Export Market Penetration Dynamics*

Joseph B. Steinberg†

July 27, 2023

Abstract

I develop a dynamic theory of exporting that synthesizes two approaches: static models, in which exporting costs depend on the number of foreign customers a firm serves in the present; and dynamic models, in which these costs depend on whether a firm exported in the past. The theory simultaneously accounts for two sets of established empirical findings: (i) larger markets attract exports from more firms and these exports are more concentrated; and (ii) new exporters are generally smaller and more likely to exit than incumbents. It also accounts for a new set of findings from Brazilian microdata showing that differences between new exporters and incumbents are more pronounced in larger markets. When calibrated to match all of these findings, the theory predicts that trade reforms cause greater but also slower trade growth in smaller markets.

JEL Classifications: F10, F12, F14, F15.

Keywords: Market penetration; export participation; new-exporter dynamics; trade liberalization.

1 Introduction

Trade flows are driven by individual firms' decisions: whether to start or stop exporting, whether to expand to new foreign markets, and how much to expand operations in existing markets. The literature on static trade models emphasizes how economic geography determines entry into exporting, and how this margin influences the distribution of exporting firms and the long-run consequences of trade reforms. The dynamic trade literature emphasizes how sunk investments in export capacity shape exporters' life cycles at the micro level, and how this margin drives trade adjustments over time at the macro level. This paper develops a new theory of exporter dynamics that synthesizes these two literatures' modeling approaches. The theory provides a unified account of the empirical regularities documented in these literatures, and also a new set of findings from Brazilian microdata about how exporters' life cycles differ across foreign markets.

The theory extends the endogenous market penetration framework of Arkolakis (2010) to a dynamic environment in which firms start exporting, gradually accumulate customers, and stop exporting in response to persistent idiosyncratic shocks. The theoretical environment consists

*I acknowledge the Social Sciences and Humanities Research Council (grant number 50319) for financial support. I thank the editor, Costas Arkolakis, and three anonymous referees for their insightful comments that have markedly improved the paper. I also thank audiences at the Bank of Canada, the 2019 Dynamic Trade Workshop, the University of Texas at Austin, Yale University, the Virtual International Trade and Macro Seminar, and the Bank of Canada-Orleans Webinar on International Trade. Finally, I thank Paul Piveteau for an excellent discussion and Sam Kortum for helpful feedback. This work is not authorized by Brazil's Foreign Trade Secretariat (SECEX) but it is covered under the doctrine of Fair Use: it is non-profit academic research; it is transformative and none of the underlying data are reproduced herein; and it has no effect on the market value of SECEX products.

†University of Toronto. Email: joseph.steinberg@utoronto.ca.

of an exporting country populated by a continuum of firms and a discrete number of foreign markets that differ in population, income per capita, and trade barriers. Firms in the exporting country differ exogenously in two ways: productivity, which affects a firm's ability to produce for all markets equally; and demand in each market, which reflects the willingness of customers in that market to buy the firm's products. Firms also differ in the number of customers they have in each market, which is endogenous. To build their foreign customer bases, firms must advertise to attract new customers and retain old ones; greater advertising expenditures are required to reach more customers. In equilibrium, firms equate the marginal cost of reaching an additional customer today to the expected present value of the benefit from selling to that customer in the future, or exit if the cost of reaching even a single customer is less than the benefit.

The model has four key properties that allow it to account for the economic mechanisms and empirical facts highlighted by both the static and dynamic trade literatures:

1. The marginal cost of reaching the first customer is strictly positive, regardless of the size of a firm's current customer base.

2. The marginal cost of reaching additional customers is increasing.

3. The marginal cost of reaching additional customers is decreasing in the size of a firm's current customer base.

4. Reaching customers is more expensive in smaller, poorer markets when measured relative to these markets' purchasing power, both on average and at the margin.

