. Fitzgerald, S. Moreira and A. Priolo (2021), "How Do Firms Build Market Share?" mimeo.

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

[4] Arkolakis, C. (2016), "A Unified Theory of Firm Selection and Growth," *Quarterly Journal of Economics* 131 (1), 89-155.

[5] Bastos, P., D. A. Dias and O. Timoshenko (2018), "Learning, Prices, and Firm Dynamics," *Canadian Journal of Economics* 51, 1257-1311.

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

[7] Berman, N., V. Rebeyrol, and V. Vicard (2019), "Demand Learning and Firm Dynamics: Evidence from Exporters," *Review of Economics and Statistics* 101, 91-106.

[8] Bils, M. (1989), "Pricing in a Customer Market," *Quarterly Journal of Economics* 104, 699-718.

[9] CMO Survey, Duke University, 2008-2018, cmosurvey.org

[10] De Loecker, J., J. (2007), "Do Exports Generate Higher Productivity? Evidence from Slovenia," Journal of International Economics 73, 69-98.

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

[12] Drozd, L. A., and J. B. Nosal (2012b), "Pricing-to-Market in Business Cycle Models," mimeo.

[13] Eaton, J., S. Kortum, and F. Kramarz (2011), "An Anatomy of International Trade: Evidence From French Firms," Econometrica 79 (5), 1453-1498.

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

[15] Fitzgerald, D. and S. Haller (2018), "Exporters and Shocks," *Journal of International Economics* 113, 154-171.

[16] Fitzgerald, D., Y. Yedid-Levi and S. Haller (2019), "Can Sticky Quantities Explain Export Insensitivity to Exchange Rates?" mimeo.

[17] 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.

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

[19] Gilchrist, S., R. Schoenle, J. Sim and E. Zakrajsek (2017), "Inflation Dynamics During the Financial Crisis," *American Economic Review* 107 (3), 785-823.

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

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

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

[23] Head, K. and T. Mayer (2014), "Gravity Equations: Workhorse, Toolkit, and Cookbook," Chapter 3 in *Handbook of International Economics*, Editor(s): Gita Gopinath, Elhanan Helpman, Kenneth Rogoff, Elsevier, Volume 4, Pages 131-195.

[24] Hottman, C., S. Redding, and D. Weinstein (2016), "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics* 131, 1291-1364.

[25] Hsieh, C.-T. and P. Klenow (2009), "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics* 124, 1403-1448.

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

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

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

[29] 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.

[30] Piveteau, P. (2021), "An Empirical Dynamic Model of Trade with Customer Accumulation," *American Economic Journal: Microeconomics* 13, 23-63.

[31] Ravn, M., S. Schmitt-Grohe, and M. Uribe (2006), "Deep Habits," *Review of Economic Studies* 73 (1), 195-218.

[32] Ruhl, K., and J. Willis (2017), "New Exporter Dynamics," *International Economic Review* 58, 703-726.

[33] Traina, J. (2018), "Is Aggregate Market Power Increasing? Production Trends Using Financial Statements," mimeo.

[34] 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
Exporter size premium (employees, mean)	1.65
Exporter size premium (revenue, mean)	1.85
Mean export share conditional on exporting	0.32
Mean number of markets per exporter	6.6

Notes: Statistics are for our cleaned data set of CIP firms. Firms are defined as exporters if they are matched to positive concorded product exports from customs data. Export intensity is calculated as total concorded product exports from customs divided by sales reported in the CIP. Values greater than 1 are replaced by 1. Source: CSO and authors' calculations.

Table 2: 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	11	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: Correlations of m^i and f^k with employment, GDP and distance

	m^i	Log emp.	Rev/worker	TFP	f^k	Sh. GDP	Log dist.
# markets per firm (m^i)	1						
Log employment	0.57	1					
Rev/worker	0.37	0.32	1				
TFP	0.17	0.04	0.56	1			
# firms per market (f^k)	-0.29	-0.10	-0.16	-0.12	1		
Sh. world GDP	-0.15	-0.05	-0.11	-0.06	0.67	1	
Log distance	0.18	0.05	0.02	0.06	-0.43	-0.05	1

Notes: For correlations between firm-level variables, an observation is a firm. For correlations between market-level variables, an observation is a market. For correlations between firm-level and market-level variables, an observation is an export spell at the firm-product-market level. Source: CSO and authors' calculations.

