An empirical dynamic model of trade with consumer accumulation[*]

Paul Piveteau[†]

April 2017

Abstract

Sunk entry costs have been identified as the main export barrier by standard dynamic models of trade. However, these large entry costs are inconsistent with the existence of many small new exporters with low survival rates in foreign markets. To reconcile these patterns, this paper develops a dynamic structural model of trade in which firms slowly accumulate consumers in foreign markets. Estimating the model using export data from individual firms and a Markov chain Monte Carlo estimator, the model correctly predicts lower survival rates for new exporters and estimates much lower entry costs of exporting - less than a third of those estimated in the absence of consumer accumulation. These results have important implications at the aggregate level. In contrast to the standard model without consumer accumulation, this model correctly replicates important facts regarding the aggregate response of international trade to simulated shocks. Moreover, out-of-sample predictions demonstrate that the model better predicts actual trade responses to an observed shock than the standard model.

[*]First version: October 2015. I am grateful to Amit Khandelwal, Eric Verhoogen, Jonathan Vogel and David Weinstein for their advice and guidance. I also would like to thank Costas Arkolakis, Matthieu Bellon, Chris Conlon, Donald Davis, Jean Jacques Forneron, Juan Carlos Hallak, Colin Hottman, Ildiko Magyari, Thierry Mayer, Antonio Miscio, Ferdinando Monte, Jean-Marc Robin, Bernard Salanie, Gabriel Smagghue, Ilton Soares, Daniel Xu and seminar audiences at Columbia University and Sciences Po Paris for comments and suggestions. Part of this research was conducted while I was visiting the economics departments of the ENS Cachan and Sciences Po Paris, I thank them for their hospitality. I also would like to thank the Alliance program and Columbia University CIBER for financial support and the CNIS and French customs for data access. All remaining errors are mine.

[†]School of Advanced International Studies, Johns Hopkins University. 1717 Massachusetts Avenue NW. Washington DC, 20036. Email: ppiveteau@jhu.edu

1 Introduction

The decision by individual firms to enter into an export market is responsible for most of the variations in aggregate trade flow across destinations and time. For instance, Bernard et al. (2007) estimate that around 80 percent of the decline of international trade with geographical distance is due to reductions in the number of exporting firms (extensive margin) rather than changes in exports within the firm (intensive margin). Therefore, understanding the determinants of export decisions and the barriers that firms face in foreign markets is critical.

Standard dynamic models of trade that quantify the nature of these trade costs, such as Das, Roberts, and Tybout (2007), highlight the prevalence of large sunk entry costs as barriers to trade. These large entry costs are necessary to explain the persistence in export decisions, the so-called hysteresis of exporters. However, the prevalence of these entry costs is incompatible with important characteristics of new exporter dynamics that have been recently documented in the literature: most new exporters start small and only a small fraction survive and expand in these foreign markets.

This paper introduces inertia in consumers' choices into a dynamic empirical model of trade to reconcile the observed hysteresis in exporting decisions and the dynamic features of new exporters. I introduce this inertia through the existence of a stock of consumers that firms accumulate throughout their experience in foreign markets. To assess the importance of this accumulation of consumers on exporter dynamics, I develop a Markov Chain Monte Carlo (MCMC) estimator that allows me to include other sources of persistent heterogeneity at the firm level such as productivity and product appeal, and estimate the model using export data from individual French firms. The estimated model correctly predicts lower survival rates for new exporters, but also estimates low sunk entry costs of exporting. On average, entry costs are less than a third the value of those estimated in a model without consumer accumulation. These results have important implications regarding the aggregate predictions of the model: aggregate trade responds slowly to shocks and the contribution of the extensive margin is larger in the long run than in the short run. Both of these patterns have been recently documented in the literature; however, they are inconsistent with the standard model.

I start by presenting three stylized facts about exporters that highlight the importance of growth in demand in these exporter dynamics. Consistent with recent studies, sales and survival rates of young exporters are low upon entry, but grow at a fast rate during the first years of exporting. Moreover, this growth is not due to variations in prices during the life of an exporter, but instead, prices tend to also increase on average with export experience. This result suggests that the growth in sales observed in the years following entry into a foreign market is mainly driven by an increase in the demand shifts received by exporters.[1]

Based on these findings, I develop an empirical dynamic model of trade in which consumers only buy from a limited set of firms, which generates inertia in their consumption choice.[2] Therefore, each firm has a different stock of consumers, depending on its history in the foreign

[1] This finding is consistent with recent papers that show the importance of demand characteristics as source firm heterogeneity (Hottman et al., 2016; Roberts et al., 2012).

[2] This extends to a dynamic setting the consumer margin first introduced in international trade by Arkolakis (2010). This inertia could be alternatively modeled with habits formation or other sources of state-dependence in demand.

market, which shapes its profit, expectations, and decisions in each market. This addition to the model has two important consequences on the dynamics of exporters: first, it implies that new exporters start with low levels of sales and profits when entering a new destination. As they survive and accumulate consumers their sales and profits increase, inducing increasing survival rates with their experience in a destination. Second, because current sales are a source of customer acquisition, firms have incentives to reduce their price to foster the accumulation of new consumers.[3]

In order to study the importance of this mechanism on exporter dynamics, I structurally estimate this model using customs data from France. I perform this estimation on the wine industry, which has the double advantage of being an important exporting industry in France, while also being composed of single-good producers. The dataset provides sales and quantities exported by individual firms on each destination market, which allows me to account for several sources of persistent heterogeneity across firms and destinations. In addition to heterogeneity in demand across destinations, the model identifies three types of heterogeneity at the firm-level: product appeal, defined as a demand shifter that is common across destinations;[4] productivity, acting as a cost shifter; and the firm's consumer base, which is identified from within-firm demand variations across destinations. Because this large number of persistent unobservables complicates the estimation, I take advantage of recent results from the statistical literature to develop a Markov Chain Monte Carlo (MCMC) estimator that accounts for this unobserved heterogeneity through particle filtering, and facilitates the solution of the dynamic problem of the firm.[5] Therefore, this estimator delivers value estimates of the entry and per-period fixed costs of exporting, which are identified by rationalizing the actual entry and exit patterns of exporters on the different export markets.

The estimation results demonstrate the importance of consumer accumulation to replicate exporter dynamics. The introduction of state dependence in demand improves the model's ability to fit the dynamics of young exporters: the model can rationalize lower survival rates for young exporters, as well as the growth of sales and survival rates as exporters become more experienced. Moreover, estimated entry costs of exporting are small relative to existing estimates. The average cost to start exporting to a foreign European destination for a wine exporters is around 34 000 euros, less than the average revenue in these destinations.[6] To confirm this result, I estimate a version of the model without consumer accumulation and obtain an estimate of the average entry cost to European destinations of 114 000 euros, more than three times the estimates of the full model. The reason for this finding is simple: as the model accounts for the fact that it takes years for firms to grow in foreign markets and become successful, large entry costs become unnecessary to rationalize the small fraction of exporters.

These results have important implications at the aggregate level. In particular, the model generates trade responses to trade shocks that are consistent with patterns documented in the

[3] Recent empirical evidence for this type of mechanism on domestic market was found by Foster et al. (2016) who studied the behavior of new firms producing homogeneous goods.

[4] Khandelwal (2010) at the product level or Hottman et al. (2016) at the micro level, also define appeal or quality as the demand shifter after controlling for prices.

[5] To my knowledge, this is the first paper to apply particle filtering within a MCMC algorithm, to account for persistent unobservable heterogeneity at a microeconomic level.

[6] Or equivalently, 3 times the median yearly revenue on these destinations.

literature. First, the model predicts a slow increase in trade as a response to a permanent positive trade shock: because of the slow accumulation of consumers, it takes time for existing and new exporters to expand and reach their new optimal stock of consumers. As a consequence of these adjustment frictions, the trade response is larger in the long-run than the short-run. In my simulations, the trade elasticity is around 1.5 after a year, but equals 6 in the long-run, consistent with the discrepancies documented in the international trade and international macroeconomics literatures. Second, the model can predict the increasing contribution of the extensive margin throughout this trade expansion, as recently documented by Kehoe and Ruhl (2013) and Alessandria et al. (2013). While they enter small in the foreign markets after the shock, new exporters record larger growth than established exporters in the following years, hence increasing their relative contribution to trade growth throughout these years.

Finally, I employ out-of-sample predictions to further confirm the importance of this consumer accumulation in explaining firms' response to shocks. During the sample period, large variations in exchange rates led to a decrease of the exported values and market shares of French wine in the Brazilian market.[7] Based on these variations in exchange rates that affected the relative price of French wine, I construct variations in aggregate demand for French wine from Brazilian consumers. This aggregate demand, in conjunction with outcomes from the model estimated on other destinations, allows me to generate predictions on entry, sales and prices in the Brazilian market, and compare them to the actual realizations of these variables. The model with consumer accumulation is able to replicate, unlike the standard model, the decrease in total trade and in the number of exporters. The decrease in estimated entry costs between the two models, reduces the option value of exporting. Therefore, as economic conditions fluctuate, the model with consumer accumulation (and low entry costs) can predict larger inflows and outflows of exporting firms, and therefore larger variations in total trade.

This paper is closely related to the literature investigating exporter and firm dynamics. Das, Roberts, and Tybout (2007) is the first study to quantify entry and per-period fixed costs of exporting by estimating an entry model of trade. Their estimation emphasizes the importance of entry sunk costs to explain the hysteresis of export decisions.[8] My paper builds on their contribution by capturing this hysteresis through state dependence in demand rather than sunk entry costs, and demonstrating the importance of this extension for a number of micro and macro-level facts. Many recent studies have documented and studied the specific dynamics of new exporters. Nguyen (2012), Albornoz et al. (2012), Berman et al. (2015) and Timoshenko (2015) emphasize the role of demand uncertainty and experimentations to explain exporter dynamics, while Rauch and Watson (2003) and Aeberhardt et al. (2014) develop models where exporters need to match with foreign customers in order to trade. Foster et al. (2016), Fitzgerald et al. (2015) and Rodrigue and Tan (2015) also introduce consumer accumulation to explain the post-entry growth of firms in domestic and foreign markets respectively. However, they do not study the participation decision in these markets. Similar to my paper, Eaton et al. (2014) also develop an entry model with accumulation of customers: they use an importer-exporter matched dataset

[7]The Brazilian devaluation in 1999 and the depreciation of the Argentinian peso in 2002, that fostered Argentina exports to Brazil, have increased the relative price of French wines.

[8]Lincoln and McCallum (2015) similarly shows the prevalence of entry costs when estimating fixed costs of exporting for US firms.

to estimate an empirical model in which exporters grow through the search of foreign distributors and as they learn their own ability.[9] However, while they do not allow for other margins of firms' growth in foreign markets, my model features other sources of time-varying heterogeneity at the firm level, such as productivity and product appeal. Therefore, I am able to investigate the importance of this new margin on exporter dynamics, and its consequences on the estimation of trade costs and the predictions of aggregate trade movements.

This article is also related to macroeconomic papers that similarly introduce a consumer margin, or study aggregate trade dynamics. Arkolakis (2010, 2015) develops a static framework in which a consumer margin at the firm level generates convex costs of participation to foreign markets and heterogeneous elasticities of trade in the cross section of firms. I extend this consumer margin to a dynamic setting to empirically investigate its consequences on exporter dynamics. Drozd and Nosal (2012) and Gourio and Rudanko (2014) show how convex adjustment costs of market shares can explain several puzzles in international macroeconomics and adjustments along the business cycle. Moreover, several recent papers have investigated the reasons for the slow response to trade, and the discrepancy between short and long-run trade elasticities.[10] This series of papers develops macroeconomic models to explain this discrepancy between elasticities through the role of entry and exit of firms, the importance of establishment heterogeneity or the existence of export-specific investment (Alessandria and Choi, 2007, 2014; Alessandria et al., 2014). My paper also explains this discrepancy by combining the role of consumer accumulation at the firm-level, and the entry of new exporters. However, whereas I do not develop a calibrated general equilibrium model, I estimate an entry model using micro-data and a full-information estimator to discipline the role of this mechanism and investigate its consequences on aggregate trade dynamics.

Finally, this study heavily builds on the literature related to the estimation of dynamic discrete choice models (DDCM). These models display a high level of nonlinearity and therefore require the development of specific techniques to facilitate their estimation. Rust (1987) and Hotz and Miller (1993) can be cited as seminal papers in the development of these techniques. More specifically, I employ a MCMC estimator recently developed by Imai et al. (2009) and Norets (2009), to solve the full solution of the DDCM.[11] Moreover, I use particle filtering to account for unobservable heterogeneity, following recent results from Andrieu, Doucet, and Holenstein (2010).[12]

In the next section, I present stylized facts about the trajectories of exporters that emphasize the importance of demand in exporter dynamics. In section 3, I build an empirical model of export entry that is consistent with these facts. I present the estimation method in section 4, and show the results of the estimation on a set of French wine makers in 5. Finally, section 6 inspects the aggregate implications of the estimated results through simulations and out-of-sample predictions, and section 7 concludes.

[9] See also Akhmetova and Mitaritonna (2012) and Li (2014) that show the importance of demand uncertainty, and Aw et al. (2011) about the impact of R&D activities on exporters.

[10] See Ruhl (2008) for a review on the discrepancy between trade elasticities in the international macro and international trade literature.

[11] An application of this method in Industrial Organization can be found in Osborne (2011).

[12] See Flury and Shephard (2011) for an application to economic models and Blevins (2015) using classical estimation.

2 Stylized facts about exporters dynamics

In this section, I present three important facts about exporters' dynamics using French customs data. First, new exporters have low survival rates upon entry, but survival increases quickly with experience. Second, exported values grow with age in foreign markets, even after controlling for survival. Third, prices also increase with exporters' age.

These facts are consistent with the empirical model I will present in the next section: first, the high level of attrition across age will require the model to account for endogenous selection. Moreover, the rise in sales, while prices increase on average, indicates that this growth is driven by a positive shift in the demand schedule of the firm: the consumer margin introduced in the model will be able to replicate this increase as exporters start small, and accumulate consumers with experience. Finally, the low mark-up charged by young firms to foster this accumulation will explain the observed increase in prices with age.

2.1 Data

The dataset used in this paper is provided by the French customs services. It records yearly values and quantities exported by French firms from 1995 to 2010.[13] Annual trade flows are disaggregated at the firm, country and eight-digit product category of the combined nomenclature (CN). This dataset is used to present stylized facts about new exporters in this section, and a restricted sample from the wine industry will be used to conduct the structural estimation described in the next sections. I perform a number of procedures to improve the reliability of the data. In particular, I correct for the existence of a partial-year bias, which overestimates the growth rate during the first exporting year,[14] and clean the dataset to improve the reliability of unit values. Appendix A describes more precisely this cleaning procedure.

Table 1 provides some information on the distributions of the number of observations along different dimensions. Similarly to what have been documented in the literature, trade flows from France are sparse across firms and destinations. This is true for firms across destinations or product categories in a given year, since the median exporting firm records two flows per year, usually concentrated within one product category or one destination. But this sparsity also appears across time as shown in the second panel of Table 1: contrary to the idea that exporting is a long-lasting activity, we can see that the median exporting spell lasts one year.[15] This is true even when exports are aggregated across product categories and exporting flows defined at the firm-destination level.

These statistics provide an overview of the prevalence of short and frequent export flows in the the export data. In order to further investigate this aspect and understand the evolution of the other characteristics of these exporting flows, I specifically look at their trajectories across ages in the next subsections.

[13] This dataset records most of the exporting and importing flows of Metropolitan French firms: there exists thresholds under which a firm does not need to report its exporting activity (In 2001, these thresholds were 1,000 euros for exports to countries outside of the European union, and 100,000 for the total trade within the EU.)

[14] See Berthou and Vicard (2015) and Bernard, Massari, Reyes, and Taglioni (2016) for papers investigating the extent and consequences of this bias.

[15] An exporting spell is defined as a set of consecutive yearly exporting flows between a firm and a foreign destination, or a 8-digit product category - firm pair and a foreign destination.

Table 1: Description of the data

Statistics	mean	p5	p25	p50	p75	p95	N
# observations							
by firm-year	10.2	1	1	2	5	35	716 398
by firm-CN8-year	2.53	1	1	1	2	9	2 887 653
by firm-dest-year	2.83	1	1	1	2	9	2 582 508
Spells duration (years)							
firm-dest-CN8 level	1.87	1	1	1	2	6	3 898 682
firm-dest level	2.24	1	1	1	2	8	1 151 712

Notes: CN8 denotes an eight-digit category after normalization. An exporting spell is defined as a set of consecutive yearly export flows to the same destination.

2.2 Specifications

To describe the trajectories of exporters upon entry, I look at the variation of their survival rates, sales and prices across different ages in foreign markets. I define the age of a firm-product-destination triplet as the number of years this firm has been successively exporting this product category to a market, a market being defined as a 8-digit product category-country pair. I regress the variables of interest (dummy for survival, logarithm of sales or prices) on a full set of age dummies. The specification is augmented with fixed effects that control for the large heterogeneity that exists across industries, destinations and years. Formally, indexing a firm by f, a destination by d, a product category by p, and a year by t, the econometric specifications are the following:

$$Y_{fpdt} = \sum_{\tau=1}^{10} \delta_\tau \mathbb{1}(\text{age}_{fpdt} = \tau) + \mu_{pdt} + \varepsilon_{fdt}, \tag{1}$$

where age_{fpdt} is defined as the number of consecutive years a firm f has been selling the product p to destination d. Y_{fpdt} will be the logarithm of export sales, the logarithm of prices (unit values),[16] or a dummy equal to one if the firm is still exporting to the market the following year. μ_{pdt} is a market×year-specific fixed effect such that the variations that identify the coefficients δ_τ comes from variations across firms of different ages, within a given destination×product category×year pair.

Trade data at the firm-product level are known to have a very large level of attrition. These low levels of survival, especially in the early years of exporting, imply that firms surviving 10 years differ substantially from firms who recently started to export. Consequently, the variation captured by these regressions when comparing old and new firms mostly comes from a selection effect comparing different set of firms, rather than changes across ages for a given set of firms. In order to partially account for this dynamic selection, I also present results when only looking at firm-product-destination triplets that survive 10 years in their specific markets. Even though this only partially accounts for selection, since surviving firms are also firms with specific trajectories, it will show that the observed relationships are not only due to dynamic selection, but also appear within a constant set of firms.[17]

[16] I use the terms unit values and prices interchangeably throughout the paper. As usual with this type of dataset, prices are obtained by dividing export values by export quantities.

[17] Another possibility to partially account for this dynamic selection would be to use firm-product fixed effects, or first difference transformations. These transformations would control for the heterogeneity across firms but

2.3 Results

Here I present three important facts about exporters, namely the growths of the survival rates, exported values, and prices with export experience in foreign markets. Regarding the growth of sales and survival rates, these facts have been extensively documented and discussed in the literature in international trade and macroeconomics,[18] and I show that these facts still hold after controlling for the partial-year bias highlighted by Berthou and Vicard (2015) and Bernard et al. (2016). However, the increase in prices has not been documented using a comprehensive trade dataset,[19] even though Foster, Haltiwanger, and Syverson (2016) documents similar patterns for the domestic prices of homogeneous goods, and Macchiavello (2010) show evidence of similar trajectories for prices of Chilean wine in the UK market.

Fact 1: Survival rates are low for new exporters, and strongly increase with their age
First of all, the probability to survive on a market, i.e. to export on this market the following year, is very low for the average exporter. Figure 1 displays the average survival rate for a firm-product pair on a foreign market, for different age or experience levels. For an exporter in its first year, the probability to export the following year is roughly 35 percent. However, this survival probability rapidly increases once exporters have survived several years: this rate is larger than 50 percent at age 2, and close to 75 percent at age 6. This result reflects the idea highlighted in the previous section that most export spells are short lived. These low, yet

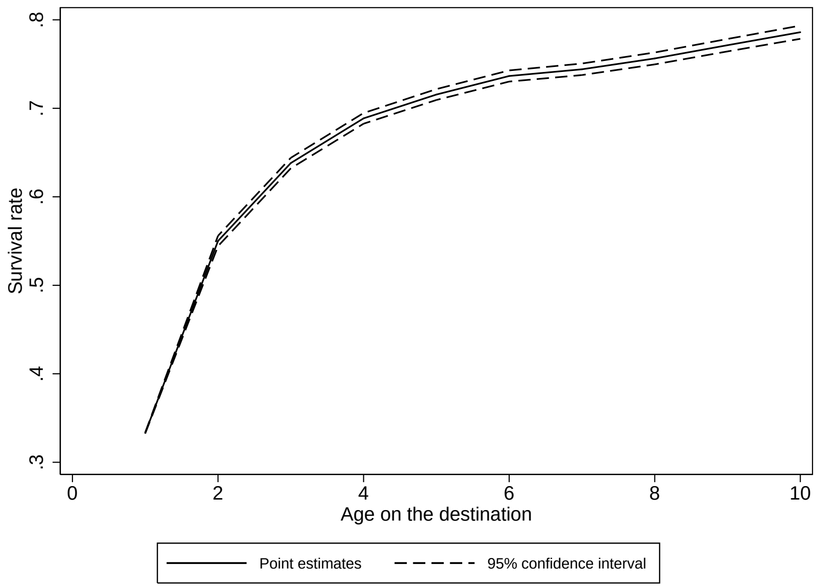

FIGURE 1: Survival rates across export ages

Notes: The figure reports the average survival rate of a firm-product category pair in a destination at different ages. Standard errors clustered at the firm-product-destination level.

rise other issues regarding identification and selection. I discuss these related specifications in appendix B.