The first property, which implies that some firms may choose not to serve any customers at all, generates the extensive margin of trade. Crucially, it generates endogenous exit as well as endogenous entry: incumbent exporters with low enough demand shocks will choose to stop exporting because the marginal cost of retaining the first old customer now exceeds the benefit. The second property, which implies that exporters with lower productivity and/or demand have fewer customers, generates the cross-sectional distribution of export sales. This property accounts for the existence of many small exporters and the concentration of sales among the largest exporters. The third property, which implies that exporters gradually accumulate customers over time, generates exporter life cycles. Combined with the first property, it implies that entrants are more likely to exit than incumbents. Combined with the second property, it implies that firms that start exporting under better conditions, which are more likely to achieve longer export spells, sell more initially and grow faster. The fourth property accounts for differences in these margins across markets. Combined with the first three properties, it implies that larger, richer markets have higher export participation rates, lower exit rates, more small exporters, and more concentrated sales; that new exporters sell less and are more likely to exit than incumbents in larger markets; and that larger markets have greater differences between sales trajectories during short versus long export spells.

The theory nests several widely used existing models as special cases. When retaining old customers is impossible, the decision about how many customers to serve becomes static. In this case, the model is equivalent to Arkolakis (2010), or more specifically, the version analyzed in Arkolakis (2016) in which firms experience growth driven purely by exogenous productivity shocks. When the marginal cost of attracting and retaining customers is constant rather than increasing, all firms serve the same number of customers conditional on choosing to export. In this case, the model is equivalent to the framework of Das et al. (2007) and Alessandria and Choi (2007) in which firms pay a large sunk cost to start exporting and a smaller fixed cost to continue exporting in the future. When both of these restrictions hold simultaneously, the model collapses to the seminal static framework of Melitz (2003), in which exporting simply requires a per-period fixed cost that does not depend on a firm's current export status. Additionally, the theory provides a micro-foundation for the new-exporter dynamics model of Alessandria et al. (2021b), in which exporters accumulate customers exogenously at the same rate in all markets, although this model does not nest within the theory precisely.

To provide motivation and empirical support for my theory, I use microdata on Brazilian exports during the period 1996–2007. The cross-sectional distribution and life-cycle dynamics of Brazilian exporters are consistent with evidence that has been documented elsewhere in the literature. There are more small exporters in large, rich markets, and sales to these markets are more concentrated among the largest exporters, as shown by Arkolakis (2010) and Eaton et al. (2011). New entrants to an export market sell less than incumbents and are more likely to exit, as in Ruhl and Willis (2017), Gumpert et al. (2020), and Alessandria et al. (2021b), and exporters that remain active for longer spells sell more on entry and grow faster, as reported by Fitzgerald et al. (2023). As described above, the model proposed in this paper accounts for all of these facts. I use these data to show that the facts about exporter dynamics vary with market characteristics, just as the cross-sectional facts do. In "easy" destinations—large, rich, and/or close markets like the United States—turnover is lower and new exporters are smaller and exit more frequently relative to incumbents than in "hard" destinations—small, poor, and/or distant countries such as Vietnam. Moreover, the differences between the sales trajectories observed over longer and shorter export spells are more pronounced in harder destinations, and the model accounts for these facts as well.

To study the model's ability to quantitatively account for the facts at hand, I calibrate it to the facts described above using indirect inference. I target only a subset of the facts described in my empirical analysis, but the calibrated model is able to replicate all the other facts as well—something that none of the existing models described above can do. In the static market penetration model of Arkolakis (2016), turnover is too frequent; new exporters survive too often relative to incumbents; and exporters' sales grow too much over the course of their export spells, particularly in markets with low export participation. In the sunk cost model of Das et al. (2007), exports are not concentrated enough among top exporters; new exporters are too large and too likely to survive compared to incumbents; and sales actually fall with time in a market rather than rising. The model of Alessandria et al. (2021b) that incorporates exogenous new-exporter dynamics fares better at generating new exporters that look less like incumbents, but still fails to produce the patterns in sales growth over export spells—and the differences in these patterns across destinations—observed in the data. All of these other frameworks, however, capture some of the observed variation across markets in sales concentration, overall turnover, and new exporters' sizes and exit rates. This indicates that while customer accumulation is key to accounting for all the facts at hand, the distribution and dynamics of firms' exogenous characteristics (productivity and demand) also play important roles.