Table 4: Distribution of duration: Export spells and export observations

Obs. level	Firm-mkt		Firm-prod-mkt	
duration	Spells	Obs.	Spells	Obs.
1	0.34	0.09	0.51	0.26
2	0.10	0.06	0.10	0.10
3	0.05	0.04	0.04	0.06
4	0.03	0.03	0.02	0.04
5	0.02	0.03	0.01	0.03
6	0.01	0.02	0.01	0.02
7+	0.06	0.17	0.02	0.11
Left cens.	0.25	0.45	0.17	0.28
Right cens.	0.14	0.10	0.11	0.11
N	55,131	187,409	262,969	526,438

Notes: Table reports share of relevant unit of observation (export spells, or export observations) with given duration, for data at the firm-product-market level, and data aggregated to the firm-market level. Source: CSO and authors' calculations.

Table 5: Regression of duration on m^i and f^k

	coeff	s.e.
m^i	0.40	(0.01)**
f^k	0.62	(0.02)**
N	188,435	
R^2	0.006	

Notes: Dependent variable is duration. Observations are at the firm-product-market spell level. Duration is top-coded at 7. Left- and right-censored spells are excluded. Source: CSO and authors' calculations.

Table 6: Building our specification: Quantity

	(1)		(2)		(3)		(4)		(5)	
	duration		tenure, exit		dur, tenure, exit		dur × tenure		baseline	
	duration						duration at tenure = 1			
2 yrs	0.47	(0.03)**			-0.53	(0.03)**	0.50	(0.04)**	0.57	(0.04)**
3 yrs	0.83	(0.03)**			-0.60	(0.04)**	0.74	(0.05)**	0.86	(0.06)**
4 yrs	1.10	(0.04)**			-0.57	(0.05)**	0.82	(0.07)**	1.04	(0.07)**
5 yrs	1.31	(0.04)**			-0.54	(0.06)**	0.98	(0.09)**	1.17	(0.10)**
6 yrs	1.51	(0.05)**			-0.45	(0.06)**	0.92	(0.11)**	1.07	(0.12)**
7+ yrs	2.26	(0.03)**			0.06	(0.06)	1.18	(0.07)**	1.42	(0.07)**
	tenure						tenure at duration = 7+			
2 yrs			0.61	(0.02)**	0.85	(0.03)**	0.86	(0.08)**	0.80	(0.09)**
3 yrs			0.93	(0.03)**	1.10	(0.04)**	1.19	(0.08)**	1.13	(0.09)**
4 yrs			1.18	(0.04)**	1.25	(0.05)**	1.31	(0.08)**	1.24	(0.09)**
5 yrs			1.37	(0.04)**	1.32	(0.05)**	1.38	(0.08)**	1.35	(0.09)**
6 yrs			1.53	(0.05)**	1.34	(0.06)**	1.29	(0.09)**	1.27	(0.09)**
7+ yrs			1.79	(0.05)**	1.49	(0.06)**	1.32	(0.08)**	1.25	(0.08)**
l-cens	2.61	(0.03)**	1.85	(0.03)**	1.63	(0.04)**	2.62	(0.03)**	2.75	(0.03)**
r-cens	1.22	(0.04)**			-0.37	(0.05)**	1.32	(0.04)**	1.54	(0.04)**
Exit			-1.13	(0.02)**	-1.25	(0.03)**				
f.e.	fpy & pmy		fpy & pmy		fpy & pmy		fpy & pmy		fpy & pmy	
N	183,831		171,683		171,683		183,831		183,831	
R^2	0.80		0.81		0.81		0.81		0.81	