[18] See for instance Ruhl and Willis (2008) for a presentation of these facts and puzzles.

[19] Simultaneously to the redaction of this paper, several other manuscripts have documented prices patterns using trade data: see in particular Rodrigue and Tan (2015) and Fitzgerald et al. (2015).

increasing, survival rates have theoretical and methodological consequences. On the theoretical side, it will be important to have a model of export entry that can replicate and explain these low survival rates: a model in which entry costs are prevalent will have difficulties explaining why so many firms exit the export market so rapidly. On the methodological side, these very low survival rates imply it will be necessary to account for this large attrition when interpreting differences across firms in a reduced form exercise, and to model this entry decision in the design of the structural model.

Fact 2: Exported values increase with firm age in a destination, even more so in the first years of exporting Turning to the variation of sales across ages, Figure 2 documents the large growth rates of exported values across ages. This figure is obtained by plotting the results from regression (1), after normalizing the average log sales at age one to be zero. When comparing exported values, exporters which are in their third year of exporting will export more than twice as much compared to a new exporter. This difference reaches an order of 7 when comparing an exporter with 10 years of experience to a new exporter. However, it is important to note that these differences are mostly due to a strong selection across exporters: old exporters, who by definition managed to survive on foreign markets, were initially larger than the average new exporter. The right panel in Figure 2 emphasizes this point by looking at the relationship when restricting the set of exporters to those surviving 10 years. Accounting for survival, the growth rate of sales with export age is strongly reduced. Nevertheless, surviving exporters still record an average growth rate of 25 percent between ages one and two. Moreover, this growth appears to continue the first six years: at this age, exporters tend to be on average twice larger compared to their first year of exporting.

In conclusion, we observe substantial growth rates of sales during the first years of exports. These growth rates are large but appear to be lower than previously described in the literature because of the correction for the partial-year effect highlighted in Berthou and Vicard (2015) and Bernard et al. (2016). Moreover, this positive relationship appears to be robust across product categories and destinations. However, it is important to emphasize that this growth could be generated by the stochastic nature of the exporting process: by focusing on surviving firms, we are looking at the "winners" of the exporting game, which could explain unusually large growth rates. Accounting for this potential mechanism will be one of the roles of the structural model introduced in the next section.

Fact 3: Export prices increase with firm age in a destination, even more so when controlling for survival. One possible explanation for the growth in sales could be productivity improvements that lead to a reduction in the prices of the good exported, and therefore an increase in its sales. On the contrary, it appears that prices also increase with the experience of the firm on the export market.

Figure 3 reports the estimated parameters of regression (1) in which the average price at age one is normalized to zero. The left figure shows that the price of an exporter with 10 years of experience is on average 9 percent higher than the price of a new exporter. Similar to sales, this effect could come from a selection effect of the exporting activity: a selection process driven by the quality of the product for instance, would imply that older firms which managed to survive,

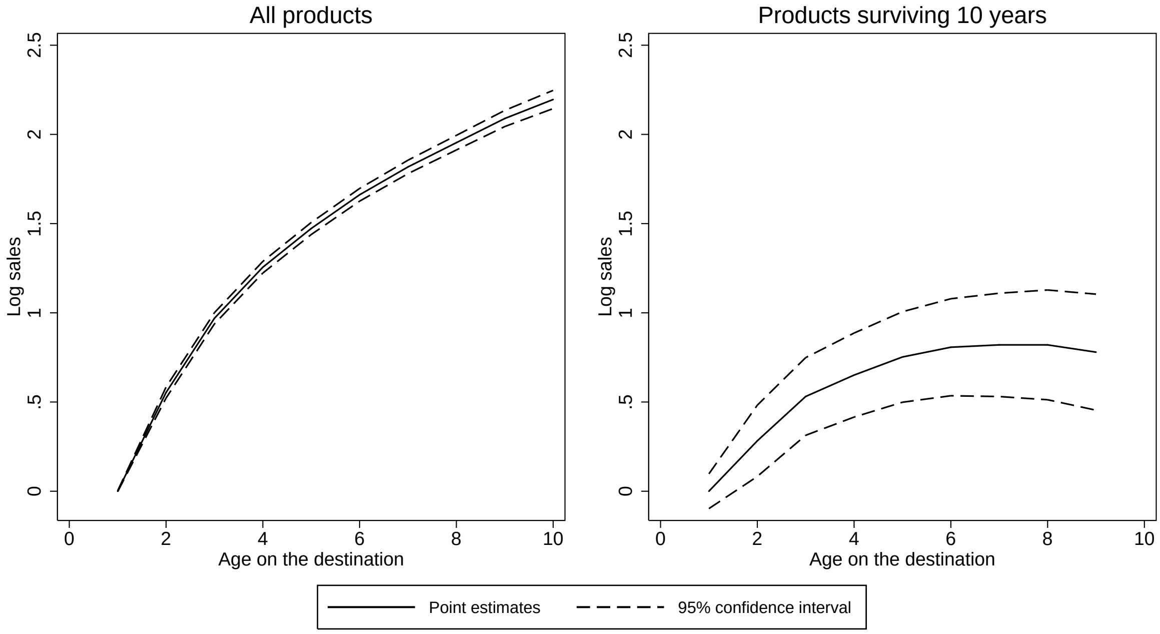

FIGURE 2: Sales across export ages

Notes: The figure reports the cumulative growth of sales of a firm-product category pair in a destination at different ages. Left panel uses the entire sample, while the right panel only uses firms that reach age 10. Standard errors clustered at the firm-product-destination level.

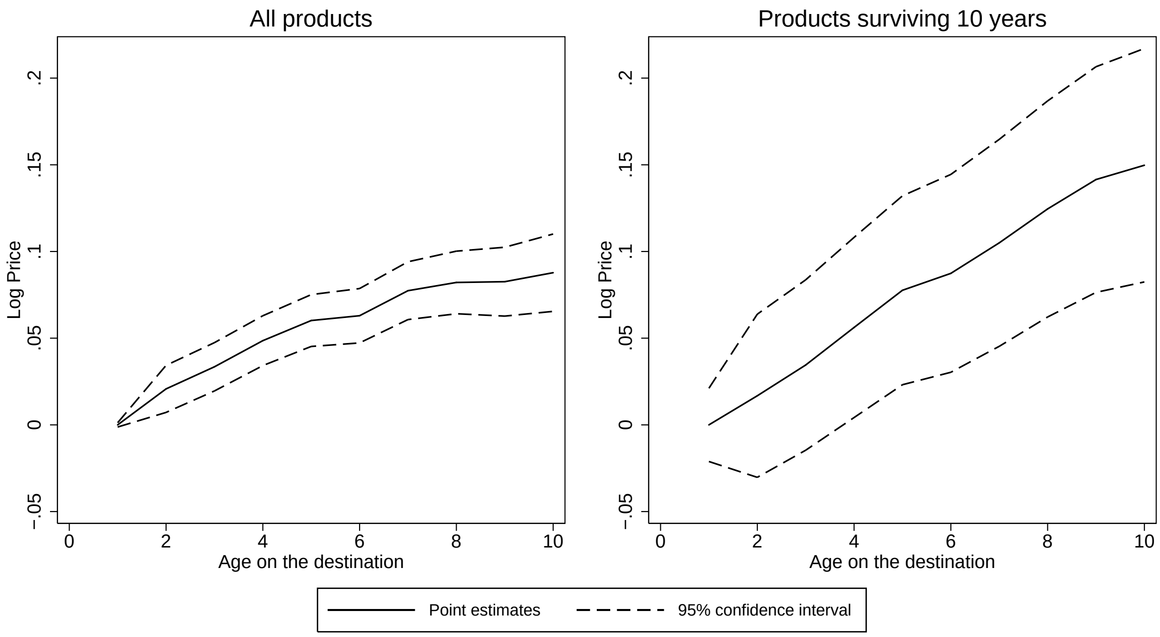

FIGURE 3: Prices across export ages

Notes: The figure reports the cumulative growth of prices of a firm-product category pair in a destination at different ages. Left panel uses the entire sample, while the right panel only uses firms that reach age 10. Standard errors clustered at the firm-product-destination level.

have higher prices than young exporters. However, when controlling for selection by looking at surviving firms (Right panel of Figure 3), it appears that the growth of prices is even larger compared to the regression using the full sample: the price after 10 years appears to be in average 15 percent larger than the price charged by the same firm at age one.

Observing a larger growth of prices when looking at a constant sample of firms has two important implications. First of all, it means that costs are the main driver of the selection process: high price firms tend to disappear more in the first years such that the positive correlation between prices and age is weakened when using the full sample. Second, it implies that this positive correlation cannot be only driven by dynamic selection. Therefore, an additional mechanism is necessary to explain why firms tend to increase their price during their exporting life. The structural model presented in the following section will introduce such a mechanism, through the dynamic pricing of the firms.

There exists other methods that can partially account for the endogenous selection across ages. However, within variations within firm-product-destination cannot be used in this context as it is not possible to separately identify the role of experience, cohort and trend effects. An alternative specification could include firm-product fixed effects such that the variation is obtained from the same firm-product pair which is selling to different destinations, with different ages. Using this specification raises potential issues from the endogenous sorting across destinations, but leads to similar, although weaker, results. Detailed results are provided in appendix B.[20]

This section introduced simple facts about exporters' dynamics that will guide the empirical model developed below. We can draw three conclusions from these figures. First, survival rates are very low in export markets and grow with the age of the firm, which is contradictory with a world where the main barrier to export is made of sunk entry costs: in such a world, exporters would tend to keep exporting once they have overcome this important barrier. Second, sales of exporters grow rapidly in the first years of exporting. These large growth rates are also present when accounting for dynamic selection across firms. Third, this increase in sales is driven by a growth in the demand of the firm: price variations cannot explain this large increase, implying the importance of demand characteristics as main drivers of this increase in sales. On the contrary, it appears that prices tend to rise with age, even more so when controlling for dynamic selection. This pattern could be explained by a dynamic behavior of the firms that foster their growth in the early years by reducing their prices.

Despite these conclusions, it is difficult to make strong causal statements by comparing firms of different ages. This brings to light a second benefit of developing and estimating a structural model to study the entry and growth of exporters: in addition to understanding the dynamic decisions of firms, it allows the model to control for the endogenous sorting and attrition of firms, and recover the different processes that drive the observable variables of the model.

[20] Recent papers have discussed this correlation between prices and experience in export markets. In particular, Fitzgerald et al. (2015) find similar results but favors a specification identifying prices trajectories across destinations, hence controlling for marginal costs, which leads to the absence of significant positive relationship between prices and age. I also find a positive but insignificant relation using this identification (see figure 15 in appendix B). Importantly, the main result of the structural estimation, regarding the importance of demand frictions when estimating fixed costs, does not require increasing prices but non decreasing prices. Moreover, since I do observe increasing prices in the cross-section of firms, and given existing evidence of these patterns using more precise micro-data (see Foster et al. (2016) or Macchiavello (2010)), I develop a model that can account for these prices dynamics.

3 Structural model of export entry

This section describes an empirical model of entry into foreign markets in which the accumulation of consumers creates a new source of dependence in the dynamic problem of the firm. This model aims to identify the different sources of firms' profit explaining export decisions. Therefore, it is crucial to allow for heterogeneity across firms and destinations, but also to allow this heterogeneity to be persistent over time. Indeed, persistent heterogeneity is the main competing hypothesis to sunk entry costs to explain persistence in export decisions. As a consequence, the model features two additional sources of persistence at the firm level - productivity and product appeal - and one persistent characteristic specific to destinations - their aggregate demand. Therefore, a potential profit for a firm-destination pair depends on four characteristics: productivity, product appeal, aggregate demand and consumer share.[21]

The introduction of consumer accumulation implies two deviations from the standard dynamic model: first, firms start small in a new market. Their sales and profit will rise in the following years as they accumulate more consumers. Second, because part of this consumer accumulation comes from sales, firms have dynamic incentives to lower their prices to foster their future demand.

I start by describing the demand schedule of the firm and how the accumulation of consumers affects the demand from foreign destinations. After introducing the costs associated with the production process, I solve the dynamic problem of the firm to study the consequences of this consumer margin on entry and pricing decisions.[22]

3.1 Demand

A large literature in industrial organization has found empirical evidence of inertia in consumption and state dependence in demand. This literature also points out the large number of mechanisms that can generate this dependence in demand, as well as the difficulty to empirically disentangle these different channels.[23] In order to keep the model tractable, I introduce state dependence in demand through the existence of a firm-specific customer base in each destination. This customer base, denoted n_{fdt}, describes the share of consumers, in a destination d at time t, that includes the product f in their consideration set,[24] which is consistent with the idea of customer margin introduced in the macroeconomic and international trade literature.[25]

Therefore, I assume that new exporters have an initial share of consumer n_0 when they enter a foreign destination. In the subsequent years, the consumer awareness of the products propagates through two mechanisms. First, the sales of a product increase its awareness in the next period.

[21] I assume that entry decisions are independent across destinations, once controlling for firms' characteristics. McCallum (2015) provides support for this assumption. See also Morales et al. (2014) which uses moments inequalities to handle such a large state space.

[22] Note that I do not study the choices made by the firms for each product it produces. Firms are seen as single-good producers, and will be considered as such in the empirical application.

[23] One can cite habits in consumption, costly search, or imperfect information as mechanisms leading to state dependence in demand (see for instance Dubé, Hitsch, and Rossi (2010) for a paper distinguishing and measuring the contribution of these different mechanisms).

[24] The marketing literature defines a consideration set as the set of products that consumers consider when making purchase decisions. See for instance Shocker et al. (1991).

[25] See for instance Drozd and Nosal (2012) and Gourio and Rudanko (2014) for macroeconomic papers, and Arkolakis (2010) in international trade.

Specifically, an euro increase in the sales of a product increases by η_1 the potential share of consumers in the next period. This mechanism can arise in situations in which consumers have imperfect information about product characteristics, and therefore use sales as a signal for the expected utility gain from consuming a good.[26] Second, consumer accumulation also comes from word-of-mouth: I assume that each aware consumer share its awareness with η_2 consumers. Both of these mechanisms generate a potential growth in the share of consumers for the firm. However, because some of these reached consumers are already aware of the existence of the product, this acquisition of new consumers is discounted by a factor $(1 - n')^\psi$ with $\psi > 0$, such that the marginal effect of sales s and consumer share n on the future share n' are

$$\begin{aligned} \frac{\partial n'}{\partial s} &= \eta_1 (1 - n')^\psi, \\ \frac{\partial n'}{\partial n} &= \eta_2 (1 - n')^\psi \end{aligned} \tag{2}$$

This specification is largely inspired from the marketing literature described in Arkolakis (2010): the accumulation of consumers has decreasing returns such that it is more difficult for an established firm to reach new consumers. For this firm, a significant share of reached consumers are already part of its consumer share, hence not contributing to its growth. Therefore, the parameter ψ describes the importance of these decreasing returns, while parameters η_1 and η_2 characterize the importance of the two sources of accumulation.

Importantly, these two margins of growth generate different optimal responses by the firm. In a world with word-of-mouth, where consumers learn from their neighbors, the growth of the consumer share can be seen as exogenous, only based on the past share of consumers. In this world, firms cannot affect this accumulation with their pricing decisions.[27] However, in a world where consumers face uncertainty regarding product characteristics and sales are seen as a signal, firms have incentives to reduce their price in order to foster the accumulation of consumers.[28]

Adding an initial condition to these differential equations, $n(0, 0) = \underline{n}$, we obtain the following law of motion for the consumer share of a firm f, at date t and destination d:

$$n_{fdt} = 1 - \left[(1 - \underline{n})^{1-\psi} - \eta_1 (1 - \psi) s_{fdt-1} - \eta_2 (1 - \psi) n_{fdt-1} \right]^{\frac{1}{1-\psi}} \tag{3}$$

Therefore, the share of consumers today n_{fdt} depends on the sales s_{fdt-1} and the share of consumers n_{fdt-1} in the previous period in this market.

This share of consumer acts as a demand shifter for the firm since it scales the demand firms receive from each destination. Moreover, we assume that each consumer displays CES preferences over its consideration set. Denoting Ω_i the consideration set of a consumer i, its utility function

[26] With CES preferences, the amount spent for a specific good is proportional to the utility gain obtained from the consumption of this good.

[27] This model does not take into account advertising as a source of growth. Trade datasets do not provide information on advertising expenditures, which makes it difficult for an empirical model to account for this channel.

[28] This distinction echoes differences between structural and spurious structural dependences (Heckman, 1981), that generate different optimal responses by firms.

is

$$U_i = \left[\sum_{f \in \Omega_i} \exp\left(\frac{1}{\sigma}\lambda_f\right) q_{if}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \qquad \sigma > 1,$$

where q_{if} is the quantity consumed of good f and λ_f the appeal of the product. This consumer i maximizes this utility function given a budget y_i devoted to this set of goods, and prices $\tilde{p}_f$. As solution of this optimization, the quantities q_{if} demanded by consumer i for a good f are

$$q_{if} = \begin{cases} \exp(\lambda_f)\tilde{p}_f^{-\sigma} P^{\sigma-1} y_i & \text{if } f \in \Omega_i \\ 0 & \text{if } f \notin \Omega_i \end{cases}$$

where P is the CES price index faced by the representative consumer.[29] Aggregating the demand of individual consumers from each destination d and time period t, the demand received by firm f from destination d at time t is:

$$q_{fdt} = q(\lambda_{ft}, X_{dt}, n_{fdt}, p_{fdt}, \varepsilon^D_{fdt}) = n_{fdt} \exp(\lambda_{ft} + X_{dt} + \varepsilon^D_{fdt}) p_{fdt}^{-\sigma} \tag{4}$$

where X_{dt} captures all the aggregate variables of the demand shifter,[30] p_{fdt} is the factory price of the good, and ε^D_{fdt} is a random demand shock.

It is important to note that the appeal of the product λ_{ft} does not vary across destinations. Given the existence of an aggregate demand shifter, this implies that firms cannot vary the relative quality or appeal of their good across destinations. Therefore, this specification can still explain that firms provide different product appeal in different destinations, as long as these differences are common across firms. This assumption is fundamental to explain the identification assumption of the model: while λ_{ft} and X_{dt} are respectively firm and destination specific, the customer share n_{fdt} will be identified through the sales of a firm in a specific destination. After describing the demand faced by firms, I now turn to the costs associated with production and international trade.

3.2 Technology and costs

The costs associated with production and international trade are similar to those traditionally assumed in the literature. I first describe the constant marginal costs of production, then the fixed costs associated with the exporting activity.

First, I assume constant marginal costs of production. These marginal costs are a decreasing function of the firm productivity ϕ_{ft}, and vary with the appeal of the product λ_{ft}. Moreover, I assume the existence of non-persistent productivity shocks ε^S_{fdt}, and I allow costs to vary with destination markets by including a set of coefficients γ_g.[31] Formally, the marginal cost function

[29] With different sets of goods, each consumer has a different price index. However, I follow Arkolakis (2010) by assuming that each consumer has probabilistically an equivalent set of goods, such that all consumers face the same price index $P = \left[\sum_{f\in\Omega} n_f \exp(\lambda_f)\tilde{p}_f^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$.

[30] $X_{dt} \equiv \log Y_{dt} - (1-\sigma)\log P_{dt} + (1-\sigma)\log(\tau_{dt}e_{dt})$ where $Y_{dt} \equiv yN_{dt}$ are total expenditures from a number of consumers N_{dt}, and τ_{dt} and e_{dt} are respectively iceberg transportation costs and exchange rates that converts the factory price to the consumer price.