After calibrating the model, I explore its aggregate implications by simulating the transition dynamics that follow permanent trade reforms. I find that trade grows more in the long run but takes longer to adjust in harder markets than in easier ones. In comparison, the static market penetration model of Arkolakis (2016) generates similar long-run predictions but cannot account for the adjustment process. On the other hand, trade adjusts slowly in the sunk cost model of Das et al. (2007) and the exogenous new-exporter dynamics model of Alessandria et al. (2021b), but these models do not generate material differences in long-run trade growth or transition dynamics across markets. My framework, by synthesizing these two approaches, generates both slow adjustments and differences in dynamics across markets.

This paper makes several contributions to the literature. A number of empirical studies, such as Ottaviano and Mayer (2007), Eaton et al. (2011), and Bernard et al. (2012), have documented that export participation and the cross-sectional distribution of export sales vary with the characteristics of foreign markets. Arkolakis (2010) accounts for these facts by developing a model in which exporting is more cost-effective in larger markets but that serving additional customers in a given market becomes more and more costly, which implies that larger markets have higher export participation rates but also more small exporters. Other studies that focus on the life-cycle dynamics of exporting firms have documented that new exporters sell less than incumbents, are more likely to exit, and grow more rapidly (Bernard and Jensen, 2004; Eaton et al., 2007; Fitzgerald et al., 2023; Ruhl and Willis, 2017; Gumpert et al., 2020; Alessandria et al., 2021b). In this paper, I show that these facts about dynamics vary with export destinations' characteristics, just as the cross-sectional facts do, and that a "dynamicization" of the Arkolakis (2010) model explains both sets of facts simultaneously.

In the quantitative literature, models in which firms face large sunk costs when entering an export market and small costs to continue exporting are often used to study the microeconomic dynamics of export participation and to analyze the macroeconomic implications of these dynamics (Das et al., 2007; Alessandria and Choi, 2007; Ruhl, 2008; Ruhl and Willis, 2017; Alessandria and Choi, 2016; Alessandria et al., 2021b). These models explain why trade flows respond less strongly to trade reforms and other shocks in the short run than in the long run, and suggest that the dynamic gains from trade may differ substantially from the long-run gains. However, these models cannot account for the gradual growth in sales that occurs over an exporter's tenure in a market or the fact that new entrants are more likely to exit, except for variants like Ruhl and Willis (2017) and Alessandria et al. (2021b) in which demand is assumed to grow exogenously with time in a market. My model of market penetration dynamics generates this growth as an endogenous outcome, and accounts for the observed variation in this growth across firms as well as across markets. It also provides new insights about how trade adjustment dynamics depend on market characteristics: in smaller, poorer markets, trade is more elastic in the long run, but also takes longer to converge.

The most similar papers in terms of modeling methodology are Fitzgerald et al. (2023) and Piveteau (2020), both of which feature endogenous customer accumulation. My model has several key advantages. First, neither paper explains why entrants are smaller than incumbents or why the relative size of entrants varies across markets; in both papers, all entrants start exogenously with the same number of customers in all markets regardless of productivity or demand for their products. Second, both papers require sunk entry costs and fixed continuation costs on top of customer accumulation costs to generate entry and exit, as well as requiring exogenous variation in these costs across firms and across markets to match the data. In my model, extensive-margin decisions are driven solely by the marginal cost of serving the first customer in a market, which varies endogenously across firms, over time, and across markets. Third, the parsimony of my approach makes it more amenable to quantitative analysis. In fact, it is tractable even in general equilibrium; I used an early version of the model in a multi-country DSGE environment to study the consequences of uncertainty about Brexit in Steinberg (2019). Most importantly, though, my approach accounts for and explains variation in exporter performance dynamics across destinations, whereas Fitzgerald et al. (2023) and Piveteau (2020) do not explore this variation at all.