Notes: Dependent variable is log quantity at the firm-product-market level. All equations include firm-product-year and product-market-year fixed effects. Column (4) includes full set of duration-tenure interactions and reports only a subset of results. Column (5) reports a subset of the coefficients from our baseline specification as reported in Table 9 in Appendix F. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

Table 7: Building our specification: Price

	(1)		(2)		(3)		(4)		(5)	
	duration		tenure, exit		dur, tenure, exit		dur × tenure		baseline	
	duration						duration at tenure = 1			
2 yrs	-0.02	(0.02)			-0.01	(0.02)	-0.02	(0.02)	-0.02	(0.02)
3 yrs	0.00	(0.02)			0.03	(0.02)	0.01	(0.03)	0.00	(0.03)
4 yrs	0.02	(0.02)			0.00	(0.03)	0.01	(0.04)	0.01	(0.04)
5 yrs	-0.02	(0.02)			0.00	(0.03)	-0.01	(0.05)	-0.02	(0.05)
6 yrs	-0.03	(0.02)			-0.00	(0.03)	0.01	(0.05)	0.00	(0.05)
7+ yrs	-0.07	(0.02)**			-0.04	(0.03)	-0.04	(0.03)	-0.05	(0.03)
	tenure						tenure at duration = 7+			
2 yrs			-0.01	(0.01)	-0.01	(0.02)	-0.02	(0.04)	0.00	(0.04)
3 yrs			-0.03	(0.02)**	-0.04	(0.02)*	-0.08	(0.04)**	-0.05	(0.04)
4 yrs			-0.02	(0.02)	-0.01	(0.02)	-0.03	(0.04)	-0.01	(0.04)
5 yrs			-0.02	(0.02)	0.00	(0.03)	-0.02	(0.04)	-0.01	(0.04)
6 yrs			-0.04	(0.03)	-0.01	(0.03)	-0.03	(0.04)	-0.03	(0.04)
7+ yrs			-0.07	(0.02)**	-0.02	(0.03)	-0.05	(0.04)	-0.04	(0.02)**
l-cens	-0.04	(0.02)**	-0.02	(0.01)**	-0.03	(0.02)	-0.04	(0.02)**	-0.05	(0.02)**
r-cens	0.01	(0.02)			0.03	(0.03)	0.01	(0.02)	0.01	(0.02)
Exit			0.02	(0.01)**	0.02	(0.01)				
f.e.	fpy & pmy		fpy & pmy		fpy & pmy		fpy & pmy		fpy & pmy	
N	183,831		171,683		171,683		183,831		183,831	
R^2	0.87		0.87		0.87		0.87		0.87	

Notes: Dependent variable is log price at the firm-product-market level. All equations include firm-product-year and product-market-year fixed effects. Column (4) includes full set of duration-tenure interactions and reports only a subset of results. Column (5) reports a subset of the coefficients from our baseline specification as reported in Table 9 in Appendix F. Robust standard errors calculated. ** significant at 5%, * significant at 10%. Source: CSO and authors' calculations.

Table 8: Structural model: parameter estimates

	σ_ν	σ_η	ρ	λ	$F^\dagger$	ω	γ	$\underline{D}^\S$	α	δ	ϕ	ψ
Param	0.63	0.09	0.87	0.05	0.07	0.07	0.68	0.08	0.41	0.62	1.46	1
s.e.	0.00	0.00	0.00	0.00	n.a.	0.00	0.00	n.a.	0.00	0.00	0.00	n.a.

Notes: † The value reported for F is the average ratio of F_t^{ik} to revenue across all participants in their first period (6 months) of participation. This includes participants for whom $F_t^{ik} = 0$. The estimate of the parameter governing F i is 0.06 and the standard error is 0.00. § The value reported for $\underline{D}$ is the average of $\underline{D}/D_{13}$ across all participants who survive 13 (6-month) periods in the market, and have nonnegative investment in period 13, where D_{13} is customer base in period 13. The estimate of the parameter governing $\underline{D}$ is 0.04 and the associated standard error is 0.00. Standard errors are calculated following Gourieroux et al. as described in Appendix E. We do not calculate a standard error for ψ, since the estimated value is at a corner.