[31] Destination markets will be geographically divided in three groups indexed by g.

is

$$c_{fdt} = c(\phi_{ft}, \lambda_{ft}, \varepsilon^S_{fdt}) = \exp(-\phi_{ft} + \alpha\lambda_{ft} + \gamma_g + \varepsilon^S_{fdt}) \tag{5}$$

In addition to these production costs, I assume that firms need to pay entry and per-period fixed cost for each destination they respectively enter or export to. These fixed costs are defined as follows

$$FC(\mathcal{I}_{fdt-1}, \nu_{fdt}) = \begin{cases} fc_g + \nu^c_{fdt} & \text{if} \quad \mathcal{I}_{fdt-1} = 1 \\ fe_g + \nu^e_{fdt} & \text{if} \quad \mathcal{I}_{fdt-1} = 0 \end{cases}$$

where $\mathcal{I}_{fdt}$ is a dummy that equals one if the firm f is active (records positive sales) in destination d at time t, and ν_{fdt} are random shocks on fixed costs. Note that these fixed costs will vary across groups g of destinations.[32] Moreover, I assume that shocks ν^c_{fdt} and ν^e_{fdt} follow a logistic distribution with respective variance parameters σ^c_ν and σ^e_ν. These shocks allow the model to rationalize all observed decisions made by the firms.

3.3 Profit and value function

From the demand received by the firm, and the costs of production, I derive the potential profit of the firm for each destination market. After describing the timing of a typical period, I define the entry problem of the firm, and the associated value functions. This dynamic problem depends on five variables that define the state space of the problem: the exogenous variables - product appeal λ, productivity ϕ and aggregate demand X - the share of consumer n, and the presence in the market in the previous year $\mathcal{I}_{-1}$.

In this model, firms decisions are limited: they decide whether to be active on the market, and the price they charge if they decide to export. Consequently, the appeal of the product, the productivity and the aggregate demand from each destination are exogenous but persistent variables that potentially capture the hysteresis of the exporting decisions. For ease of exposition, I denote these variables $\xi \equiv (\lambda, \phi, X)$ such that, ignoring the subscripts and the parameters of the model, the profit function of a firm is

$$\begin{aligned} \Pi(\xi, n, p, \varepsilon, \mathcal{I}_{-1}, \nu) &= q(\xi, n, p, \varepsilon^D) \left[p - c(\xi, \varepsilon^S)\right] - FC(\mathcal{I}_{-1}, \nu) \\ &= \pi(\xi, n, p, \varepsilon) - FC(\mathcal{I}_{-1}, \nu) \end{aligned}$$

This profit function is made of a variable profit and fixed costs. Despite a CES demand, this variable profit could be negative because of the dynamic nature of the pricing decision of the firm: some firms could set a price lower than their marginal costs to foster future demand. The second part of the profit function comes from the fixed costs of exporting $FC(\mathcal{I}_{-1}, \nu)$ that depend on the past presence of the firm on the market and the profit shock ν, which allow the empirical model to explain entry and exit decisions of firms that cannot be rationalized by variable profits.

This profit is earned if the firm decides to be active on the market at this period. To study this decision, figure 4 defines the timeline of a typical period, which provides the timing at which

[32]For instance I assume that fixed costs are equal for all European destinations.

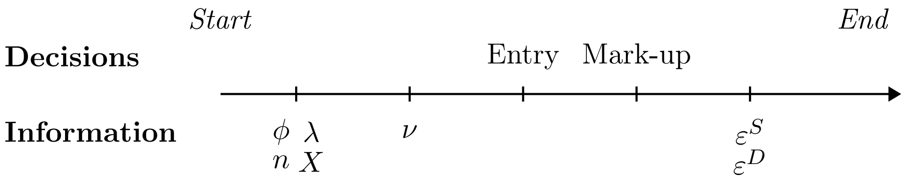

FIGURE 4: Timeline of a period

decisions are made and the information sets available to the firms when they take decisions.

As described in figure 4, the firm observes at the beginning of the period its exogenous variables, λ, ϕ, n and X. After realization of the profit shock ν, it decides whether to export in the market. If the firm decides to export, it optimally chooses the mark-up to charge over their marginal costs.[33] Finally, sales and prices are obtained after observing the realization of the non-persistent shocks ε.[34]

Therefore, denoting μ the multiplicative mark-up of the firm such that $p = \mu c$, the value function of the firm can be defined as the following:

$$
\begin{aligned}
& V(\xi, n, \mathcal{I}_{-1}) = E_\nu \max \Big\{ V_I(\xi, n) - FC(\mathcal{I}_{-1}, \nu) \;\; ; \;\; V_O(\xi) \Big\} \\
\text{with} \qquad & V_I(\xi, n) = \max_{\mu} \Big\{ E_\varepsilon \Big\{ \pi(\xi, n, \mu, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1) \Big\} \Big\}, \\
& V_O(\xi) = \beta EV'(\xi, n_0, 0), \\
& EV'(\xi, n', \mathcal{I}) = \int_{\xi'} V(\xi', n', \mathcal{I}) \, dF(\xi'|\xi).
\end{aligned}
$$

For each market, the firm chooses between exporting $V_I(\xi, n) - FC(\mathcal{I}_{-1}, \nu)$ and being inactive $V_O(\xi)$. By being inactive, the firm makes no profit today but retains the possibility to update its decision in the next period. In contrast, when exporting, it obtains a present profit that depends on the shocks ε and the mark-up chosen by the firm. Moreover, the firm has a continuation value, $EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1)$, characterized by a stock of consumer n' and lower fixed costs to pay in the next period. This continuation value is constructed from the transition of the exogenous variables $F(\xi'|\xi)$, and the expected value of $V(\xi, n', I)$.

3.4 Firms' decisions: entry and pricing.

After defining the problem of the firm, I can now derive the optimal entry and pricing decisions of the firm. Because the accumulation of consumers is based on the sales of the firm, the optimal price charged by the firm deviates from a standard constant mark-up. Instead, firms optimally reduce their mark-up to account for the accumulation of consumers. Because this pricing decision is taken once the firm has decided to enter, I start by describing the optimal mark-up charged

[33] Choosing the mark-up rather than the price facilitates the computation of the solution, while allowing for structural shocks ε in demand and costs.

[34] Theses assumptions are mostly driven by the construction of the empirical model. The realization of the shock ν before the entry decisions allow the model to rationalize all entry decisions. Similarly, the realizations of the shocks ε after the markup decision generate structural errors that can explain observed sales and prices variations. Another advantage of this timing assumption is that the mark-up chosen by the firm generates exogenous movements in prices, which allows me to identify the price-elasticity of demand.

by the firm. By backward induction, I then infer the expected profit of the firm and solve for the value and probability of exporting.

Optimal mark-up The firm's choice of mark-up is made after entry, in order to maximize the sum of the present profit and the continuation value of exporting:

$$V_I(\xi, n) = \max_{\mu}\left\{ E_\varepsilon \left\{ \pi(\xi, n, \mu, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1) \right\} \right\}$$

such that the optimal price chosen by the firm is:

$$p(\xi, n) = \frac{\sigma}{\sigma - 1} \frac{1}{1 + \beta \int \omega(\varepsilon) \eta_1 (1 - n')^\psi \frac{\partial EV'(\xi, n', 1)}{\partial n'} dF(\varepsilon)} c(\xi, n) \tag{6}$$

with $\omega(\varepsilon) = \frac{\exp(\varepsilon^D + (1-\sigma)\varepsilon^S)}{\int \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) dF(\varepsilon)}$.[35] The optimal mark-up charged by the firm has two components. First, the firm applies the standard CES mark-up $\frac{\sigma}{\sigma-1}$ based on the price-elasticity of demand. Second, the firm applies a discount factor based on the dynamic incentives it has to lower its price to attract more consumers in the future. This factor depends on two elements: first, how much this increase in sales increases its consumer share tomorrow, $\eta_1(1-n')^\psi$; this element induces lower mark-ups for small or young firms that benefit from higher returns of accumulation. Second, the extent of this discount also depends on the impact of this increase in the future consumer share on the continuation value $\frac{\partial EV'(\xi, n', 1)}{\partial n'}$. This effect is not linear but hump shaped with the profitability of the firm:[36] young firms that are unlikely to survive have no incentives to invest in future consumers, while firms that can use extra consumers to increase their survival probability get the largest benefits from increasing their consumer share. Finally, note that this equation defines the unique optimal price charged by the firm but only through an implicit function, since the future share n' depends on the price charged by the firm.[37]

Consequently, the accumulation of consumers implies heterogeneous mark-ups across firms, depending on their current share of consumers, and their expectations on future profits. Having described the optimal mark-up of the firm, I can now study its entry decision.

Entry condition Knowing the expected option values of being active or inactive, I can now study the entry decision of the firm. Firm pick the most profitable option, after observing the shock ν that affects the fixed costs of being active on a market. From the logistic distribution, the expected value of the firm before observing the shock ν is

$$V(\xi, n, \mathcal{I}_{-1}) = \sigma_\nu \log \left[\exp\left(\frac{1}{\sigma_\nu} (V_I(\xi, n) - f) \right) + \exp\left(\frac{1}{\sigma_\nu} V_O(\xi) \right) \right] \tag{7}$$

in which f and σ_ν can be respectively fe or fc and σ_ν^e or σ_ν^c, depending on the value of $\mathcal{I}_{-1}$. This equation closes the dynamic problem of the firm, by providing the fixed point that defines

[35] See appendix C for derivations.

[36] This comes directly from the probability of exit that makes the value function of the firms convex for low profitability firms, and concave for higher profit firms.

[37] In the estimation procedure, the dynamic problem of the firm is solved on a grid. Therefore, I do not use this formula to find the optimal mark-up but instead, pick the mark-up in the grid that maximizes the value function of the firm. See appendix D for details.

the value function $V(\xi, n, \mathcal{I}_{-1})$. Moreover, the probability for a firm to be active, before the realization of the fixed cost shock ν, is,

$$P(\mathcal{I} = 1|\xi, n, \mathcal{I}_{-1}) = \frac{1}{1 + \exp\left(-\frac{1}{\sigma_\nu}(DV(\xi, n) - f)\right)} \tag{8}$$

with $DV(\xi, n) = V_I(\xi, n) - V_O(\xi)$. This last equation predicts the probability of entry of a firm, conditional on its current characteristics, described by ξ, n and $\mathcal{I}_{-1}$. While n and $\mathcal{I}_{-1}$ are endogenous, ξ are exogenous and unobservables variables. Therefore, to finish the derivation of the model, it is necessary to describe the evolutions of these exogenous variables across time. These evolutions will be important to compute the expectation of the value functions, $EV'(\xi, n, \mathcal{I}_{-1})$, as well as disciplining the variations of sales and prices across times in the empirical application.

3.5 Evolution of exogenous variables

Most of the hysteresis in exporting decisions is likely to come from the persistence over time of firms characteristics Therefore, it is necessary to allow these processes to be time-varying and persistent. Therefore, I assume that exogenous variables follow AR(1) processes, with flexible parameters, such that:

$$\begin{aligned}
\lambda_{ft} &= \rho_\lambda \lambda_{ft-1} + \sigma_\lambda \varepsilon^\lambda_{ft} \\
\phi_{ft} &= \mu_\phi + \rho_\phi \phi_{ft-1} + \sigma_\phi \varepsilon^\phi_{ft} \\
X_{dt} &= \mu_{Xg} + \rho_X X_{dt-1} + \sigma_X \varepsilon^X_{dt}
\end{aligned} \tag{9}$$

where the ε shocks follow a normal distribution with zero mean and unit variance. Note that, by normalization, λ is centered around zero: since both X and λ enters linearly in the demand function, it is not possible to separately identify their respective means. Moreover, because X_{dt} describes the aggregate demand from a destination d, I allow the mean μ_{X_g} of this process to change across destinations group to capture different trends in aggregate demand.

Finally, I need to impose distributional assumptions on the initial conditions of these unobservables. I assume that the distributions of product appeal and productivity are stable over time such that the initial distributions are constrained by a stationary assumption. However, I assume that the variation in aggregate demand across destinations does not arise from a stationary distribution, such that $X_{d0} \sim N(\mu_{X_0}, \sigma_{X_0})$. Finally, I assume that the initial share of consumers follow a Beta distribution with parameters 1 and 5.[38]

3.6 Restricted model

In order to assess the importance of consumer accumulation on estimated trade costs and aggregate response to trade, I also estimate a restricted version of the model that does not feature this mechanism. This restricted model is equivalent to assuming that exporters have a consumer share n_{fdt} equal to one when they are active on the market. As a consequence, firms do not have incentives to deviate from the CES pricing, and mark-ups are similar across all firms.

[38] This matters for firms that record positive sales the year before the beginning of the sample. Given the small number of firms in this case, and the length of the panel I use (14 periods), this assumption has no consequence on the estimation.

This restricted version of the model can be seen as the canonical model used in the literature. In this model, firm-level heterogeneity and entry costs of exporting explain the hysteresis in exporting. This model can be seen as a dynamic version of Melitz (2003), as estimated by Das et al. (2007). Estimating this restricted model is essential to assess the importance of consumer accumulation on the outcomes of the estimation and the aggregate implications of the model.

4 Estimation

In this section, I describe the procedure to estimate the parameters of the model. I start by describing the likelihood of the problem, based on the three structural equations linked with the observable variables (sales, prices and participation to export). I then turn to the algorithm, showing the advantages of a MCMC estimator to account for persistent and unobserved heterogeneity and solve the dynamic problem of the firm. Finally, I provide intuition behind the identification of parameters and unobservables of the model.

4.1 Likelihood

I start by presenting the likelihood that is obtained from the three main equations of the model: the price and demand equations that feature the stock of consumers and the dynamic mark-up charged by the firm, and the entry probability that describes the exporting decision on each destination.

First of all, the demand and price equations (4), (5) and (6) are taken in logarithm to obtain

$$\begin{aligned}\log s_{fdt} &= \log n_{fdt} + \lambda_{ft} + X_{dt} + (1-\sigma)\log p_{fdt} + \varepsilon^D_{fdt}\\ \log p_{fdt} &= -\phi_{ft} + \alpha\lambda_{ft} + \log\mu(\xi, n_{fdt}) + \gamma_d + \varepsilon^S_{fdt}\end{aligned}$$

This block constitutes the first part of the likelihood. Assuming that ε follows a bivariate normal distribution with variance Σ, I define this likelihood block as $L_\varepsilon(s_{fdt}, p_{fdt}|\xi_{fdt}, n_{fdt}; \Theta)$, with Θ being the full set of parameters, such that

$$\begin{aligned}L_\varepsilon(s_{fdt}, p_{fdt}|\xi_{fdt}, n_{fdt}; \Theta) = G_\Sigma\Big(&\log s_{fdt} - \log n_{fdt} - \lambda_{ft} - X_{dt} - (1-\sigma)\log p_{fdt}\\ &; \log p_{fdt} + \phi_{ft} - \alpha\lambda_{ft} - \log\mu(\xi_{fdt}, n_{fdt}) - \gamma_d\Big)\end{aligned} \tag{10}$$

where G_Σ is the density function of a bivariate normal distribution with means zero and variance matrix Σ.

The second block of the likelihood is based on the entry decision of the firm. Equation (8) defines the probability to enter for a firm, based on its set of unobservables ξ, its stock of consumer n and its past exporting activity. I denote this function $L_\nu(\mathcal{I}_{fdt}|\xi_{fdt}, n_{fdt}, I_{fdt-1}; \Theta)$ that is obtained from the binary choice made by the firm

$$\begin{aligned}L_\nu(\mathcal{I}_{fdt}|\xi_{fdt}, n_{fdt}, \mathcal{I}_{fdt-1}; \Theta) = &\left[1 + \exp\left(\frac{1}{\sigma_\nu}\left(-DV(\xi_{fdt}, n_{fdt}) + f\right)\right)\right]^{-\mathcal{I}_{fdt}}\\ &\times\left[1 + \exp\left(\frac{1}{\sigma_\nu}\left(DV(\xi_{fdt}, n_{fdt}) - f\right)\right)\right]^{\mathcal{I}_{fdt}-1}\end{aligned} \tag{11}$$

where function $DV(\xi_{fdt}, n_{fdt})$ and f are defined as previously. Therefore the total likelihood for a given observation $D_{fdt} \equiv \{s_{fdt}, p_{fdt}, \mathcal{I}_{fdt}\}$ is the product of the two densities $L_{\varepsilon}(.)$ and $L_{\nu}(.)$.

To obtain the unconditional likelihood, that does not depend on the unobservables, it is necessary to integrate out this set of unobservables. Since these unobservables are persistent over time, the likelihood of the entire dataset D is obtained by repeatedly integrating the unobservables from period T to 0:

$$L(D|\Theta) = \int_{n_{-1}} \int_{\xi_0} ... \int_{\xi_T} \prod_{f,d} L(D_{fdT}|D_{fdT-1}, \xi_{fdT}) \times ... \times L(D_{fd0}|D_{fd-1}, \xi_{fd0}, n_{fd-1})$$
$$dF(\xi_{fdT}|\xi_{fdT-1}) \times ... \times dF(\xi_{fd0})\, dF(n_{fd-1})$$

where $F(\xi_{fd0})$ and $F(n_{fd-1})$ are defined by the initial unobservables density function, and D_{fd-1} the observables previous to the estimation sample. After describing the likelihood of the problem, I now turn to the estimation procedure to obtain the posterior distribution of parameters Θ.

4.2 Algorithm

To estimate the model, I develop a Markov Chain Monte Carlo (MCMC) estimator to tackle the two important difficulties in evaluating the likelihood: integrating the numerous integrals and solving the dynamic problem of the firm.[39]

In order to circumvent these difficulties, I employ a MCMC estimator, taking advantage of recent Bayesian techniques to sample the posterior distribution of the parameter Θ, conditional on the data. The choice of a Bayesian estimator relies on two recent methods developed in the Bayesian literature. First, I employ a particle filter to perform the integration of the unobservables. In particular, I follow recent methods described in Andrieu, Doucet, and Holenstein (2010) that use particle filtering to update the set of unobservables through a Gibbs sampler. This sampling technique allows me to develop a MCMC estimator in which parameters and unobservables are alternatively sampled conditional to each other.[40] Second, to overcome the computational burden of solving the value functions in the likelihood, Imai, Jain, and Ching (2009) and Norets (2009) show how to take advantage of the iterative feature of the MCMC estimator, by only updating the value functions in the Bellman equation once at each iteration. The intuition is that the value function can be approximated at the early stages of the Markov chain: by using this value function as initial value in the next iteration, the value function will converge toward the fixed point defined by the contraction mapping as the Markov chain converges and explores the posterior distribution of Θ.

Overall, the MCMC estimator explores the posterior distribution of the parameters Θ. This distribution is proportional to the product of the likelihood and the prior distribution such that

$$P(\Theta \,|\, D) \propto \int_{\xi} L(D \,|\, \xi, \Theta)\, dF(\xi \,|\, \Theta)\, P(\Theta) \tag{12}$$

[39]The literature on dynamic discrete choices model, starting from Rust (1987), is mostly devoted to this second problem. This problem can be largely simplified using the mapping between conditional choice probabilities and value functions, as highlighted in Hotz and Miller (1993). However, in my application, state variables are not observed, hence complicating the estimation of conditional choice probabilities.

[40]I describe in details in appendix D the estimation procedure.

To avoid the influence of priors in the parameters estimation, I assume flat priors except for values of parameters that do not satisfy theoretical or stationarity constraints.[41] Therefore, the goal of the Markov Chain is to repeatedly sample from the posterior distribution according to (12). Given the large number of parameters (32), this is achieved by sequentially updating blocks of parameters and unobservables: in particular, I divide parameters in three blocks. One consists of the parameters from L_ε with those of the law of motion of n(.), one of the parameters of the different laws of motion of ξ, and a final block consists of the dynamic parameters from L_ν. Therefore, a typical iteration of the Markov Chain consists of the sampling of these three blocks of parameters, conditional to unobservables, the sampling of unobservables conditional to parameters, and an iteration of the value function according to its definition in (7).

Two important points are worth point out regarding the algorithm. First, the value functions that allow the computation of the objects $DV(.)$ and $\mu(.)$ are obtained on a grid that is updated throughout the algorithm. The specific values of $DV(.)$ and $\mu(.)$ are then obtained by interpolation to be evaluated at any point in the state space.[42] Second, due to the complexity of the estimation procedure, I do not estimate the value of β, the discount rate of future periods. This parameter is difficult to identify in dynamic discrete choice models and I therefore set its value to 0.95.

After describing the details of the estimation procedure, I provide, in the next section, intuition about the sources of identification of the parameters and the unobservables.

4.3 Identification intuition

Despite the complexity of the algorithm, estimating this model using micro data and a full information estimator provides simple intuitions of parameters' identification.

To describe the sources of identification, it is useful to distinguish the identification of unobservables and parameters. Let's assume first that the parameters are known. In this case, the identification of the unobservables mostly come from a variance decomposition of the demand shifters and prices: the demand shifter is decomposed between a firm-year component (the product appeal λ_{ft}), a destination-year component (the aggregate demand X_{ft}), and a firm-destination-year component (the consumer base n_{fdt}). Once the product appeal is known, the productivity ϕ_{ft} is identified from price variations across firms. Therefore, the identification of the unobservables mostly comes from a decomposition of observables variables, which is straightforward if the parameters are known. Moreover, the hierarchical structure and the entry decisions bring additional information to identify the posterior distribution of these unobservables. For instance, if a firm is not exporting one year, the information from previous and future years will help identify the potential value of the unobservables. Similarly, if a firm only exports to one destination at a given year, the fact that it does not export somewhere else provides information about its product appeal or productivity.