Together, these contributions fill a gap highlighted by Alessandria et al. (2021a) in their recent review of the trade dynamics literature: "[T]he literature has largely avoided the treatment of a firm's dynamic decisions across multiple destinations. The literature on (static) quantitative trade and firm heterogeneity has focused on the impact of geography on [exporting] costs. Merging these two approaches is a relatively unexplored, but promising, avenue of future research."

2 Empirical evidence from Brazil

I begin my analysis by using Brazilian microdata to document a set of facts about the distribution and dynamics of exporting firms. Some of these findings have already been established in the literature, while others are entirely novel. Collectively, they provide a comprehensive overview of how exporters' performance varies over time, across markets, and along both of these dimensions together. Importantly, they show that both the cross-sectional distribution and the life-cycle dynamics of exporting firms vary with export markets' characteristics. This provides both motivation and empirical support for the theory developed in the next section.

The data source is a record of all Brazilian firms' monthly foreign sales from 1996 to 2008. For each transaction, the dataset includes the destination country, the value of the shipment in U.S. dollars, the year and month of the transaction, an eight-digit product code, and a unique firm identifier. I restrict attention to manufacturing industries and aggregate the data to the firm-year-destination level. I exclude destinations that are served by fewer than 20 firms per year following Fernandes et al. (2016), and I exclude the year 2008 to avoid issues related to the Great Recession that might affect exporter dynamics. I combine these data with destination-level information from the CEPII Gravity Database on population, income per capita, and trade costs (measured as residuals from a standard gravity regression).[1]

First, I study how the following factors vary across export markets: the distribution of exporters' sales; the likelihood of exit; and the performance of new exporters relative to incumbents. Second, I study how sales and survival grow after exporters enter a market, and how these trajectories differ across markets. In the online appendix, I report additional results about the differences in performance within individual exporters' portfolios of destinations; these results are tangential to the main points of this paper but provide further evidence on other dimensions of exporter performance that vary across markets.

The data used in these analyses cannot legally be distributed, but all programs and intermediate datasets needed to reproduce the results reported in this section are available on GitHub at `https://github.com/joesteinberg/dyn-mkt-pen`. I have also analyzed similar data on Mexican and Peruvian exporters from the World Bank's Exporter Dynamics Database (Fernandes et al., 2016). These publicly available datasets are of somewhat lower quality than the Brazilian data, containing fewer firms and covering shorter time periods. Nevertheless, all the results documented in this section about Brazilian exporters also apply to Mexican and Peruvian exporters. This corroborates the findings I report in this section, indicating that they are robust to variation in conditions in the exporting country. These additional results are also available in the online appendix.

2.1 Concentration, turnover, and new-exporter dynamics across markets

The first part of my empirical analysis studies exporter performance at the market level. For each export destination in my sample, I compute three cross-sectional measures: the number of firms that export to that destination, the share of sales accounted for by the top 5% of exporters in that destination, and the average number of other markets served by firms that export to that destination. I also compute three measures of exporter dynamics: the exit rate, the average sales

[1] I have also studied exporter performance at the industry level, including industry effects to control for variation in the industrial composition of exports across destinations, but this requires restricting attention to a very small set of industry–destination pairs with at least 20 firms per year. I have also estimated specifications in which tariffs, distance, and other gravity variables are included directly as independent variables rather than indirectly through a gravity residual. The results from these alternative specifications are all in line with the results reported in this paper.

of entrants relative to incumbents, and the exit rate of entrants relative to incumbents. Each of these measures is computed at the destination-year level.

Panel (a) of Table 1 reports summary statistics for each of these measures. They all vary considerably across destinations. The most popular export market, Argentina, has an export participation rate more than 100 times greater than the least popular market, Vietnam. The top-five share ranges from 0.29 to 0.84, indicating that exports are highly concentrated among the largest exporters in some markets and more evenly distributed in others. The average number of other markets exporters serve varies from less than eight to almost 30, which shows that some markets are served only by firms with a large portfolio of other markets. The last three columns of the table show that exporter dynamics also vary across markets. Exit rates range from as low as 19 percent to as high as 55 percent. Finally, entrants are smaller and more likely to exit than incumbents in all markets, as documented by Ruhl and Willis (2017) and Gumpert et al. (2020), but these differences are more pronounced in some markets and more muted in others.