Figure 1: Estimated quantity trajectories

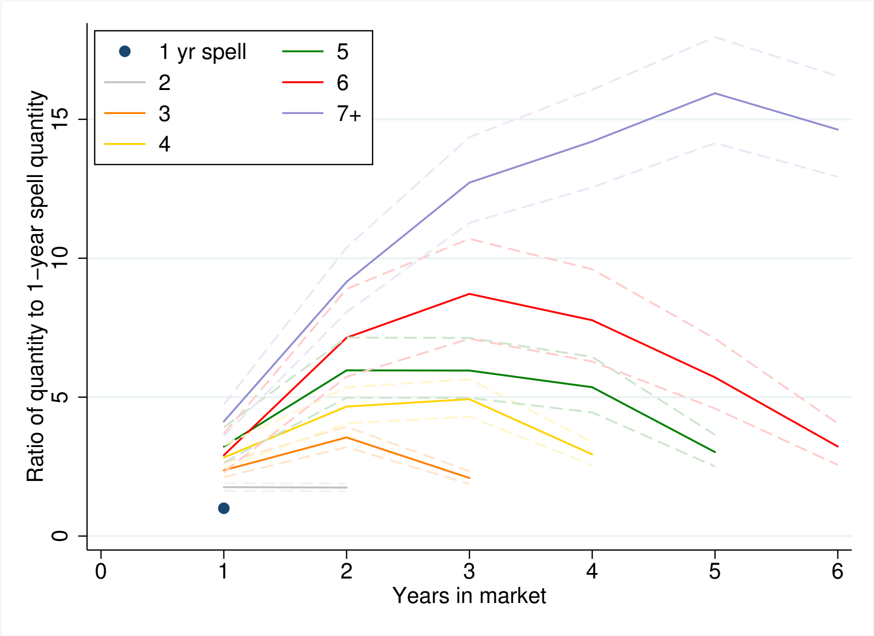

Notes: Figure shows evolution of quantities at the firm-product-market level with market tenure, allowing trajectories to differ by export spell duration. Trajectories are conditional on firm-product-year and market effects. 95% confidence intervals are plotted. Corresponding table is Table 9 in Appendix F. Source: CSO and authors' calculations.