Turning to parameters identification, they can be divided in three groups. The identification of the 17 parameters related to the laws of motion of the unobservables can be easily identified

[41] I exclude from the support of Θ (or equivalently assigned a prior probability of zero for these values), negative values for the variance parameters, as well as values beyond -1 and 1 for the autocorrelation parameters. I also impose the fixed costs parameters (f, fe) and the parameter ψ to be positive.

[42] I provide extensive details in appendix D about the implementation of the algorithm.

once knowing the values of these unobservables. Regarding the 7 parameters entering the demand and pricing equations, their identification is similar to traditional demand and supply equations: correlation between sales and prices, and prices with destination dummies and product appeal, while the parameters of the variance matrix are obtained from the variance of the unexplained variation in prices and sales. Even though I do not employ instruments for prices, despite their correlation with demand shocks, the price-elasticity is identified by the mark-up term that delivers variations in prices, exogenous to demand shocks.[43][44] Finally, the 8 parameters related to the entry problem are obtained by comparing potential profits and firms' observed decisions: the number of exporters identifies the per-period fixed costs, the persistence in exporting the entry costs, and the remaining variance in exporting decisions identifies the required variance of these fixed costs' shocks.

Consequently, the identification of the unobservables conditional to the parameters, and of the parameters conditional to the unobservables are quite straightforward. The goal of the MCMC estimator is to repeatedly sample each component conditional to the other, in order to obtain their joint distribution. After a necessary period of convergence, the Markov Chain describes the posterior distributions of the parameters from which confidence intervals can be obtained.[45]

5 Results

I implement the estimation on a set of wine exporters from France; the choice of this industry is based on two criteria. First, wine producers only export wine. Therefore, it is reasonable to assume that the entry decisions into foreign destinations are made at the firm level, and it is possible to aggregate sales and prices at the firm level for each destination. Second, the wine industry is a large industry in France and, therefore, I can obtain a large enough sample of exporters with a large set of destinations. In appendix A.2, I describe the specific selection procedure to obtain the estimation sample of 200 firms, and provide statistics to describe this sample.

I start by describing the fit of the model relative to the exporters' dynamics presented earlier. Then I present the estimated values of the parameters, and in particular the decrease in entry costs induced by the introduction of the consumer margin.

5.1 Fit of the model

I report in this section the fit of the model regarding the survival rates, sales and prices of the firm-destination pair at different ages. Figure 5 reports the predictions of the model relative to the data. I also report the results of the restricted version of the model, which does not contain a consumer margin.

As reported in figure 5, the full model with consumer accumulation can reproduce the growth in sales across ages (top left figure). This ability is not surprising as the full model allows for destination-specific growth in sales through consumer accumulation. As the sales of exporters

[43] See Piveteau and Smagghue (2015) which develops an instrument for prices using trade data.

[44] This identification strategy does not apply to the restricted model. Therefore, I set the price elasticity in the restricted model such that both models deliver the same average markup.

[45] Specifically, I perform 60 000 iterations of the MCMC. See appendix D for details.

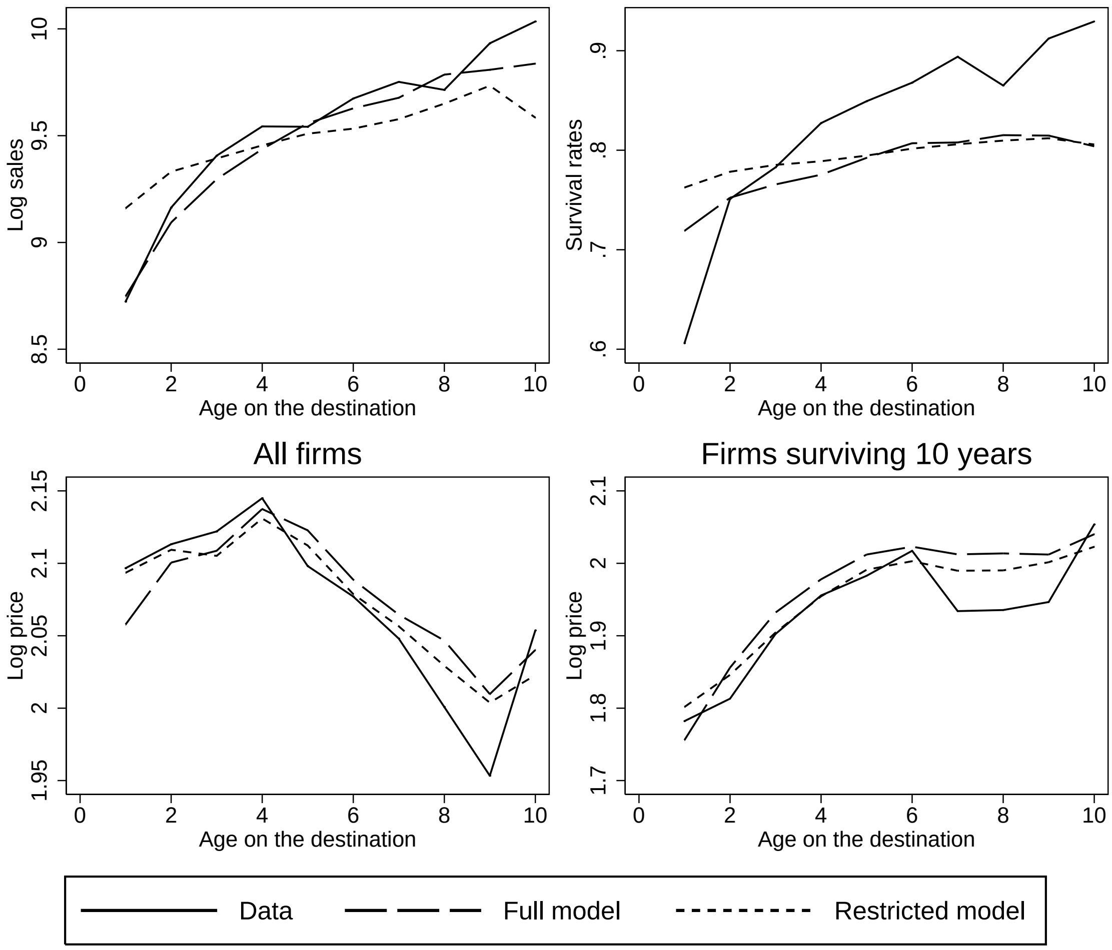

FIGURE 5: Predictions of survival rates, sales and prices across ages.

increase with age, their profit also increases. Consequently, the full model can perform better at explaining lower survival rates for young exporters: the average prediction gives a survival rate of 70 percent the first year, and more than 80 percent after 5 years in the foreign markets. However, this growth in sales is not sufficient to fully explain the low survival rates of young exporters, and, therefore, does not entirely solve the puzzle linked with young exporters dynamics. This negative result comes from the use of a full-information estimator: the large variance in exit decisions in the cross-section of firms induces the model to estimate large variance in fixed costs, hence limiting differences in survival rates across ages.

In comparison, the restricted model cannot explain this rise in sales and even less in survival rates: in the restricted model, the predicted survival rate is constant across ages, between 77 and 80 percent, which is similar to the average survival rate in the sample. However, the predictions on prices appear quite similar across models (bottom figures). Both of them can reproduce the decrease in prices with age. When looking at firms surviving 10 years, we can see that the full model can do slightly better in explaining the rise in price with the age of the firm.[46]

After describing the fit of the model, I now turn to the description of the estimated values of the parameters.

5.2 Estimated parameters

The results of the estimation of the model are reported in table 2. I report for each parameter the median of its posterior distribution, and its 90 percent confidence interval.

First, looking at the law of motion of the consumer margin, we note that the initial share of consumers at entry (n_0) is relatively small, equal to 3 percent, which leaves a large potential for firms to grow through the accumulation of consumers. This growth is driven both by the past sales of the firm (η_1), as well as the past shares of consumers (η_2), since theses coefficients are all significantly larger than zero. Moreover, we can see that the degree of concavity of this law of motion is significant, with a median of the posterior distribution of the coefficient ψ equal to 0.38.

Second, the other unobservables of the model - appeal, productivity and aggregate demand - depict strong degrees of persistence. The coefficients of autocorrelation of the AR(1) processes are estimated to be in average 0.98, 0.83 and 0.81, respectively for the product appeal, the productivity of the firm, and the aggregate demand of the destination. Moreover, the estimated price-elasticity are in the ballpark of estimates found in the IO literature, but smaller than elasticities found in the trade literature. This perfectly makes sense as this price elasticity is defined for a constant set of consumers. This elasticity will be magnified by the adjustment of the consumer base when prices change.[47]

Finally, because I estimate a structural model of entry, the model is able to deliver euro estimates of the entry and per-period fixed costs paid by an exporter. We see that the obtained fixed costs are relatively low, with the estimated entry cost to an European destination being around 34 000 euros.[48] In addition, a firm will have to pay 17 000 euros every year to keep

[46]Prices do not display a significant trend with age in the structural sample. Therefore, it is quite difficult to distinguish and interpret the differences between the two models.

[47]I come back to this point in section 6 when simulating a tariff reduction.

[48]Prices are normalized across years using a national consumer price index, such that the values are expressed

TABLE 2: Estimated parameters

Parameter		**Estimate**	**90% Confidence Interval**	
			Lower bound	*Upper bound*
Continuation fixed costs	Europe	13 634	12 401	14 973
(in euros)	Americas	9 681	9 067	10 334
	Asia/Oceania	12 063	10 938	13 091
Entry fixed costs	Europe	34 044	31 699	35 766
(in euros)	Americas	26 121	24 333	27 895
	Asia/Oceania	30 177	27 822	32 091
Variance of continuation costs	σ_ν^c	10 310	9 469	11 012
Variance of entry costs	σ_ν^e	4 705	4 366	5 247
Law of motion of n	n_0	0.032	0.030	0.033
	$\underline{n}$	0.006	0.006	0.006
	$\eta_1(10^{-5})$	0.265	0.259	0.280
	η_2	0.234	0.230	0.236
	ψ	0.375	0.353	0.383
Law of motion of appeal	ρ_λ	0.98	0.98	0.99
	σ_λ	0.13	0.12	0.15
Law of motion of productivity	ρ_ψ	0.83	0.81	0.86
	σ_ψ	0.08	0.07	0.08
	μ_ψ	-0.25	-0.29	-0.21
Law of motion of agg. demand	ρ_X	0.81	0.81	0.82
	σ_X	0.38	0.37	0.39
	μ_{X1}	2.35	2.31	2.37
	μ_{X2}	2.01	2.00	2.02
	μ_{X3}	2.25	2.22	2.28
	μ_{X_0}	14.03	13.90	14.15
	σ_{X_0}	0.23	0.16	0.35
Elasticity cost of appeal	α	0.92	0.89	0.95
Price-elasticity	σ	2.04	2.03	2.05
Cost dummies	γ_2	0.19	0.17	0.21
	γ_3	0.26	0.24	0.28
Variance matrix	Σ_{11}	1.18	1.15	1.20
	Σ_{12}	0.15	0.14	0.16
	Σ_{22}	0.15	0.14	0.15

exporting to this destination. As an element of comparison, the average export value of a firm in my sample to an European destination is 42 000 euros, while the median value is 13 000. One of the reasons for these relatively low numbers is the small variance parameter of these entry costs' shocks, whose the median of the posterior distribution is 4 705. This low number reflects the ability of the model to correctly predict the entry of firms, such that a large variance of these entry costs' shocks is not necessary to rationalize entry decisions.

In order to confirm the small magnitudes of these entry fixed costs relative to the literature, I compare in table 3 these parameters with the estimates of the restricted version of the model, which does not feature a consumer margin. The comparison between the two models highlights that entry costs are much larger in the version without consumer margin. For instance, the average entry costs to export to Europe jump from 34 000 to 114 000 euros. As a consequence, the variance of these entry shocks also increase between the two models. However, estimates of the continuation fixed costs are roughly similar in the two models.

TABLE 3: Estimated parameters (comparison between models)

Parameter		**Estimates**	
		Full model	*Restricted model*
Continuation fixed costs	Europe	13 634	12 834
	Americas	9 681	15 860
	Asia/Oceania	12 063	17 724
Entry fixed costs	Europe	34 044	114 031
	Americas	26 121	99 823
	Asia/Oceania	30 177	113 110
Variance of continuation costs	σ_ν^c	10 310	34 134
Variance of entry costs	σ_ν^e	4 705	23 160
Elasticity cost of appeal	α	0.92	0.35

This decrease in entry costs is explained by two main factors. First, since the model accounts for the fact it will take time for firms to grow, become successful and make large profits in foreign markets, high entry costs are not necessary to deter entry into these markets. Second, the consumer margin captures some of state dependence in exporting status, reducing the role played by entry costs in explaining the hysteresis in export decisions. This result will be very important when looking at the models' predictions in response to shocks. Estimating large entry costs to export implies a substantial option value of exporting: large entry costs make entering so difficult that firms will hesitate to exit this market in case of adverse shocks. I study these consequences in the next section when comparing the predictions of these models under simulated and observed trade shocks.

Another important difference between these two models arises from the estimates of the cost of appeal. In the full model with consumer margin, appeal is very costly, with a cost elasticity of 0.92. However, the model without consumer margin identifies product appeal with a low impact on prices, with an average estimate of 0.35. This difference shows that the consumer margin, by capturing some of the variance in sales, modifies the definition of product appeal: in the presence

as euros from the year 2000.

of consumer accumulation, product appeal is more related to product characteristics rather than distribution network for instance. As a consequence, this product appeal is more connected to marginal costs than the one identified in the restricted model.[49]

5.3 Other outcomes of the model

I now discuss the evolution with export experience of two important objects introduced in this model: the consumer shares and the mark-up charged by firms. Figure 6 provides the distribution of consumer shares for each age of the firm. Remember that when firms enter, they all have an initial share of approximately 3 percent, which explains why the figure provides distributions from ages 2 to 10. Figure 6 illustrates that the distribution tends to shift toward the right as age increases. One can see that most of the firms have a small consumer share at age 2: only a small fraction of them are larger than 25 percent. However, as age increases, more and more firms reach a larger size. Therefore, at age 10, a significant number of them has a consumer share that is larger than 50 percent. However, there is also a large amount of heterogeneity within ages: some firms are large at ages 2 or 3, but a large fraction of them are still small in terms of consumer shares when reaching years 9 or 10. As a result, the overall distributions appear to flatten as age increases, rather than translate toward the left. This implies that the process of consumer accumulation is not identical across firms, and very much relies on the individual sales of the firm rather than an exogenous increase of consumers with age. Some firms will never reach a large fraction of consumers, because it is not profitable for them to do so.

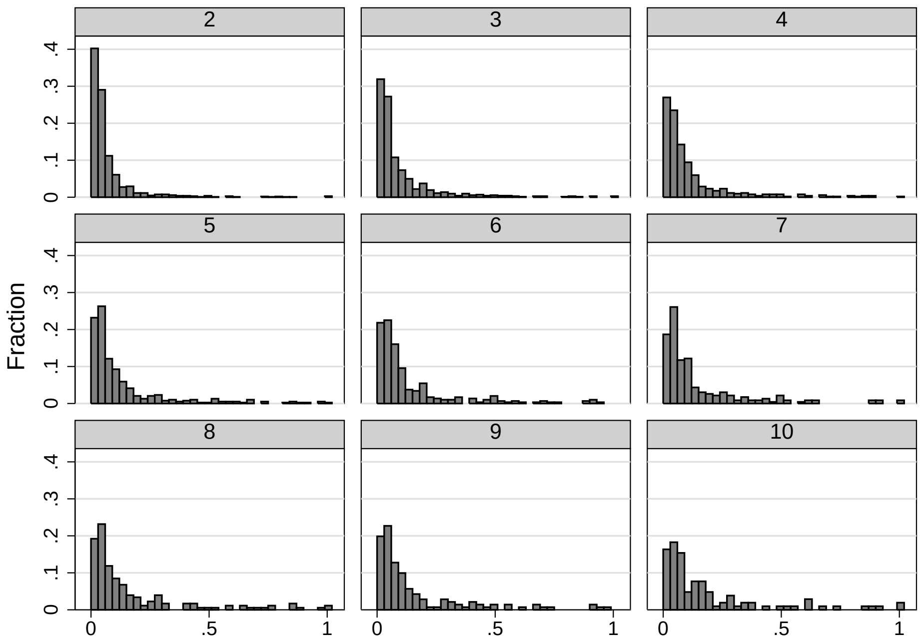

FIGURE 6: Distribution of consumer shares by age

I then turn to the distributions of mark-ups charged by the firms, that can be used by firms to

[49]The full set of parameter estimates for the restricted model can be found in table 9 in appendix E.

foster consumer accumulation. Figure 7 reports the distributions of mark-ups, separately for each age from 1 to 9. One can see that, similar to the consumer shares, there is a large heterogeneity in mark-ups across ages, but also within ages: the model does not imply a mechanical correlation between mark-ups and age. However, we can see that firms tend to price more aggressively at a young age, in comparison to more established firms. The reason is twofold: first, these firms are small and therefore benefit from large returns of higher sales on consumer accumulation. Second, because these firms are small and young, additional consumers are crucial to survive in the following years.

However, it also appears that the pricing behavior does not vary mechanically with ages: some old firms appear to price almost as aggressively as young firms. This result is due to the fact that survival rates do not change that much across ages, hence forcing firms to maintain low prices after a few years to increase their likelihood to survive. Overall, when disentangling the different contributions in prices dynamics, the dynamic pricing behavior leads to a 6% increase in prices between age 1 and 10.

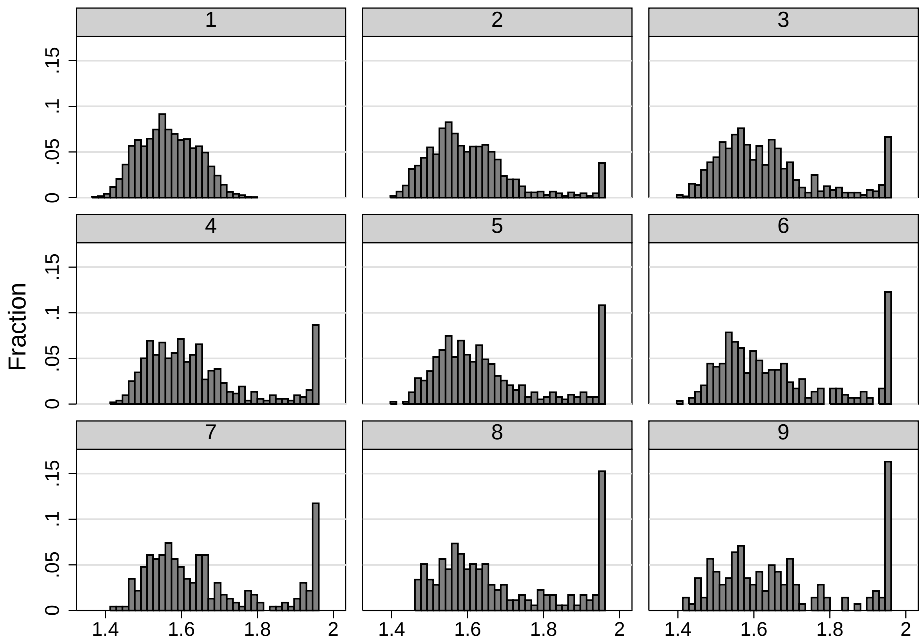

FIGURE 7: Distribution of mark-ups by age

In the next section, I explore the implications of the estimated models on patterns of international trade.