Table 1: Market-level measures of exporter performance

Statistic/coefficient	Num. exporters	Top-five share	Avg. num. dests.	Exit rate	Entrant rel. size	Entrant rel. exit rate
(a) Summary statistics						
Mean	591	0.60	16.70	0.36	0.38	0.26
Min.	23	0.29	7.87	0.19	0.09	0.06
Max.	3,721	0.84	27.90	0.55	1.15	0.38
Std. dev.	829	0.14	4.95	0.07	0.24	0.06
(b) Associations with market characteristics						
log GDPpc	0.581	0.051	−1.663	−0.005	−0.085	0.008
	(0.061)§	(0.008)§	(0.385)§	(0.005)	(0.017)§	(0.004)*
log population	0.422	0.047	−1.163	−0.009	−0.042	0.003
	(0.045)§	(0.007)§	(0.218)§	(0.003)‡	(0.013)‡	(0.004)
log trade barrier	−1.096	−0.065	2.686	0.033	0.103	−0.019
	(0.074)§	(0.009)§	(0.478)§	(0.005)§	(0.018)§	(0.005)§
Num. observations	627	627	627	627	627	627
R^2	0.70	0.42	0.42	0.73	0.13	0.40

Notes: Panel (a) reports summary statistics. Panel (b) reports associations with market characteristics. All specifications include year fixed effects. Standard errors are clustered at the market level. §, ‡, and † denote significance at 0.1%, 1%, and 5% levels, respectively. Variable definitions are as follows. Num. exporters: number of firms that export to a given market. Top-5 share: fraction of total exports by firms with exports above 95th percentile. Avg. num. markets.: average number of other markets served by firms that export to a given destination. Exit rate: Fraction of firms that export to a given market in one year but not the next year. Entrant rel. size: average exports of new entrants (firms in their first year of exporting to a given market) divided by average exports of incumbent firms. Entrant rel. exit rate: exit rate (defined above) for entrants minus the exit rate of incumbents.

It is well known that export participation is higher and that sales are more concentrated among top exporters in larger, richer markets (see, e.g., Ottaviano and Mayer, 2007; Arkolakis, 2010; Eaton et al., 2011). To verify the existence of these relationships in the Brazilian data and to determine whether exporter dynamics also vary with market characteristics, I estimate regressions of the form

$$M_{j,t} = \alpha + \beta \log L_{j,t} + \gamma \log Y_{j,t} + \delta \log \tau_{j,t} + f_t + \epsilon_{j,t} \quad (1)$$

where the dependent variable, $M_{j,t}$, is a measure of exporter performance (e.g. the top-five share or the exit rate) in market j in year t. The independent variables are the market's characteristics: population, $L_{j,t}$; income per capita, $Y_{j,t}$; and trade costs, $\tau_{j,t}$. The variable f_t is a year fixed effect that controls for multilateral trends such as Brazilian business cycles and exchange rate depreciation.

Panel (b) of Table 1 reports the results from these estimations. The first column shows that export participation is increasing in population and income per capita and decreasing in trade barriers. This is not terribly surprising, but it suggests that the export participation rate can be thought of as a convenient, one-dimensional summary of the difficulty of exporting to a market, which will prove useful in the analyses that follow. The next two columns show that the cross-section of exporting firms varies with the characteristics of export markets as documented in other studies. In "easy" markets with large, rich populations, and/or low trade barriers, exports are more concentrated and the average exporter serves only a few other destinations. The last three columns of the table show that exporter dynamics, too, vary with market characteristics. In easier markets turnover is lower and new exporters are smaller and more likely to exit relative to incumbents, whereas in harder markets turnover is higher and new-exporter dynamics are less pronounced.