Figure 2: Estimated price trajectories

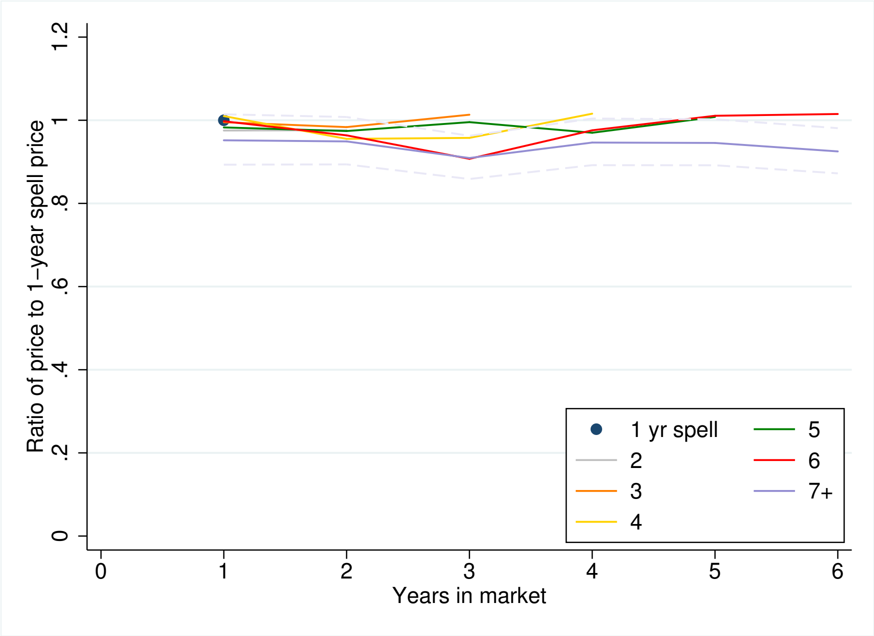

Notes: Figure shows evolution of prices at the firm-product-market level with market tenure, allowing trajectories to differ by export spell duration. Trajectories are conditional on firm-product-year and market effects. 95% confidence interval for spells of 7+ years is plotted. Corresponding table is Table 9 in Appendix F. Source: CSO and authors' calculations.

Figure 3: Estimated exit trajectories

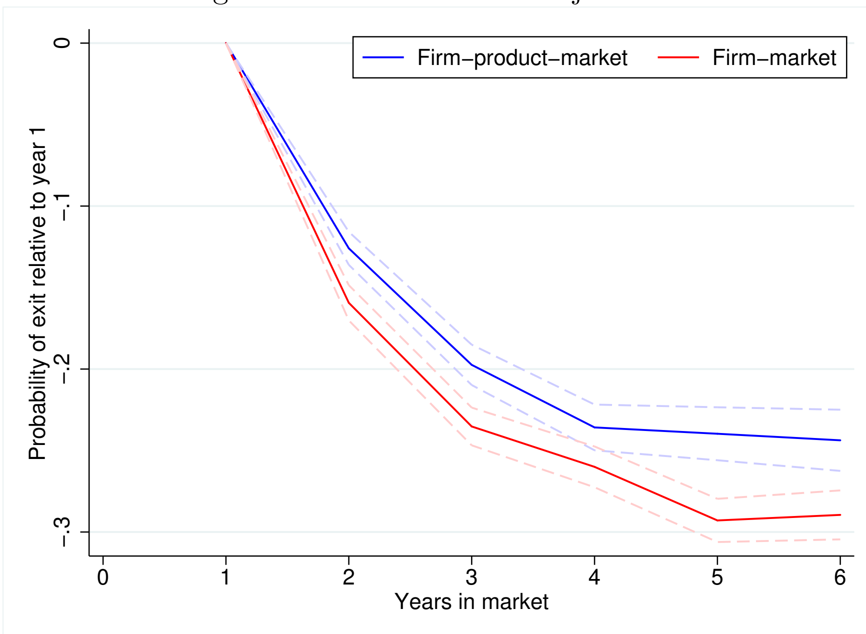

Notes: Figure shows reduction in probability of exit at the firm-market and firm-product-market levels with compared to probability of exit in the first year in a market. Trajectories are conditional on firm-year and market and firm-product-year and market effects, respectively. 95% confidence intervals are plotted. Corresponding Table is Table 17 in Appendix F. Source: CSO and authors' calculations.