6 Aggregate implications

In this section, I use simulations and out-of-sample predictions to demonstrate the importance of the model regarding aggregate trade responses to shocks. The introduction of the consumer margin generates a sluggish response of trade flows, as it will take time for firms to reach new

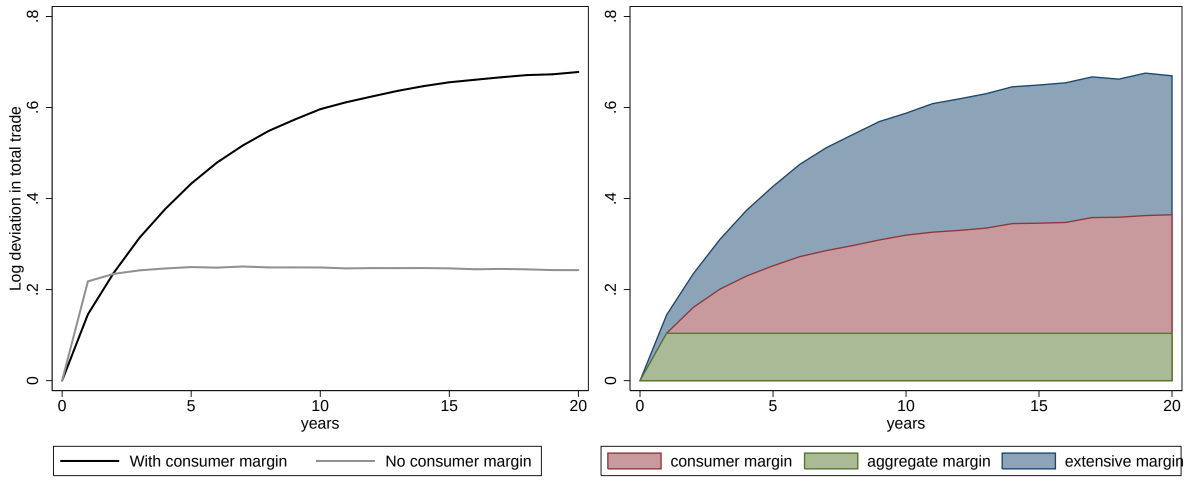

FIGURE 8: Effect of a permanent 10 points tariffs decrease.

consumers. Moreover, low entry costs imply a stronger response of firms' entry and exit to shocks. As a consequence, the model can replicate, in response to a positive trade shock, a discrepancy between the short and long run trade elasticities. Moreover, I show that the model generates out-of-sample predictions that better match the actual behaviors of French exporters in response to exchange rate movements in the Brazilian markets.

6.1 Sluggish trade response

The accumulation of consumers by firms generates frictions in growing on foreign markets. As a consequence, the trade response to shocks will be slow at the microeconomic and aggregate level. This documented pattern can explain the discrepancy between values of the trade elasticity at different horizons. International macro economists use elasticities around 1 or 2 in order to match trade responses to price variations at a high frequency. However, international trade economists use elasticities ranging from 6 to 8 to explain variations in trade flows across countries, or trade responses after a trade liberalization episode.[50]

In order to quantitatively evaluate the ability of the model to generate this discrepancy in trade elasticities between horizons, I simulate a decrease of 10 points on tariffs applied to French exports to the US. I simulate the trajectories of the 200 firms from my sample following this tariff reduction, and compare them to a counterfactual scenario without tariff decrease. I apply this experiment to the full model, as well as the standard model that does not feature consumer accumulation. Figure 8 reports the log-deviation relative to the counterfactual scenario without tariff change, of the total trade to the US, and the decomposition of this growth between different margins.[51]

As we can see from figure 8, the predictions of the two models are significantly different. In the model without consumer margin, trade increases instantaneously as the shock occurs: with lower tariffs, exporters prices decrease and trade increase, and new exporters enter the market. After these first years, no further adjustment occurs. In comparison, the model with consumer

[50] See Ruhl (2008) or Alessandria et al. (2013) for studies of this discrepancy.
[51] The decomposition follows a methodology identical to Hummels and Klenow (2005).

margin depicts a slower adjustment to trade as it takes up to 15 years to observe the full effect of the reduction in tariff. The long-run effect is roughly 4 times the effect recorded after one year with a trade elasticity of 1.5 after one year and 6 in the long run. Therefore, the model with consumer margin can generate trade elasticities that are consistent with those found in the literature, both in the short and long run.

Moreover, the decomposition in the right figure highlights the contribution of each margin in this slow adjustment: in the first year, most of the increase in trade is due to the decreasing tariff that leads to lower prices and larger sales. However, in the following years, both the intensive and extensive margins amplifies the trade elasticity as it takes time for existing and new exporters to reach their optimal number of consumers. Moreover, due to decreasing returns of consumer accumulation, the model can also explain why the extensive margin displays a increasing contribution throughout the trade expansion. This result is consistent with recent findings documented in Kehoe and Ruhl (2013) and Alessandria et al. (2013): in their empirical study, the latter manuscript finds that the contribution of the extensive margin goes from zero to around 50% during a trade expansion following a devaluation episode. By contrast, I find a relative contribution of the extensive margin of respectively 25 and 45 percent in the short and long run.[52]

6.2 Out-of-sample predictions: export response to exchange rate variations in Brazil.

In order to further demonstrate the relevance of the model with consumer margin, I compare its out-of-sample predictions relative to the standard model. I take advantage of additional destinations, that have not been previously used in the estimation, to test the ability of the model to correctly predict the exporting behavior of the French exporters contained in my sample.

I apply this methodology to the Brazilian wine market during my sample period.[53] The choice of the Brazilian market is based on two reasons: first, it is a large market such that a large enough number of French wine producers export to Brazil. Second, the Brazilian wine market has recorded during the sample period two important shocks that affected the Brazilian demand for French wine: the devaluation of the Brazilian currency, the real, in 1999, that has been followed by a strong depreciation of the currency in the following years, and the Argentinian devaluation in 2002, which led to a strong growth in wine export to Brazil. These shocks respectively generated a strong increase in the price of French wines in local currency and an important drop of the price index on the Brazilian wine market.

Therefore, I take advantage of these variations in exchange rates, which can be arguably seen as exogenous to French exporters behavior, as sources of variation in the aggregate demand received by French firms. The model relies on five state variables that characterize the entry and sales of exporters: the appeal λ_{ft} and productivity ϕ_{ft} of the firms, their consumer shares n_{fdt}, the aggregate demand from a destination X_{dt} and their previous export activity $\mathcal{I}_{fdt-1}$. Because the quality and productivity of the firms are common across destinations, I can use the estimated individual qualities and productivities from the estimation procedure. Moreover, the

[52] See figure 17 in appendix E that describes the relative contribution of each margin.

[53] My sample period goes from 1997 to 2010. However, I stop my predictions in 2007, since the trade collapse generated a strong decrease in trade that is difficult to account for.

variables n_{fdt} and $\mathcal{I}_{fdt-1}$ are obtained from the predictions of the model, such that only initial conditions are required for these variables. Therefore, I can construct the aggregate demand X_{dt} from Brazil, using variations in real exchange rates and the Brazilian GDP,[54] and feed this variable in the model to deliver predictions of entry, sales and prices in the Brazilian market for each of the 200 firms I used in the estimation.

The results of these predictions are displayed in figures 9 10. These figures compare the realized data and the median predictions from the two models of the total trade and number of exports from the 200 firms in the sample.[55] Figure 9 reports the decrease in wine export to Brazil that occurs between 1998 to 2003. This decrease is explained by the Brazilian devaluation in 1999, and the growth in Argentinian export led by their devaluation in 2002. However, total exports increase after 2003 as a result of the improvement in economic conditions in Brazil at this period. Regarding the predictions of the models, we can see that the model without consumer margin does not react very much to the changes in exchange rates. This variation in relative prices does reduce sales, but not in the same magnitude as in the data. However, the model with consumer margin can better predict the reduction in trade due to the devaluation of the Brazilian and Argentinian currencies.

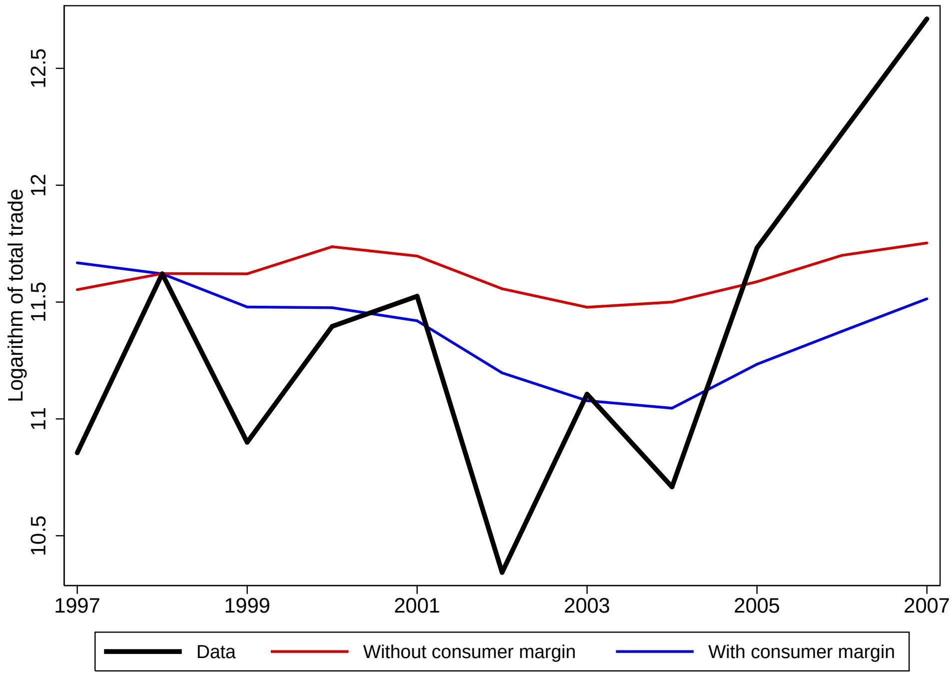

FIGURE 9: Total exports of wine to Brazil from selected firms

Figure 9 highlights the reason for this difference in trade responses across models. In addition to matching the correct number of exporters to Brazil, the full model with consumer accumulation

[54]From the demand equation used in the model, $X_{dt} = \log Y_{dt} - (1-\sigma)\log P_{dt} + (1-\sigma)\log(\tau_{dt}e_{dt})$. Therefore, I use the Brazilian GDP, the BRA/FRA exchanges rates and exchange rates movements of main wine exporters to Brazil to construct variations in the price index. To obtain the level of $X_{BRA,t}$, I set $X_{BRA,t}$ such that the sales of the median prediction equals the realized sales on the market in 1998, the year before the shock. Appendix F provides details.

[55]In appendix F, I provide these figures with the inclusion of the confidence intervals of the predictions.

does predict that firms will leave the Brazilian market.[56] This adjustment of the extensive margin does not appear in the restricted model: because of large entry costs, firms will prefer to lose money temporarily when a negative shock occurs, in order to keep the option value of exporting in the next years. On the contrary, with consumer accumulation and lower entry costs, some firms decide to leave the market because they know it will be affordable to potentially reenter in the future.

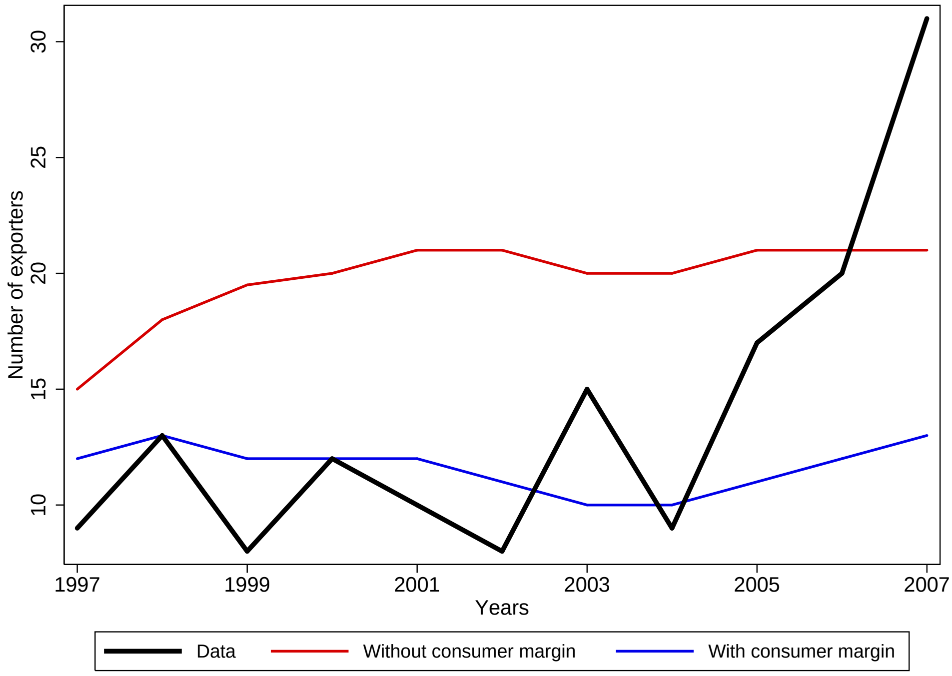

Figure 10: Number of wine exporters to Brazil from selected firms

Overall, it appears that the predictions of the model with consumer margin, unlike the standard model, can quantitatively replicate the decrease in total trade during this period. This result comes from the larger response of firms entry and exit, due to the lower level of the entry costs of exporting in this model.

7 Conclusion

In this paper, I develop and estimate a dynamic empirical model of trade that features state dependence in demand through the accumulation of consumers in foreign markets. Estimating the model using a set of French wine exporters, I show that accounting for this dependence is critical to understand the entry and exit decisions of firms in foreign markets, but also for the estimation of the costs of exporting: on average, estimated entry costs are less than a third of those estimated in the standard model without consumer accumulation. Moreover, I demonstrate using simulations and out-of-sample predictions that this consumer margin, and the associated

[56] Even though it is not clear on figure 9 that only reports the number of exporters included in the estimation sample, the total number of French exporters in Brazil did decrease during this period.

fall in entry costs, matters for aggregate predictions: the model can generate a slow response of aggregate trade to shocks and can correctly replicate the contribution of the extensive margin throughout a trade liberalization episode.

These results shed new light on the nature of the barriers to trade at the firm level. While existing models emphasize the role of large sunk entry costs as the main barrier to trade to explain the persistence in export markets, this paper shows that dependence in demand is responsible for a significant share of this persistence. In fact, the ability to reach a large and stable demand for a product appears to be one of the primary sources of success for firms in foreign markets. Therefore, this study improves our understanding of the determinants of trade dynamics at the microeconomic and aggregate levels, which has important implications for countries aiming to improve the export performance of their industries.

References

Aeberhardt, R., I. Buono, and H. Fadinger (2014): "Learning, incomplete contracts and export dynamics: theory and evidence from French firms," *European Economic Review*, 68, 219–249.

Akhmetova, Z. and C. Mitaritonna (2012): "A Model of Firm Experimentation under Demand Uncertainty with an Application to Multi-Destination Exporters," *University of New South Wales Working Paper*.

Albornoz, F., H. F. C. Pardo, G. Corcos, and E. Ornelas (2012): "Sequential exporting," *Journal of International Economics*, 88, 17–31.

Alessandria, G. and H. Choi (2007): "Do Sunk Costs of Exporting Matter for Net Export Dynamics?" *The Quarterly Journal of Economics*, 122, 289–336.

——— (2014): "Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization," *Journal of International Economics*, 94, 207–223.

Alessandria, G., H. Choi, and K. Ruhl (2014): "Trade Adjustment Dynamics and the Welfare Gains from Trade," Working Paper 20663, National Bureau of Economic Research.

Alessandria, G., S. Pratap, and V. Z. Yue (2013): "Export dynamics in large devaluations," *Manuscript*.

Andrieu, C., A. Doucet, and R. Holenstein (2010): "Particle markov chain monte carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72, 269–342.

Arkolakis, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of political economy*, 118, 1151–1199.

——— (2015): "A Unified Theory of Firm Selection and Growth," *The Quarterly Journal of Economics*.

Aw, B. Y., M. J. Roberts, and D. Yi Xu (2011): "R&D Investment, Exporting, and Productivity Dynamics," *The American Economic Review*, 101, 1312–1344.

Berman, N., V. Rebeyrol, and V. Vicard (2015): "Demand learning and firm dynamics: evidence from exporters," *Manuscript*.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007): "Firms in International Trade," *The Journal of Economic Perspectives*, 105–130.

Bernard, A. B., R. Massari, J.-D. Reyes, and D. Taglioni (2016): "Exporter dynamics and partial-year effects," *Manuscript*.

Berthou, A. and V. Vicard (2015): "Firms' Export Dynamics: Experience Versus Size," *The World Economy*, 38, 1130–1158.

Blevins, J. R. (2015): "Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models," *Journal of Applied Econometrics*.

Broda, C. and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 121, 541–585.

Das, S., M. J. Roberts, and J. R. Tybout (2007): "Market entry costs, producer heterogeneity, and export dynamics," *Econometrica*, 75, 837–873.

Drozd, L. A. and J. B. Nosal (2012): "Understanding international prices: Customers as capital," *The American Economic Review*, 102, 364–395.

Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2010): "State dependence and alternative explanations for consumer inertia," *The RAND Journal of Economics*, 41, 417–445.

Eaton, J., M. Eslava, D. Jinkins, C. Krizan, M. Kugler, and J. Tybout (2014): "A Search and Learning Model of Export Dynamics," *Manuscript*.

Fitzgerald, D., S. Haller, and Y. Yedid-Levi (2015): "How Exporters Grow," *Manuscript*.

Flury, T. and N. Shephard (2011): "Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models," *Econometric Theory*, 27, 933–956.

Foster, L., J. Haltiwanger, and C. Syverson (2016): "The slow growth of new plants: Learning about demand?" *Economica*, 83, 91–129.

Gourio, F. and L. Rudanko (2014): "Customer Capital," *Review of Economic Studies*, 81, 1102–1136.

Heckman, J. J. (1981): "Heterogeneity and State Dependence," *NBER Chapters*, 91–140.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016): "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 131, 1291–1364.

Hotz, J. and R. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60, 497–529.

Hummels, D. L. and P. Klenow (2005): "The Variety and Quality of a Nation's Exports," *American Economic Review*, 95, 704–723.

Imai, S., N. Jain, and A. Ching (2009): "Bayesian estimation of dynamic discrete choice models," *Econometrica*, 77, 1865–1899.

Kehoe, T. J. and K. J. Ruhl (2013): "How important is the new goods margin in international trade?" *Journal of Political Economy*, 121, 358–392.

Khandelwal, A. (2010): "The long and short (of) quality ladders," *The Review of Economic Studies*, 77, 1450–1476.

Li, S. (2014): "A structural model of productivity, uncertain demand, and export dynamics," *Manuscript*.

Lincoln, W. F. and A. H. McCallum (2015): "The Rise of Exporting By US Firms," International Finance Discussion Paper 2015-53, Board of Governors of the Federal Reserve System.

Macchiavello, R. (2010): "Development uncorked: Reputation acquisition in the new market for Chilean wines in the UK," *Manuscript*.

McCallum, A. H. (2015): "The Structure of Export Entry Costs," *Manuscript*.

Melitz, M. J. (2003): "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 71, 1695–1725.

Morales, E., G. Sheu, and A. Zahler (2014): "Extended Gravity," *Manuscript*.

Nguyen, D. X. (2012): "Demand uncertainty: Exporting delays and exporting failures," *Journal of International Economics*, 86, 336–344.

Norets, A. (2009): "Inference in dynamic discrete choice models with serially orrelated unobserved state variables," *Econometrica*, 77, 1665–1682.

Osborne, M. (2011): "Consumer learning, switching costs, and heterogeneity: A structural examination," *Quantitative Marketing and Economics*, 9, 25–70.

Pierce, J. R. and P. K. Schott (2012): "Concording US Harmonized System Codes over Time," *Journal of Official Statistics*, 28, 53–68.

Piveteau, P. and G. Smagghue (2015): "Estimating firm product quality using trade data," *Manuscript*.

Rauch, J. E. and J. Watson (2003): "Starting small in an unfamiliar environment," *International Journal of Industrial Organization*, 21, 1021–1042.

Roberts, M. J., D. Y. Xu, X. Fan, and S. Zhang (2012): "A Structural Model of Demand, Cost, and Export Market Selection for Chinese Footwear Producers," Working Paper 17725, National Bureau of Economic Research.

Rodrigue, J. and Y. Tan (2015): "Price and Quality Dynamics in Export Markets," *Manuscript*.

Ruhl, K. J. (2008): "The international elasticity puzzle," *Manuscript*.

Ruhl, K. J. and J. L. Willis (2008): "New exporter dynamics," *Manuscript*.

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

Shocker, A. D., M. Ben-Akiva, B. Boccara, and P. Nedungadi (1991): "Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions," *Marketing letters*, 2, 181–197.

Timoshenko, O. A. (2015): "Learning versus sunk costs explanations of export persistence," *European Economic Review*, 79, 113–128.

Van Beveren, I., A. B. Bernard, and H. Vandenbussche (2012): "Concording EU trade and production data over time," Working Paper 18604, National Bureau of Economic Research.

Appendices

A Constructions of the samples

The dataset used in the paper is initially disaggregated at the monthly level. From this raw dataset, a number of steps are implemented to improve the reliability and consistency of the data. First, I describe the operations implemented for the first empirical exercise, that uses a wide set of products. Then, I describe the procedures implemented to obtain the final sample used in the structural estimation.

A.1 Data appendix for the reduced-form exercise

I implement two important steps to prepare the data for the regressions displayed in the reduced-form exercise. First, I clean outliers and product categories that do not provide a meaningful and consistent unit of count across years. Second, I correct for the partial-year bias.