Figure 4: Model fit: Quantities

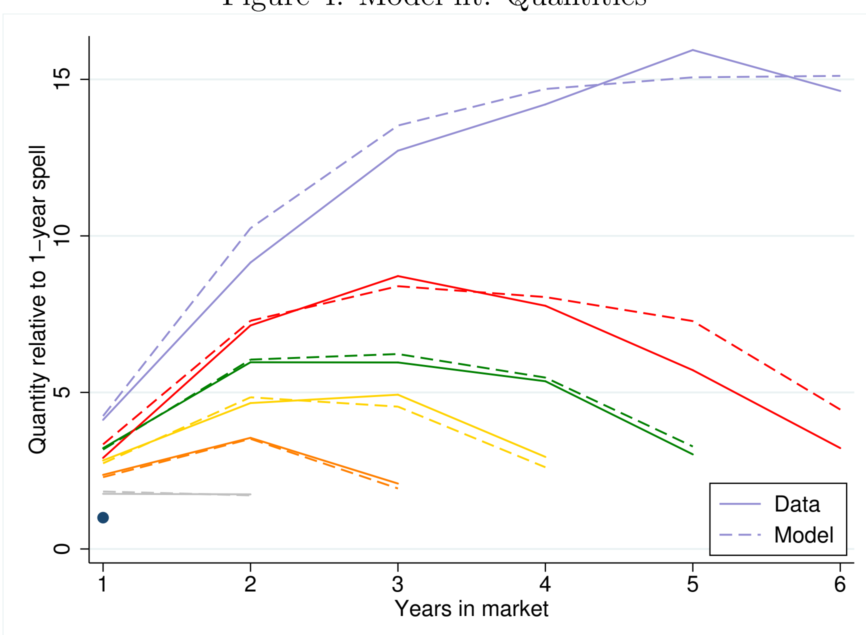

Notes: Figure shows evolution of quantities with age, for spells of different duration. Data is from Figure 1, while Model refers to the corresponding fitted values from our estimated model. All quantities are expressed relative to the quantity in a 1-year spell. Source: CSO and authors' calculations.

Figure 5: Model fit: Prices

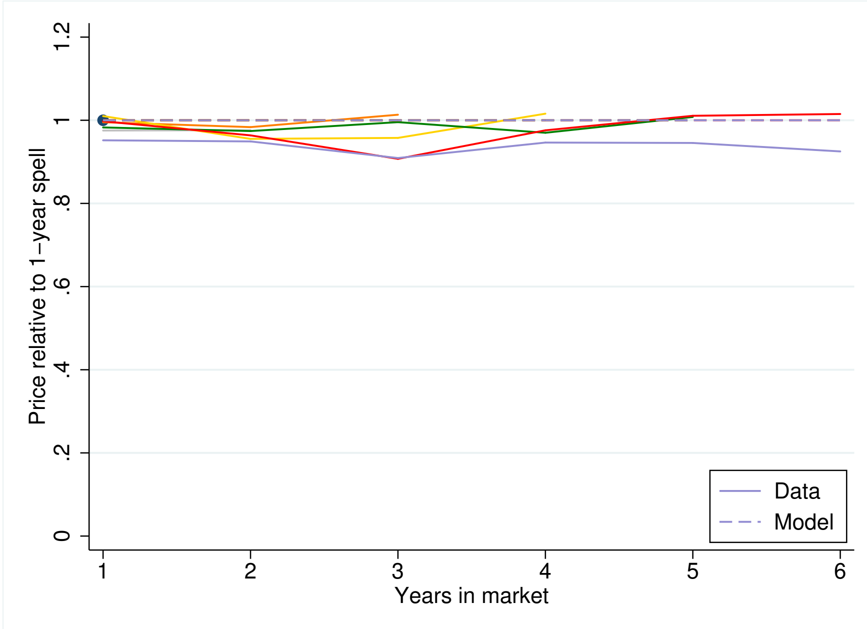

Notes: Figure shows evolution of prices with age, for spells of different duration. Data is from Figure 2, while Model refers to the corresponding fitted values from our estimated model. All prices are expressed relative to the price in a 1-year spell. Source: CSO and authors' calculations.

Figure 6: Model fit: Exit

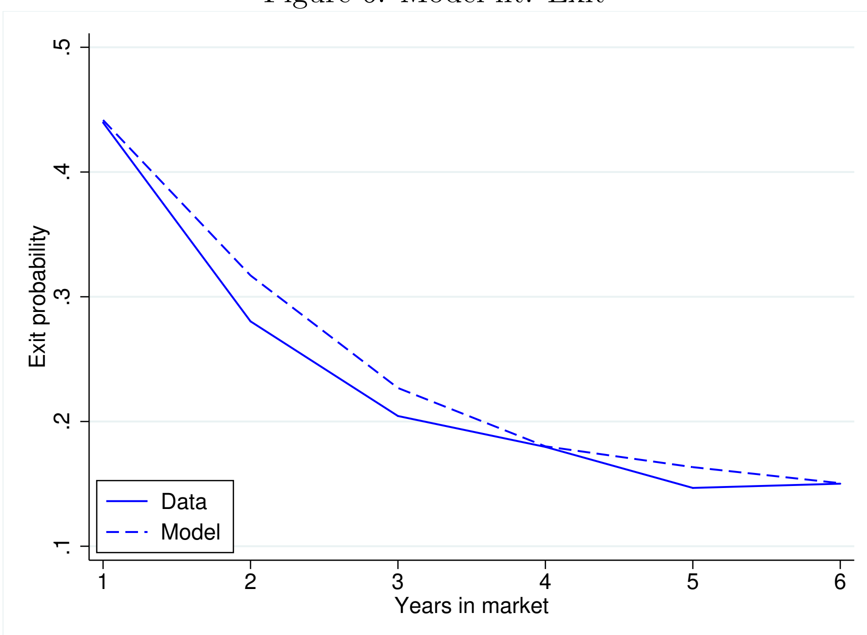

Notes: Figure shows evolution of probability of exit with tenure. Data (at the firm-market level) is from Figure 3 combined with the 1-year exit rate of 0.44, while Model refers to the corresponding fitted values from our estimated model. Source: CSO and authors' calculations.

Figure 7: Selling expenses as a share of revenue

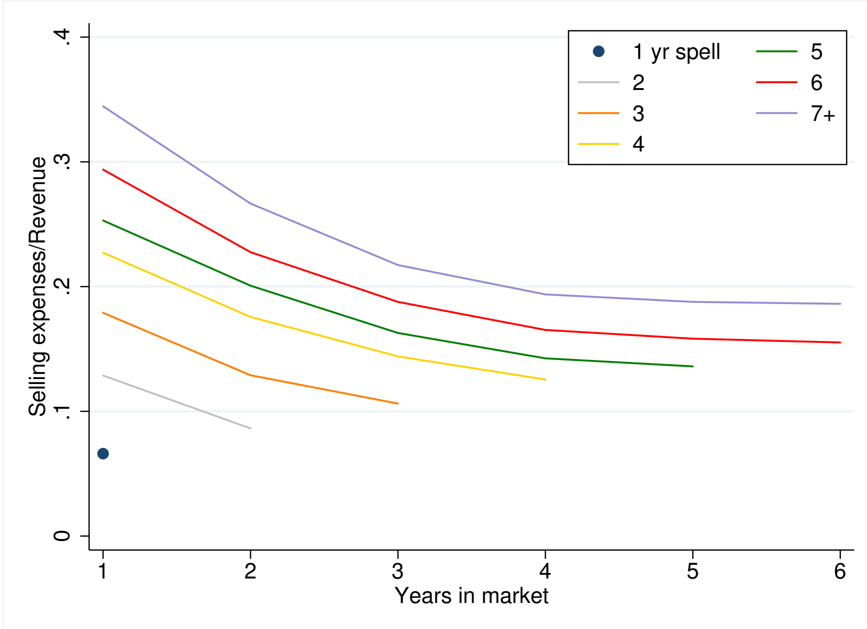

Notes: Figure shows average ratio of selling expenses to revenue predicted by the model for export spells of different length. Source: Authors' calculations.