Cleaning and harmonization I make three different operations to clean the dataset from potential outliers or measurement errors.

- First, I use the algorithm from Pierce and Schott (2012) and Van Beveren, Bernard, and Vandenbussche (2012) to account for changes in product categories at the eight digit level. This algorithm allows me to obtain categories that are consistent across the sample years (1996-2010).
- Second, I drop product categories that meet one of the following criteria:
  - the counting unit is changing across years.
  - the counting unit is not identical within the category (because of the previous step, the current product category can contain eight digit categories with different units).
  - the counting unit is weight. The reason for this exclusion relies on the use of weight for many categories as the default unit. While this can be a relevant unit for some goods, it is often used for product categories that gather non homogeneous product.
- Finally, because unit values, constructed as export values divided by quantities, are a source of measurement errors, I winsorize them at the eight-digit product category×country×year level. Specifically, I set at the values of the 5th and 95th percentiles the prices that are beyond these two thresholds.

Correction for partial-year bias As described in Berthou and Vicard (2015) and Bernard, Massari, Reyes, and Taglioni (2016), a firm will sell less in average during its first calendar year as exporter. This is because calendar years do not necessarily match the beginning of the exporting activity. In order to correct for this potential bias, I reconstruct the dataset to align calendar exporting years of each exporter. The idea is to define a new year for each spell of export, setting the first month of this year as representative of a regular year, and constructing exporting spells based on this new starting month.

Specifically, the following procedure is applied to each firm-destination-product triplet: for the earliest observation in 1996, if no observation is seen in 1995, a new spell is defined: the month of this first flow is probabilistically drawn based on the number of flows observed during the following 12 months. Then, the year is set to 1996 or 1997 depending on whether the initial month is earlier or later than July. The following observations are adjusted accordingly to preserve the duration between monthly export flows, as long as there is no discontinuity in the exporting activity according to the newly defined calendar years. In case of discontinuity, the next observation becomes a new reference point, and the same procedure is applied for this observation and the following ones.

Once this adjustment is implemented, I aggregate the data at the yearly-level. Specifically, I sum values exported within each newly created calendar year at the firm-product-category level. Moreover, I obtain yearly prices using an export-weighted average of monthly prices. In case of missing prices, I assume a weight of zero for this observation. If this observation is the only observation within a firm-destination-product- year combination, I drop all the observations within the firm-destination-product triplet.

This procedure leaves me with sales and prices measured at the firm-product-destination-year level, with no missing observation in prices, and adjusted for the existence of partial-year of exporting.

A.2 Data appendix for the structural estimation

The procedure to clean the data for the structural estimation is different than the reduced-form exercise. I describe in this subsection the choice of the wine industry and the set of destinations I use for implementing my estimation. Then, I describe the cleaning procedure implemented on the wine producers and provide summary statistics on the final sample of firms used in the estimation.

A.2.1 Wine industry

The decision to implement this estimation on wine exporters relies on two constraints. First of all, I study the entry decision made at the firm level. This level of analysis is explained by the fact that brands and reputation are often defined by the firm that produces the good. Therefore, this requires to study firms that display a small level of heterogeneity in terms of goods. A car producer for instance, that also exports car pieces, or engines for other vehicles, is difficult to analyze as a single-product firm. However, a wine producer mostly export wines, and specifically bottles of wine, whose prices are easy to define, and aggregate at the firm level. For these reasons when defining my sample, I exclusively use wine producers that do not export any other goods outside of wine. A large share of the trade in wine is made by wholesalers who export other types of items, and for which the study at the level of the firm is irrelevant. In addition to this homogeneity constraint, my estimation procedure requires enough firms which export to several destinations. As a major exporting industry from France, the wine industry meets both of these conditions: a large number of exporters, exporting a precisely defined good.

In addition to imposing restrictions on the set of firms included in the final sample, I only use a restricted set of destinations.

A.2.2 Selection of destinations

I select 15 different destinations on which I analyze the behaviors of French exporters. These destinations have been selected among the 20 most popular destinations for wine exports from France, excluding countries with large import/export platforms such as Denmark and Singapore, while reflecting some heterogeneity in terms of location. Moreover, I divide these destinations in three groups, for which I estimate different entry and fixed costs of exporting, as well as different trend in aggregate demand. The list of these destinations can be found in table 4.

TABLE 4: List of destination countries included in the structural sample

	Group 1 *Europe*		**Group 2** *Americas*	**Group 3** *Asia/Oceania*
Great-Britain	Germany	Belgium	(Brazil)	Australia
Netherlands	Italy	Spain	Canada	China
Ireland	Sweden	Switzerland	United States	Japan

Note that I do not include Brazil in the structural sample. The observations related to this destination will be used in the out-of-sample exercise and are excluded so that it does not affect the estimation procedure.

A.2.3 Aggregation

Because the estimation is conducted at the firm-destination-year level, it is necessary to aggregate the sales and quantities exported across products exported by the firm. The choice of the wine industry is crucial here since bottles of wines are quantities that can be easily aggregated. An industry producing differentiated goods would have made this aggregation less straightforward.

The aggregation of prices and sales are the following:

$$p_{fdt} = \sum_{h=1}^{H_{fdt}} w_{fhdt} \frac{s_{fhdt}}{q_{fhdt}} \qquad \text{with} \quad w_{fhdt} \equiv \frac{s_{fhdt}}{\sum_h s_{fhdt}}$$
$$s_{fdt} = \sum_{h=1}^{H_{fdt}} s_{fhdt}$$

where H_{fdt} is the number of 8-digit observations for each firm-destination-year triplet. Moreover, note that there is a certain number of missing quantities in the data. Therefore, I assign a weight w_{fhdt} equal to zero to the observations that have quantities or values exported equal to one or zero. When this observation is the only one at the firm-destination-year level (no other product is sent to this market by this firm this year), I dropped all the observations related to this firm from the sample.

A.2.4 Partial-year bias

Similar to the sample used in the reduced form exercise, I will correct for the partial-year bias, by redefining the entry months of all entering exporters. As a consequence, I shift all the subsequent flows to maintain the same sequence in the exports of the firm. Therefore, exports during the first year will look similar to the subsequent years of exporting.

A.2.5 Cleaning

I clean the data to avoid the potential existence of outliers in prices. In order to do so, I run a regression of the logarithm of prices, on sets of time, destinations and firm-specific dummies. Formally, I estimate

$$\log p_{fdt} = \alpha_f + \beta_d + \gamma_t + \varepsilon_{fdt}$$

and I define $\log \hat{p}_{fdt} = \hat{\alpha}_f + \hat{\beta}_d + \hat{\gamma}_t$. Therefore I can flag prices that deviate from these predicted prices. In particular, I consider outliers prices that deviate from a factor 2 of its predicted value ($p_{fdt} > 2\hat{p}_{fdt}$ or $p_{fdt} < 1/2\hat{p}_{fdt}$). As a cleaning procedure, I dropped all the observations of a firm which has at least one outlier among its observations.

Finally, a last criterion for a firm to be included in the final sample is based on the number of observations. Many firms export one year to one market during the sample period, and this does not provide enough information to analyze their exporting behavior. Therefore, I only keep firms that recorded at least 15 exporting events. Note that with 14 destinations and 14 years of data, the maximum number of observations by a given firm is 196. This selection process could present a problem as it is likely to affect the estimates of entry and fixed costs of exporting, by only looking at successful firms. However, this procedure will tend to select firms that survive several years, rather than short-lived exporters: as a consequence, it tends to go against the theory of consumer accumulation that can accommodate small and short-lived exporters relative to the standard model.

A.2.6 Final sample

Once these cleaning steps were implemented, I randomly sampled 200 firms among the set of firms available. Moreover, in order to have enough exporters that have activity in Brazil, and conduct the out-of-sample predictions exercise, I required that 100 of these 200 firms have some exporting activity in Brazil during the sample period.

TABLE 5: Description of the sample used in the structural estimation

Statistics:	*pc5*	*median*	*pc95*	*mean*	N
# observations per firm	15	36.5	97.5	44.2	200
av. # destinations per firm-year	1.65	3.64	8.29	4.16	2118
av. # years per firm-destination	2.5	5	9.5	5.29	1626

Table 5 provides information regarding the number of observations provided by the sampled firms, as well as the number of destinations they export to in an average year. One can see that the firms selected are relatively large, with a minimum number of export episodes equal to 15 by the sampling procedure. However, the median firm only records 36 export episodes, while the maximum number of episodes in the dataset is 196 (14×14). Moreover, they are relatively diversified in terms of destinations since the median firm exports to 3.64 destinations in an average year.

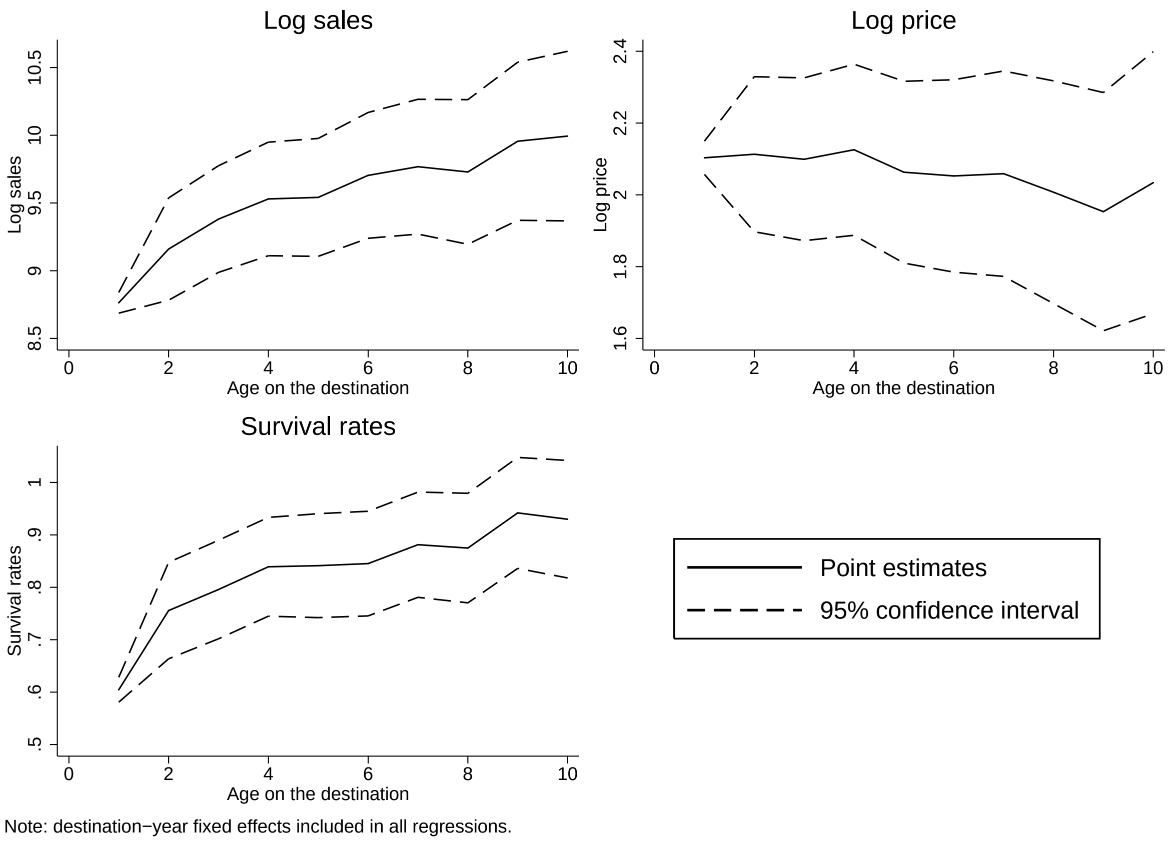

FIGURE 11: Sales, prices and survival rates across ages (Wine producers)

Notes: The figure reports the average log sales, log prices and survival rates of wine producers in a destination at different ages. The estimates are obtained from the regression of these dependent variables on a set of age dummies and destination×year fixed effects. The age in a destination is defined as the number of years a firm has been successively exporting to this country. 95 percent confidence intervals are constructed using standard errors estimates clustered at the firm-destination level.

In order to inspect how this sampling procedure affects the trajectories of the exporters, I replicate the regressions on age dummies I perform in section 2. Figure 11 reports the results of these regressions for sales, prices and survival rates.[57] The patterns of sales and prices are very similar to the ones observed using the comprehensive sample: sales appear to increase in the early years, with the an average growth rate of 30 percent the first year. Meanwhile, the variations in prices are small and insignificant across ages. However, we can see that the survival rates in the structural sample are larger than the ones displayed in the exhaustive data. While the

[57] Table 6 provides the tables related to these regressions.

survival rate was close to 35 percent in the full sample, it is around 60 percent in this restricted sample. This arises because of the requirement made during the selection of exporters: because the estimation procedure requires firms with several observations, this tends to eliminate firms with very large attrition rates that do not records many episodes of exporting activity. Note that this difference in survival rates between exhaustive and restricted samples will play against the story I develop in this paper. Large attrition rates will be consistent with a story that emphasizes strong dependence in demand rather than an important role for sunk costs of entry.

Table 6: Age regressions using the structural sample

	No fixed effects			Year x destination fixed effects		
	(1) Log sales	(2) Log prices	(3) Surv. rates	(4) Log sales	(5) Log prices	(6) Surv. rates
Age 2	0.432***	0.0166	0.144***	0.397***	0.00989	0.151***
	(0.0396)	(0.0189)	(0.0178)	(0.0399)	(0.0191)	(0.0175)
Age 3	0.684***	0.0236	0.180***	0.618***	-0.00412	0.191***
	(0.0493)	(0.0251)	(0.0193)	(0.0517)	(0.0261)	(0.0197)
Age 4	0.827***	0.0556	0.220***	0.767***	0.0224	0.235***
	(0.0598)	(0.0312)	(0.0200)	(0.0632)	(0.0338)	(0.0208)
Age 5	0.841***	0.0142	0.233***	0.778***	-0.0401	0.237***
	(0.0710)	(0.0370)	(0.0224)	(0.0766)	(0.0396)	(0.0236)
Age 6	0.987***	-0.00312	0.248***	0.941***	-0.0506	0.241***
	(0.0809)	(0.0433)	(0.0235)	(0.0876)	(0.0466)	(0.0251)
Age 7	1.045***	-0.0118	0.277***	1.005***	-0.0443	0.277***
	(0.0947)	(0.0485)	(0.0243)	(0.104)	(0.0553)	(0.0254)
Age 8	1.027***	-0.0473	0.253***	0.966***	-0.0958	0.270***
	(0.109)	(0.0549)	(0.0281)	(0.121)	(0.0646)	(0.0296)
Age 9	1.265***	-0.0996	0.313***	1.193***	-0.150*	0.337***
	(0.120)	(0.0590)	(0.0261)	(0.137)	(0.0742)	(0.0295)
Age 10	1.318***	-0.000281	0.301***	1.231***	-0.0693	0.325***
	(0.135)	(0.0711)	(0.0305)	(0.154)	(0.0884)	(0.0323)
Observations	6121	6121	5511	6121	6121	5511
R^2	0.078	0.001	0.060	0.177	0.136	0.120

Notes: Firm x destination clustered standard errors between parentheses. * p<0.05, ** p<0.01, *** p<0.001

B Additional age regressions

In this section, I describe alternative specifications to look at the correlation between sales or prices and age in an export market.

B.1 Additional specifications

Firm-destination-product fixed effects

A natural way to control for heterogeneity across firms, which could drive the correlation across ages, is to include firm-destination-product fixed effects such that the regression becomes

$$X_{fpdt} = \sum_{\tau=1}^{10} \delta_\tau \mathbb{1}(\text{age}_{fpdt} = \tau) + \mu_{pdt} + \mu_{fpd} + \varepsilon_{fdt}.$$

However, including this set of fixed effects will make it impossible to identify a trend in prices across ages. To understand why, first consider a sample of firms on a given market *pdt*. Because of the market-level fixed effect, their average price is normalized to zero. Now consider this same set of firms a year later. If none of these firms exited, it means that their average price is normalized to zero. More generally, the fact that age is a treatment that is homogenous across firms makes the identification of any trend impossible. However, because in the data, some firms will exit the market, it means that this treatment is not entirely symmetrical across firms, such that some identification is possible. But this identification will entirely rely on firms that exit and re-enter, with an age that will be one in the future. As a consequence, the inclusion of this set of fixed effects will not control for selection, but instead will make the entry and exit of firms the only source of identification. Figures 12 and 13 report the results of this specification for sales and prices. As we can see, even sales are not increasing with age with this specification.

Identification across destinations

An alternative way to identify an increase in sales and prices across age is to compare similar products sold to different destinations, and, therefore, having different export experiences. In terms of specifications, it means including a set of firm-product fixed effects such that the variation identifying the changes with age occurs across destinations. However, this specification is also potentially problematic since it compares old destinations, for which the firms has chosen to export first, and young destinations that have been chosen more recently by the firm. Therefore, it is not clear that the age across these flows are the only differences. To verify this claim, I run the following specification and display the results for sales and prices in figures 14 and 15.

$$X_{fpdt} = \sum_{\tau=1}^{10} \delta_\tau \mathbb{1}(\text{age}_{fpdt} = \tau) + \mu_{pdt} + \mu_{fp} + \varepsilon_{fdt}$$

We can see that all figures maintain the increasing in trends of sales and prices, even though price regressions are not as significant as in the main specification. However, one can see that the endogenous sorting of the destinations seem to play a role in shaping this relationship: using

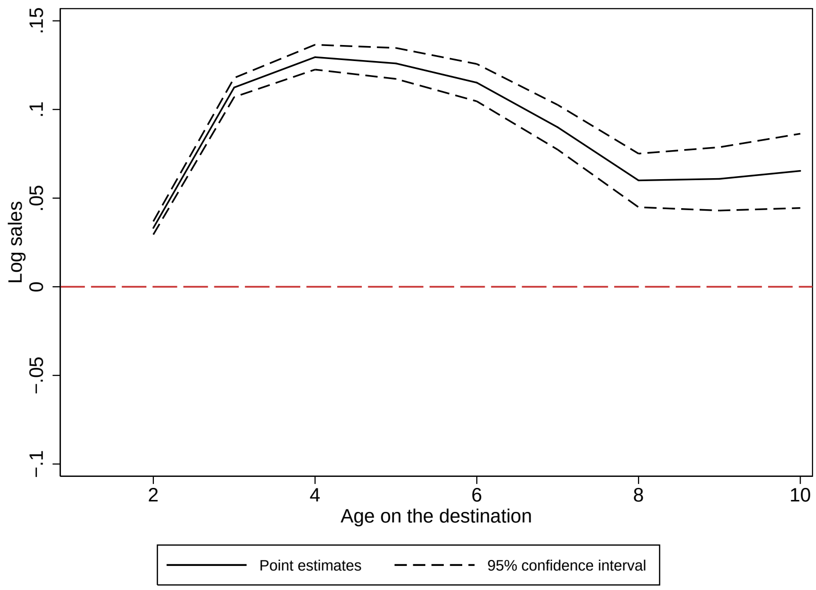

FIGURE 12: Sales across export ages, within variation

Notes: The figure reports the cumulative growth of sales compared to age one, of a firm-product category pair in a destination at different ages. The regression uses logarithm of sales as dependent variable, and includes product category×destination×year and firm×product category×destination fixed effects. 95 percent confidence intervals are constructed using standard errors clustered at the firm-product-destination level.

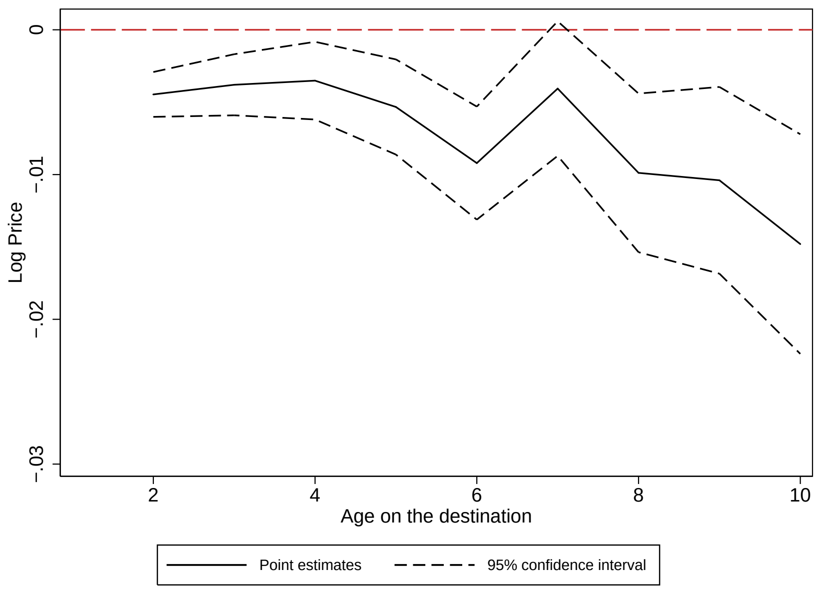

FIGURE 13: Prices across export ages, within variation

Notes: The figure reports the cumulative growth of prices compared to age one, of a firm-product category pair in a destination at different ages. The regression uses logarithm of sales as dependent variable, and includes product category×destination×year and firm×product category×destination fixed effects. 95 percent confidence intervals are constructed using standard errors clustered at the firm-product-destination level.

a constant set of firms tends to increase the growth in sales. Therefore, it is difficult to imagine that this specification accounts for the dynamic selection across age, but instead could pick up the endogenous sorting across destinations.