Figure 8: Simulated impact of trade liberalization on export behavior

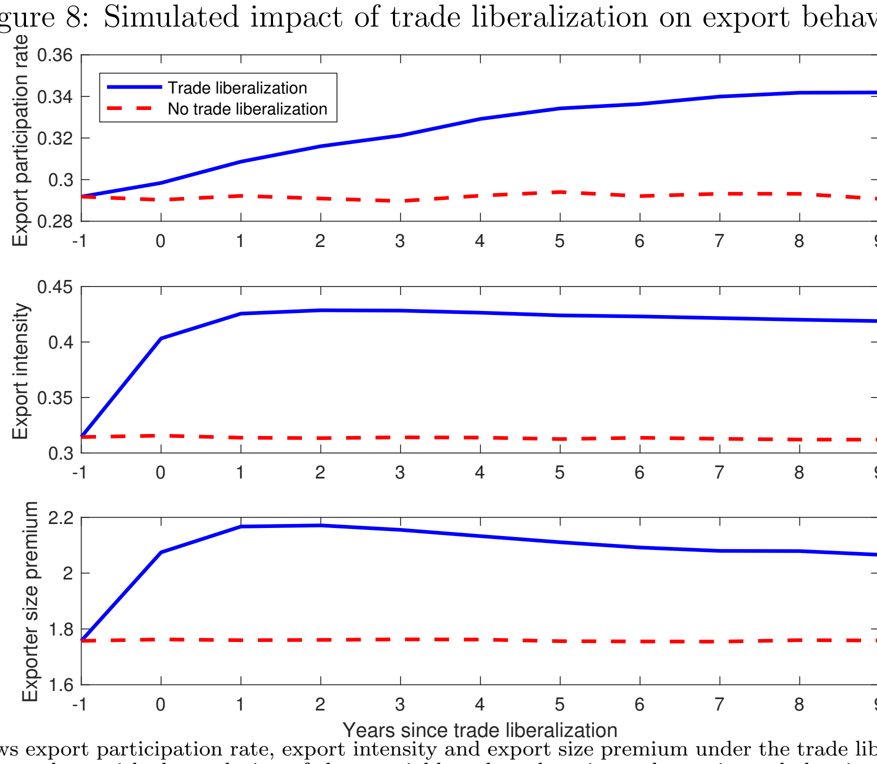

Notes: Figure shows export participation rate, export intensity and export size premium under the trade liberalization exercise described in the text, along with the evolution of these variables when there is no change in trade barriers. Source: Authors' calculations.

Figure 9: Simulated impact of trade liberalization on measured TFPR

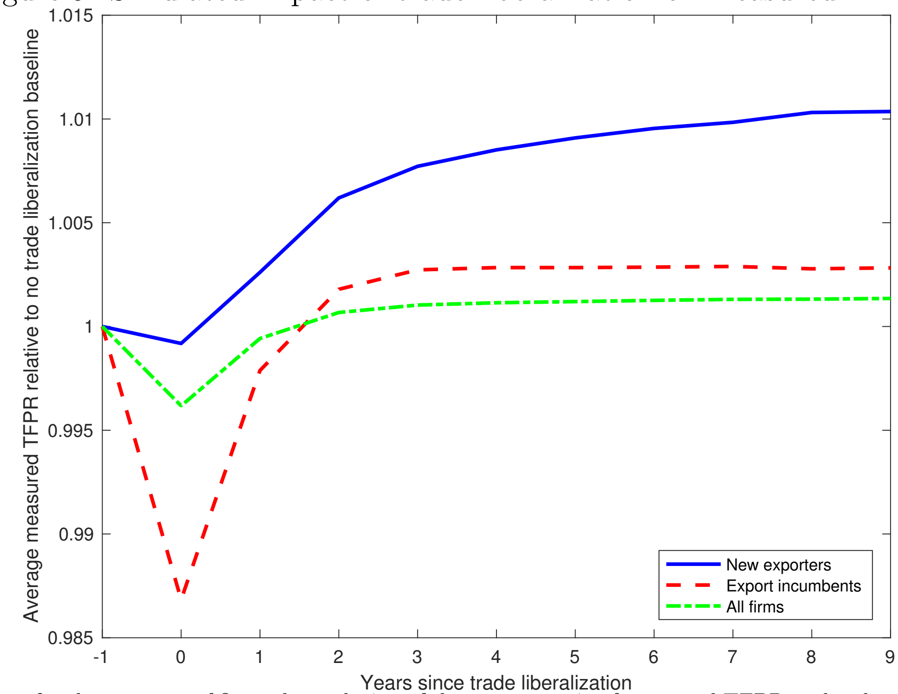

Notes: Figure shows for three groups of firms the evolution of the average ratio of measured TFPR under the trade liberalization exercise described in the text to measured TFPR when there is no change in trade barriers. Source: Authors' calculations.