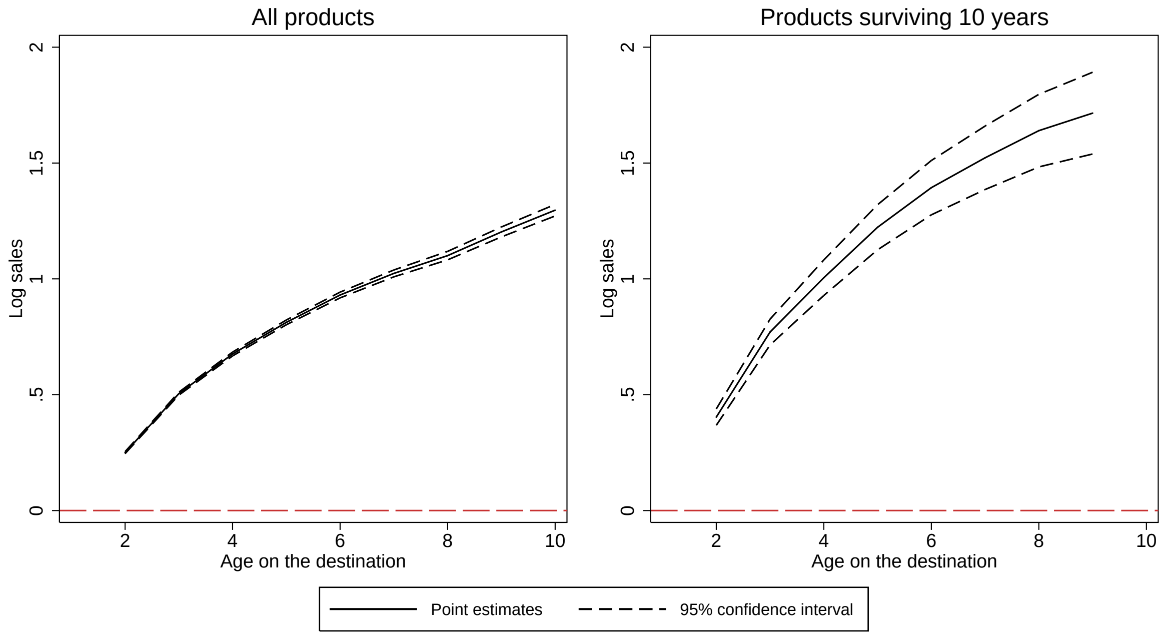

Figure 14: Sales across export ages, across destinations

Notes: The figure reports the cumulative growth of sales compared to age one, of a firm-product category pair in a destination at different ages. The regression uses logarithm of sales as dependent variable, and includes product category×destination×year and firm×product category fixed effects. 95 percent confidence intervals are constructed using standard errors clustered at the firm-product level.

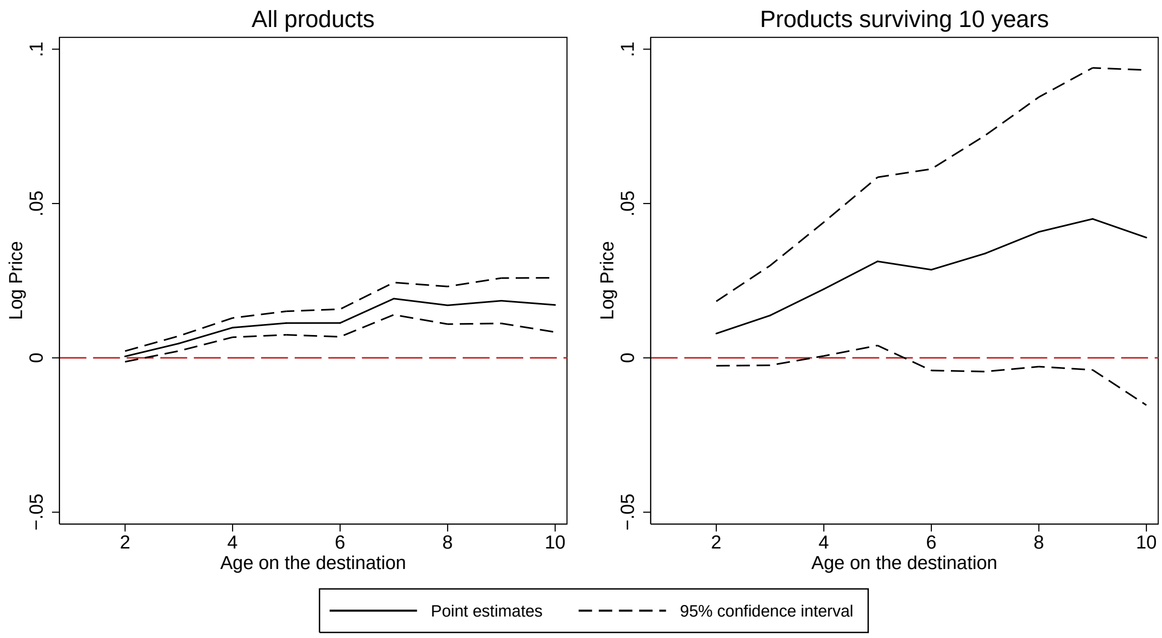

Figure 15: Prices across export ages, across destinations

Notes: The figure reports the cumulative growth of sales compared to age one, of a firm-product category pair in a destination at different ages. The regression uses logarithm of prices as dependent variable, and includes product category×destination×year and firm×product category fixed effects. 95 percent confidence intervals are constructed using standard errors clustered at the firm-product level.

B.2 Tables of results

TABLE 7: Age regressions (main specification)

	All products			Prod. surviving 10 years	
	(1) Surv. rates	(2) Log sales	(3) Log prices	(4) Log sales	(5) Log prices
Age 2	0.217***	0.553***	0.0207***	0.282***	0.0168
	(0.000651)	(0.00206)	(0.00110)	(0.0159)	(0.00999)
Age 3	0.305***	0.969***	0.0335***	0.531***	0.0345**
	(0.000818)	(0.00301)	(0.00149)	(0.0239)	(0.0106)
Age 4	0.356***	1.256***	0.0486***	0.651***	0.0562***
	(0.000969)	(0.00394)	(0.00186)	(0.0321)	(0.0117)
Age 5	0.382***	1.473***	0.0602***	0.753***	0.0776***
	(0.00112)	(0.00493)	(0.00226)	(0.0405)	(0.0126)
Age 6	0.403***	1.661***	0.0630***	0.807***	0.0874***
	(0.00130)	(0.00602)	(0.00269)	(0.0490)	(0.0136)
Age 7	0.411***	1.818***	0.0774***	0.820***	0.105***
	(0.00151)	(0.00735)	(0.00325)	(0.0574)	(0.0146)
Age 8	0.423***	1.954***	0.0821***	0.820***	0.125***
	(0.00175)	(0.00892)	(0.00394)	(0.0660)	(0.0156)
Age 9	0.438***	2.089***	0.0826***	0.780***	0.142***
	(0.00202)	(0.0107)	(0.00473)	(0.0745)	(0.0165)
Age 10	0.453***	2.196***	0.0878***	0.636***	0.150***
	(0.00234)	(0.0129)	(0.00566)	(0.0831)	(0.0174)
Observations	5716155	6293940	6293940	417427	417427
R^2	0.326	0.438	0.874	0.545	0.919

Notes: Firm x product x destination clustered standard errors between parentheses. Year x product x destinations fixed effects are included in all regressions. * p<0.05, ** p<0.01, *** p<0.001

Table 8: Age regressions with alternative specifications

	Firm x product f.e.				Firm x product x dest. f.e.	
	All products		Prod. surviving 10 years			
	(1) Log sales	(2) Log prices	(3) Log sales	(4) Log prices	(5) Log sales	(6) Log prices
Age 2	0.250***	0.000473	0.403***	0.00789	0.0332***	-0.00446***
	(0.00215)	(0.000886)	(0.0181)	(0.00532)	(0.00190)	(0.000790)
Age 3	0.504***	0.00467***	0.770***	0.0137	0.112***	-0.00380***
	(0.00320)	(0.00124)	(0.0284)	(0.00824)	(0.00276)	(0.00108)
Age 4	0.676***	0.00980***	1.005***	0.0223*	0.130***	-0.00351*
	(0.00420)	(0.00159)	(0.0391)	(0.0111)	(0.00358)	(0.00137)
Age 5	0.812***	0.0113***	1.223***	0.0313*	0.126***	-0.00533**
	(0.00525)	(0.00195)	(0.0496)	(0.0139)	(0.00445)	(0.00168)
Age 6	0.930***	0.0113***	1.394***	0.0286	0.115***	-0.00921***
	(0.00632)	(0.00228)	(0.0598)	(0.0167)	(0.00537)	(0.00199)
Age 7	1.024***	0.0192***	1.523***	0.0338	0.0900***	-0.00406
	(0.00756)	(0.00266)	(0.0699)	(0.0195)	(0.00644)	(0.00237)
Age 8	1.100***	0.0170***	1.640***	0.0408	0.0600***	-0.00988***
	(0.00931)	(0.00311)	(0.0799)	(0.0223)	(0.00771)	(0.00280)
Age 9	1.203***	0.0185***	1.715***	0.0450	0.0609***	-0.0104**
	(0.0109)	(0.00375)	(0.0900)	(0.0250)	(0.00911)	(0.00329)
Age 10	1.296***	0.0171***	1.685***	0.0390	0.0654***	-0.0148***
	(0.0127)	(0.00449)	(0.0995)	(0.0277)	(0.0107)	(0.00387)
Observations	5192771	5192771	324977	324977	4022658	4022658
R^2	0.667	0.945	0.739	0.961	0.810	0.967

Notes: Firm x product x destination clustered standard errors between parentheses. Year x product x destinations and firm x products fixed effects are included in all regressions. * p<0.05, ** p<0.01, *** p<0.001

C Derivations

The firm chooses the optimal mark-up μ to maximize the value of exporting:

$$\begin{aligned}
\mu &= \operatorname{argmax} V_I(\xi, n, \mu) \\
&= \operatorname{argmax} E_\varepsilon \Big\{ \pi(\xi, n, \mu, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1) \Big\} \\
&= \operatorname{argmax} \int_\varepsilon \pi(\xi, n, \mu, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1) dF(\varepsilon)
\end{aligned}$$

Therefore, the first order condition of the problem is

$$\begin{aligned}
&\int_\varepsilon \frac{\partial \pi(\xi, n, \mu, \varepsilon)}{\partial \mu} + \beta \frac{\partial EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1)}{\partial \mu} dF(\varepsilon) = 0 \\
\Leftrightarrow &\int_\varepsilon \frac{\partial \pi(\xi, n, \mu, \varepsilon)}{\partial \mu} + \beta \frac{\partial EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1)}{\partial \mu} dF(\varepsilon) = 0
\end{aligned}$$

First, profit function is

$$\begin{aligned}
\pi(\xi, n, \mu, \varepsilon) &= n \exp(\lambda + X + \varepsilon^D) \mu^{-\sigma} (\mu - 1) c(\xi, n, \varepsilon^S)^{1-\sigma} \\
\Rightarrow \frac{\partial \pi(\xi, n, \mu, \varepsilon)}{\partial \mu} &= \big[(1-\sigma)\mu^{-\sigma} + \sigma \mu^{-\sigma-1}\big] \, n \exp(\lambda + X + \varepsilon^D) c(\xi, n, \varepsilon^S)^{1-\sigma}
\end{aligned}$$

Second, the continuation value can be rewritten $EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1) = EV'(\xi, n'(s, n), 1)$ where s are the sales of the firm. Therefore,

$$\begin{aligned}
\frac{\partial EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1)}{\partial \mu} &= \frac{\partial EV'(\xi, n'(s, n), 1)}{\partial \mu} \\
&= \frac{\partial s}{\partial \mu} \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \\
&= (1-\sigma)\mu^{-\sigma} n \exp(\lambda + X + \varepsilon^D) c(\xi, n, \varepsilon^S)^{1-\sigma} \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'}
\end{aligned}$$

Therefore, the first order condition can be rewritten

$$
\begin{aligned}
&\int_{\varepsilon} n \exp(\lambda + X + \varepsilon^D) c(\xi, n, \varepsilon^S)^{1-\sigma} \Big[(1-\sigma)\mu^{-\sigma} + \sigma \mu^{-\sigma-1} \\
&\qquad\qquad + (1-\sigma)\mu^{-\sigma} \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big] dF(\varepsilon) = 0 \\
\Leftrightarrow \quad & \int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) \Big[\sigma + (1-\sigma)\mu \Big(1 + \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big) \Big] dF(\varepsilon) = 0 \\
\Leftrightarrow \quad & \mu(\sigma - 1) \int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) \Big(1 + \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big) dF(\varepsilon) \\
&\qquad\qquad = \sigma \int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) dF(\varepsilon) \\
\Leftrightarrow \quad & \mu = \frac{\sigma}{\sigma - 1} \frac{\int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) dF(\varepsilon)}{\int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) \Big(1 + \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big) dF(\varepsilon)} \\
\Leftrightarrow \quad & \mu = \frac{\sigma}{\sigma - 1} \frac{1}{\int_{\varepsilon} \frac{\exp(\varepsilon^D + (1-\sigma)\varepsilon^S)}{\int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) dF(\varepsilon)} \Big(1 + \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big) dF(\varepsilon)} \\
\Leftrightarrow \quad & \mu = \frac{\sigma}{\sigma - 1} \frac{1}{\int_{\varepsilon} \omega(\varepsilon) \Big(1 + \beta \frac{\partial n'}{\partial s} \frac{\partial EV'(\xi, n', 1)}{\partial n'} \Big) dF(\varepsilon)}
\end{aligned}
$$

with $\omega(\varepsilon) = \frac{\exp(\varepsilon^D + (1-\sigma)\varepsilon^S)}{\int_{\varepsilon} \exp(\varepsilon^D + (1-\sigma)\varepsilon^S) dF(\varepsilon)}$.

D Details of the algorithm

I describe in this section of the appendix the MCMC algorithm I implement. I start by describing how the Markov chain is initialized, before describing a given iteration of the chain, involving the update of the unobservables and parameters.

D.1 Initial values

I start by describing how the unobservables are obtained, before describing the initial parameters. I start by setting an initial value of 2.2 for σ,[58] that allows me to obtain $\log s_{fdt} + \sigma p_{fdt} = \log n_{fdt} + X_{dt} + \lambda_{ft}$. I can then decompose this term using firm-year and destination-year fixed effect. In order to obtain $\phi_{dt}^{(0)}$, I run the regression $\log p_{fdt} - \frac{\sigma}{\sigma - 1}$ on $\lambda_{ft}^{(0)}$. This allows me to obtain $\alpha^{(0)}$, and the residual is regressed on firm-year fixed effects to obtain $\phi_{ft}^{(0)}$. Having in hand initial values for the unobservables, I can use linear regressions to obtain the AR(1) coefficients for the unobservables, and use nonlinear least square to estimate $\underline{n}^{(0)}$, $n_0^{(0)}$, $\eta_1^{(0)}$ and $\eta_2^{(0)}$ after arbitrarily setting $\psi^{(0)} = 0.5$. Finally, I set values for the fixed costs parameters, and the variance parameter of the fixed cost shocks. I arbitrary set $f^{(0)} = fe^{(0)} = s_v^{(0)} = 1\,000$ for the three different groups of countries.

After setting these initial values, I implement 5000 iterations that does not account for the dynamic problem of the firm. Therefore, I sample unobservables and parameters assuming a constant mark-up and only taking advantage of the realized sales and prices. This step allows me to obtain initial conditions for the parameters and unobservables that are closer to their true values, although biased because they do not account for the dynamic problem.

Given this initial set of parameters and unobservables, I can start the iterative procedure described below.

D.2 Creation of the grid

In order to solve for the value function as a function of Θ, I need to create a grid describing the state space of the problem. Note that the state space is made of (λ, ϕ, n, X). Consequently, I need a grid that is relatively more precise for values of the unobservables that are more prevalent. Consequently, I create the four-dimensional grid as following

- $\lambda_g \sim N(0, 5\,\mathrm{std}(\lambda_{ft}))$
- $\phi_g \sim N(\mathrm{mean}(\phi_{ft}), 5\,\mathrm{std}(\phi_{ft}))$
- $X_g \sim N(\mathrm{mean}(X_{ft}), 5\,\mathrm{std}(X_{ft}))$
- $n_g \sim U[\underline{n}\ ;\ 1]$

Note that this grid will be updated when the standard deviations or averages of the current unobservables are 20 percent larger or smaller than the ones used for the current grid, such that the grid will follow the potential change in the distribution of the unobservables. I will set the size of the grid to be 20 on each dimension, such that the value function will be iterated at 20^4 different grid points.

[58] I start from $\sigma = 2.2$ because it is the elasticity obtained by Broda and Weinstein (2006) for the wine industry.

Moreover, in order to solve the optimal mark-up of the firm, I also need to specify a set of grid points for the optimal mark-up term. I create a set of grid points mk_g of size $g_m = 20$, such that $\mu = \frac{\sigma}{\sigma-1}\frac{1}{mk_g}$, with $mk_g \equiv \{1 \cup \{\exp(\log(0.01) + \log(1.5)\frac{i}{g_m-2})\}_{i=0..g_m-2}\}$.

D.3 Iteration

Three different objects will be updated at each iteration of the Markov Chain:

- the value function $V(\Theta^{(s)})$,
- the set of unobservables $\xi^{(s)}_{fdt} = (\lambda^{(s)}_{ft}, \phi^{(s)}_{ft}, X^{(s)}_{dt})$,
- the parameter vector $\Theta^{(s)}$.

I perform 60,000 iterations of the Markov chain, discarding the first 30,000 iterations. In the next paragraphs, I describe each of these following steps. I start by describing the step that aims to compute the value functions since they define objects that are used in the other steps. I then turn to the sampling of unobservables, and the sampling of parameters.

Update of the value function The value functions are obtained from the Bellman equation, iterated from the previous iteration of the value functions. From section 3, we have

$$V_I(\xi, n) = \max_{\mu}\left\{E_\varepsilon\left\{\pi(\xi, n, \mu, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, \mu), 1)\right\}\right\}$$

Therefore, the value function is updated the following way:

$$\begin{aligned} V(\xi_g, n_g, \mathcal{I}, \Theta^{(s+1)}) = s_v \log\Bigg[&\exp\left(\frac{1}{s_v}EV_O(\xi_g)\right) \\ &+ \exp\left(\frac{1}{s_v}\max_{mk \in mk_g}\left\{E_\varepsilon\pi(\xi_g, n_g, mk_g, \Theta^{(s+1)}) - f + EV_I(\xi_g, mk_g)\right\}\right)\Bigg] \end{aligned} \tag{13}$$

with

$$\begin{aligned} EV_I(\xi_g, mk_g) &= \frac{\sum_{\xi\,\in\,\xi_g}\sum_{n\in n_g}\ V(\xi, n, I, \Theta^{(s)})\ P_n(n\,|\,\xi_g, mk_g)\ P_\xi(\xi\,|\,\xi_g)}{\sum_{\xi\,\in\,\xi_g}\ \sum_{n\in n_g} P_n(n\,|\,\xi_g, mk_g)\ P_\xi(\xi\,|\,\xi_g)}, \\ EV_O(\xi_g) &= \frac{\sum_{\xi\,\in\,\xi_g}\ V(\xi, n_0, 0, \Theta^{(s)})\ P_\xi(\xi\,|\,\xi_g)}{\sum_{\xi\,\in\,\xi_g}\ P_\xi(\xi\,|\,\xi_g)} \end{aligned}$$

$P_\xi(.|.)$ being the transition probability of the unobservables at the current parameters, and $P_n(n\,|\,\xi_g, mk_g)$ the probability of obtaining a share n in the next period given the current unobservables ξ_g and the mark-up decision mk_g.[59] In practice, I iterate several times the Bellman equation, in order to reduce the error coming from the use of the previous value functions. In this case, I iterate not using the (s)-th value function anymore, but the current value function.

In addition to updating the value function, I define, during this iteration, two objects that will be used in the sampling of parameters and unobservables. First, I save the optimal mark-up

[59] This probability is obtained from the shock ε that makes the sales of the firms, and therefore the future share of consumers, non-deterministic.

chosen by the firm. This object, evaluated on the grid, is defined as

$$mk_g^* \equiv \text{argmax}\Big\{E_\varepsilon\Big\{\pi(\xi, n, mk, \varepsilon) + \beta EV'(\xi, n'(\xi, n, \varepsilon, mk), 1)\Big\}\Big\}.$$

Second, I create the difference in expected value functions, $DEV()$, that is defined as

$$DEV(\xi_g, n_g) = EV_I(\xi_g, mk_g^*) - EV_O(\xi_g).$$

This object will be convenient when computing the difference in value functions for each firm.

Sampling of unobservables I sample unobservables using the particle Gibbs sampler described in Andrieu, Doucet, and Holenstein (2010). The idea of this sampler is to use a set of particles to approximate the likelihood function, and draw a specific particle within this set proportionally to the integrated likelihood function. An important point of this sampler is that the current unobservables need to survive all the resampling steps of the algorithm. Moreover, I use a backward sampler when choosing the specific set of particles, in order to further improve the mixing of the sampler.

To further describe the sampling of parameters, I take the example the sampling of the unobservables λ_{ft} and ϕ_{ft}, conditional to the current unobservables at iteration s, $X_{dt}^{(s)}$. These unobservables are proportional to their prior distribution, $F_\lambda()$ and $F_\phi()$, and the conditional likelihood $L(D_{fdt}|D_{fdt-1}, \lambda_{ft}, \phi_{ft}, X_{dt}^{(s)})$. The steps are the following:

- Starting from period 0, and for each firm, I generate r=1..200 particles $(\lambda_{f0}^r, \phi_{f0}^r)$ from their prior distribution.
- I compute the likelihood of each of these particles for each firm:

$$L_{f0}^r \equiv \prod_d L(D_{fd0}|D_{fd-1}, \lambda_{f0}^r, \phi_{f0}^r, X_{d0}^{(s)})$$

  using extrapolations from the functions $DEV(\xi_g, n_g)$ and $mk^*(\xi_g, n_g)$ to obtain the difference in value functions and the mark-up necessary to compute the likelihood.

- for each period t=1..T:
  - I resample 200 $(\lambda_{ft-1}^r, \phi_{ft-1}^r)$ proportionally to L_{ft-1}^r and replace $(\lambda_{ft-1}^{200}, \phi_{ft-1}^{200})$ by $(\lambda_{ft-1}^{(s)}, \phi_{ft-1}^{(s)})$.
  - generate 200 new particles, for each firm, from the prior distribution based on the resampled $(\lambda_{ft-1}^r, \phi_{ft-1}^r)$.
  - I compute the particle-specific likelihood $L_{ft}^r \equiv \prod_d L(D_{fdt}|D_{fdt-1}, \lambda_{ft}^r, \phi_{ft}^r, X_{dt}^{(s)})$
- I retain one specific particle $(\lambda_{ft}^{(s+1)}, \phi_{ft}^{(s+1)})$ by using backward sampling from period T to period 0.
  - Sample one particle for each firm $(\lambda_{fT}^{(s+1)}, \phi_{fT}^{(s+1)})$ proportionally to L_{fT}^r
  - for each period t=T-1..0:
    * sample $(\lambda_{ft}^{(s+1)}, \phi_{ft}^{(s+1)})$ proportionally to $\frac{L_{ft}^r\, F(\phi_{ft+1}^{(s+1)}|\phi_{ft}^r)\, F(\lambda_{ft+1}^{(s+1)}|\lambda_{ft}^r)}{\sum_r L_{ft}^r\, F(\phi_{ft+1}^{(s+1)}|\phi_{ft}^r)\, F(\lambda_{ft+1}^{(s+1)}|\lambda_{ft}^r}$

Sampling of parameters The sampling of parameters is made more complicated by the fact that functions $DEV()$ and $mk()$ need to be reevaluated for a new Θ, rather than for new unobservables. Consequently, for each block of parameters, a Metropolis-Hastings sampler is used. Moreover, it is necessary to iterate the value functions for this new parameter Θ, similarly to the step updating the value functions.

Formally, the sampling of a given block of parameter Θ takes the following steps:

- A new parameter Θ^* is drawn using a proposal function.
- The value function $V(\xi_g, n_g, I, \Theta^*)$ is obtained from equation (13) and the functions $DEV(\xi_g, n_g)$ and $mk_g(\xi_g, n_g)$ are obtained.
- I obtain by interpolation DV_{fdt} and μ_{fdt}, allowing me to compute the likelihood function for the parameter Θ^*.
- $\Theta^{(s+1)}$ is set to be Θ^* with probability $\max\left\{1, \frac{\prod_t \prod_d \prod_f L_{fdt}(D, \xi_{fdt}^{(s+1)}; \Theta^*)}{\prod_t \prod_d \prod_f L_{fdt}(D, \xi_{fdt}^{(s+1)}; \Theta^{(s)})}\right\}$.

The entire set of parameters is divided in blocks that are updated separately:

- Parameters from demand/supply equations, and the law of motion of n(.) are jointly sampled using a random walk proposal function that targets an acceptance rate of 0.2.
- Dynamic parameters (fixed costs and their variance) are sampled using a random walk proposal function that targets an acceptance rate of 0.2.
- The last block of parameters, from the law of motions of the exogenous unobservables, are sampled using only the law of motion, neglecting the impact of these parameters on the dynamic problem of the firm:[60]
  - Parameters of the law of motion of λ are sampled from the set of λ_{ft} using a random walk proposal function that targets an acceptance rate of 0.4.
  - Parameters of the law of motion of ϕ are sampled from the set of ϕ_{ft} using a random walk proposal function that targets an acceptance rate of 0.3.
  - Parameters of the law of motion of X are sampled from the set of X_{ft} using Gibbs sampling.
  - Parameters of the initial distribution of X are sampled from the set of X_{f0} using Gibbs sampling.

Overall, this procedure is doable thanks to parallelization using GPU computing. On average, an iteration of the Markov chain takes 5 seconds, which implies a total computing time of three to four days for 60 000 iterations.

[60]this implies a wider posterior distribution of these parameters relative to the true posterior distribution.

E Additional results

E.1 Full results of the restricted model

Table 9: Estimated parameters - restricted model

Parameter		**Estimate**	**90% Confidence Interval**	
			Lower bound	*Upper bound*
Continuation fixed costs	Europe	12 834	11 976	13 787
(in euros)	Americas	15 860	14 617	17 213
	Asia/Oceania	17 724	15 887	19 764
Entry fixed costs	Europe	114 031	101 885	126 942
(in euros)	Americas	99 823	89 146	111 330
	Asia/Oceania	113 110	100 024	126 799
Variance of continuation costs	σ_ν^c	34 134	30 422	38 327
Variance of entry costs	σ_ν^e	23 160	20 232	26 280
Law of motion of appeal	ρ_λ	0.98	0.97	0.98
	σ_λ	0.31	0.28	0.34
Law of motion of productivity	ρ_ψ	0.95	0.94	0.96
	σ_ψ	0.11	0.10	0.11
	μ_ψ	-0.07	-0.09	-0.06
Law of motion of agg. demand	ρ_X	0.96	0.94	0.98
	σ_X	0.09	0.07	0.11
	μ_{X1}	0.48	0.22	0.76
	μ_{X2}	0.55	0.25	0.86
	μ_{X3}	0.58	0.30	0.87
	μ_{X_0}	12.06	11.60	12.51
	σ_{X_0}	0.92	0.66	1.29
Elasticity cost of appeal	α	0.35	0.33	0.37
Cost dummies	γ_2	0.28	0.26	0.30
	γ_3	0.28	0.25	0.30
Variance matrix	Σ_{11}	2.06	2.00	2.12
	Σ_{12}	0.26	0.25	0.27
	Σ_{22}	0.15	0.14	0.15

E.2 Additional figures

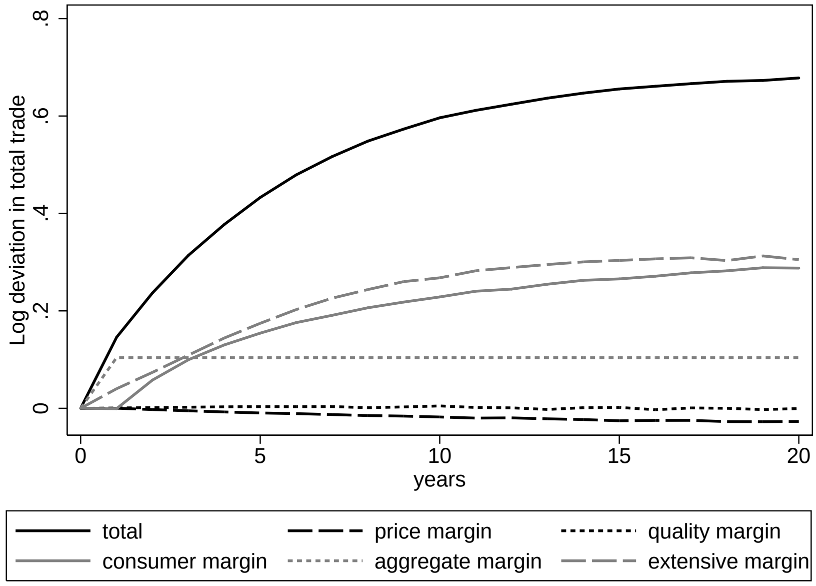

FIGURE 16: Effect of permanent 10 points tariffs decrease (All margins).

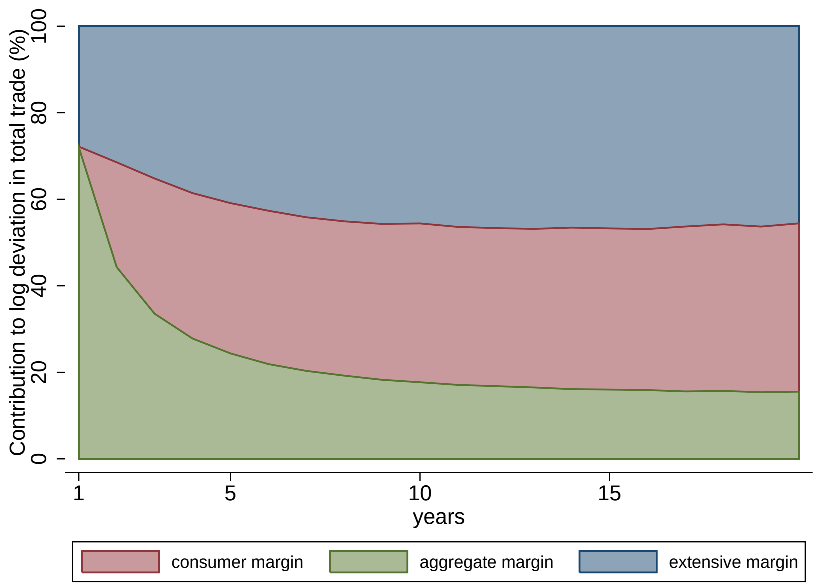

FIGURE 17: Relative share of growth margins during trade expansion.

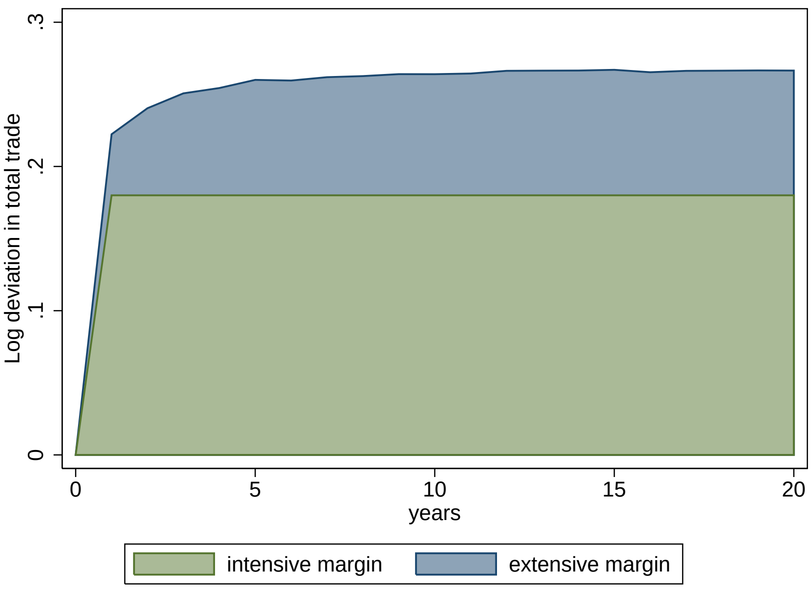

FIGURE 18: Effect of permanent 10 points tariffs decrease (Restricted model).

F Details on out-of-sample predictions

In order to perform out-of-sample predictions, I construct variations in the aggregate demand from Brazil, this variable being defined from the model as $X_{dt} = \log Y_{dt} - (1-\sigma)\log P_{dt} + (1-\sigma)\log(\tau_{dt}e_{dt})$. Importantly, I only need to construct a proxy for variations in this variable. Therefore, in addition to using the Brazilian GDP, and the exchange rate between France and Brazil, I also need to construct a proxy for variations of the price index on the Brazilian market. In order to do so, I use variations in exchange rates from the five main countries exporting to Brazil. Table 10 describes these countries and their respective market shares.

TABLE 10: Top market shares

Country	**Average market share**
France	22.1 %
Italy	20.4 %
Chile	19.6 %
Portugal	15.6 %
Argentina	13.5 %

Notes: Calculations made from BACI. Average market share is the average market share among the Brazilian imports, over the period 1997-2007, for the 4-digit category 2204 '*Wine of fresh grapes*'.

Note that the next largest wine exporter to Brazil has a market share of less than 2 percent and is therefore not included in the computation of the price index. Therefore, I construct a CES price index, using these five exporters, their respective exchange rates and the estimated price elasticity. The obtained variations in aggregated demand for French wine is described in figure 19, which highlights the impact of the Brazilian and Argentinian devaluations.

I then perform 500 simulation of trajectories using the median product appeals and productivities obtained from the estimation, and the constructed variation in aggregated demand. These trajectories differ because I need to simulate demand and supply shocks (ε) and fixed costs shocks (ν) in order to obtain predictions for each firm. Predictions reported in the text are based on the median trajectories, and I report in figures 20 and 21 the 90% confidence interval of these predictions.

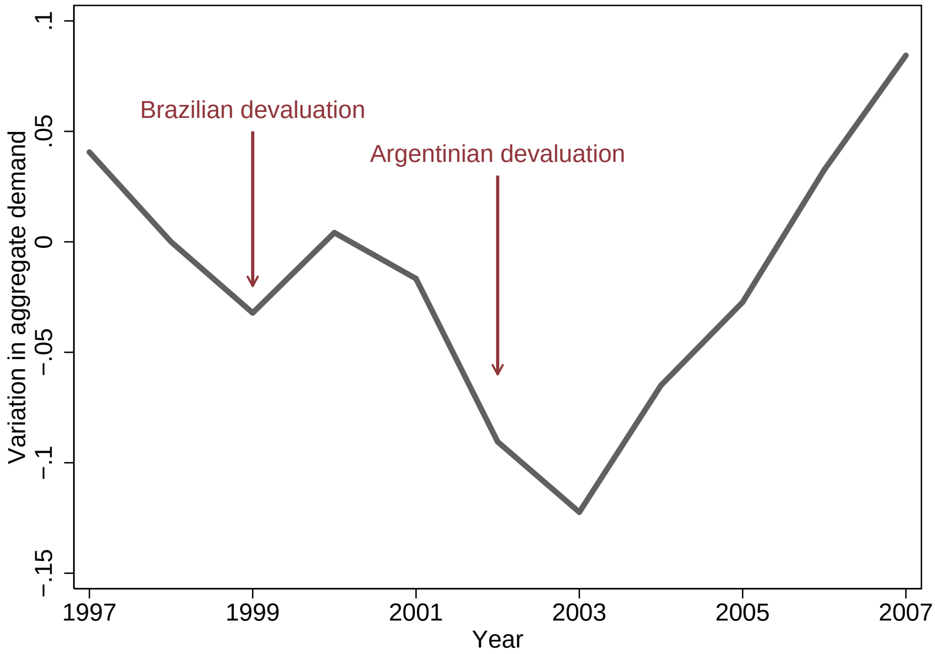

FIGURE 19: Computed variations in aggregate demand for French wine from Brazil.

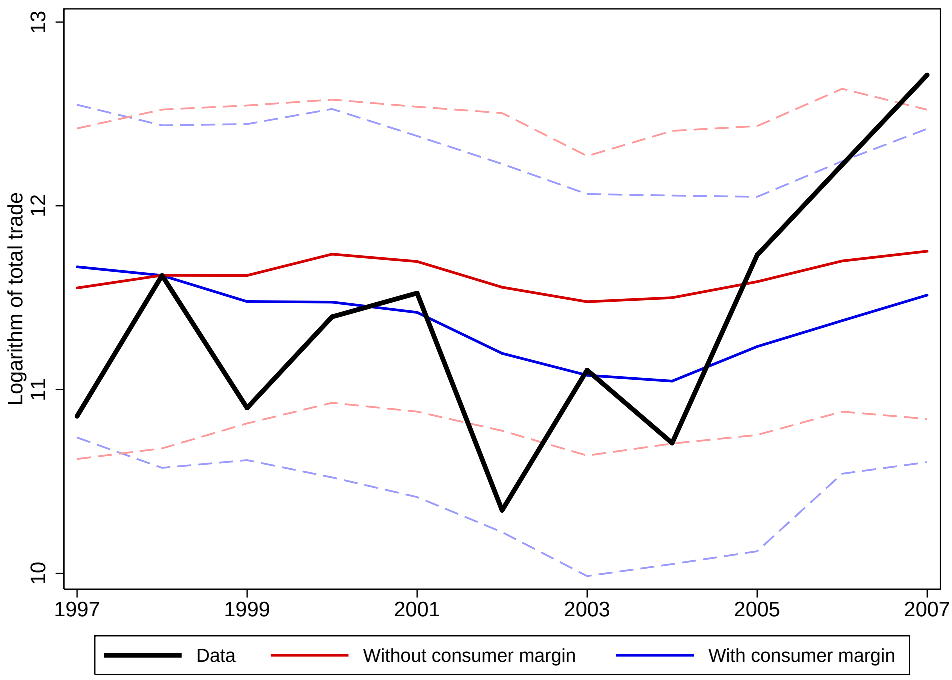

FIGURE 20: Total exports of wine to Brazil from selected firms

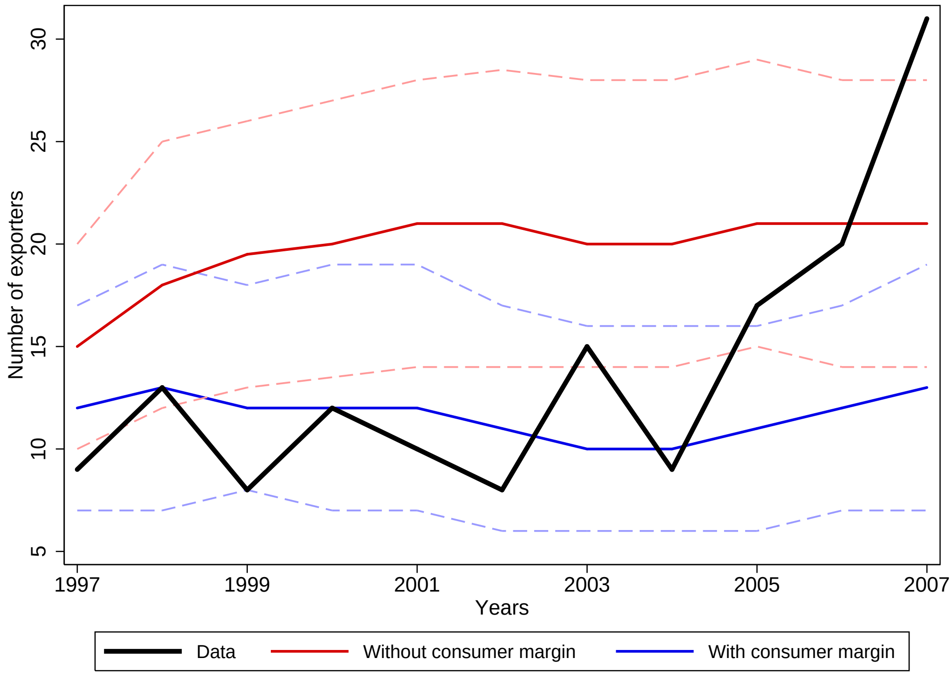

Figure 21: Number of wine exporters to Brazil from selected firms