The results described above confirm several established facts and also provide new insights about variation in exporter dynamics across markets. The two lists below formally summarize these results. In order to highlight this paper's empirical contributions, new facts documented for the first time herein are explicitly distinguished from facts that have previously been documented by other studies. It bears mentioning that some previously-established facts have been documented only for advanced economies—for example, the studies listed next to the items in Facts 1 look at data from France and Belgium—and my analysis shows that these facts apply to emerging economies such as Brazil as well.

Facts 1: Cross-sectional

1. Larger, richer, and lower-cost markets attract more exporters (Ottaviano and Mayer, 2007; Arkolakis, 2010; Eaton et al., 2011).
2. Exports are more concentrated among top exporters in markets with greater export participation (Arkolakis, 2010; Eaton et al., 2011).
3. The average exporter to a market with greater export participation serves a larger number of other markets than the average exporter to a market with lower export participation (Ottaviano and Mayer, 2007).

Facts 2: New-exporter dynamics

1. New exporters sell less than incumbents (Ruhl and Willis, 2017; Gumpert et al., 2020).
2. New exporters exit more frequently than incumbents (Ruhl and Willis, 2017; Gumpert et al., 2020).
3. Exporters are less likely to exit from markets with greater export participation (new finding).
4. New exporters sell less relative to incumbents in markets with greater export participation (new finding).
5. New exporters are more likely to exit relative to incumbents in markets with greater export participation (new finding).[2]

2.2 Sales and survival over export spells

The second part of my empirical analysis studies exporter performance at the firm level. I begin this part of the analysis by studying how exporters grow over time after they enter a new market. Following Fitzgerald et al. (2023), I group firms by the number of consecutive years that they export to a particular destination before exiting—the duration of an export "spell"—and then estimate the sales trajectories of firms in each group. Formally, I estimate the following regression:

$$\log ex_{i,j,t} = \sum_{m=1}^{6} \sum_{n=1}^{m} \beta_{m,n} \mathbb{1}_{\{\text{spell}_{i,j,t}=m\}} \mathbb{1}_{\{\text{tenure}_{i,j,t}=n\}} + \gamma \mathbb{1}_{\{\text{censored}_{i,j,t}\}} + f_{i,t} + f_{j,t} + \epsilon_{i,j,t} \tag{2}$$

where $ex_{i,j,t}$ is firm i's exports to market j in year t, $\text{spell}_{i,j,t}$ indicates the eventual length of the firm's current export spell in that market, $\text{tenure}_{i,j,t}$ indicates the number of years the firm has consecutively exported so far during that spell, and $\text{censored}_{i,j,t}$ indicates whether the current spell was censored by the end of the dataset.[3] I top-code spell length at 6 years (the shortest observation window for a destination in my dataset). I include firm–year and market–year fixed effects, denoted by $f_{i,t}$ and $f_{j,t}$ respectively.

Panel (a) of Figure 1 shows the results of this analysis, which mirror those of Fitzgerald et al. (2023) for Ireland. Longer export spells are associated with greater sales upon entry. Sales in the first year of a three-year spell are about 55% higher than a one-year spell's sales, and sales in the first year of a six-year spell are almost 100% higher. Longer spells are also associated with more sales growth. Sales grow by about 30% over the course of a three-year spell, versus almost 100% over the course of a six-year spell. Note that for spells of five years or less, sales fall in the last year of the spell; for two-year spells, sales actually fall immediately after entry. This suggests that these

[2] This is similar to the findings of Gumpert et al. (2020), who show that new exporters are more likely to exit from small, distant markets but do not compare new exporters' survival to that of incumbents. The fact that overall exit rates and new exporters' exit rates vary in opposite ways with export participation makes this distinction nontrivial.

[3] As Fitzgerald et al. (2023) point out, the most successful exporters may remain in a market long after the end of the dataset. Without controlling for the fact that these exporters' spells are thus censored, the coefficient $\beta_{6,6}$ would likely be biased upwards.

Figure 1: Effects of tenure and spell length on exporters' sales

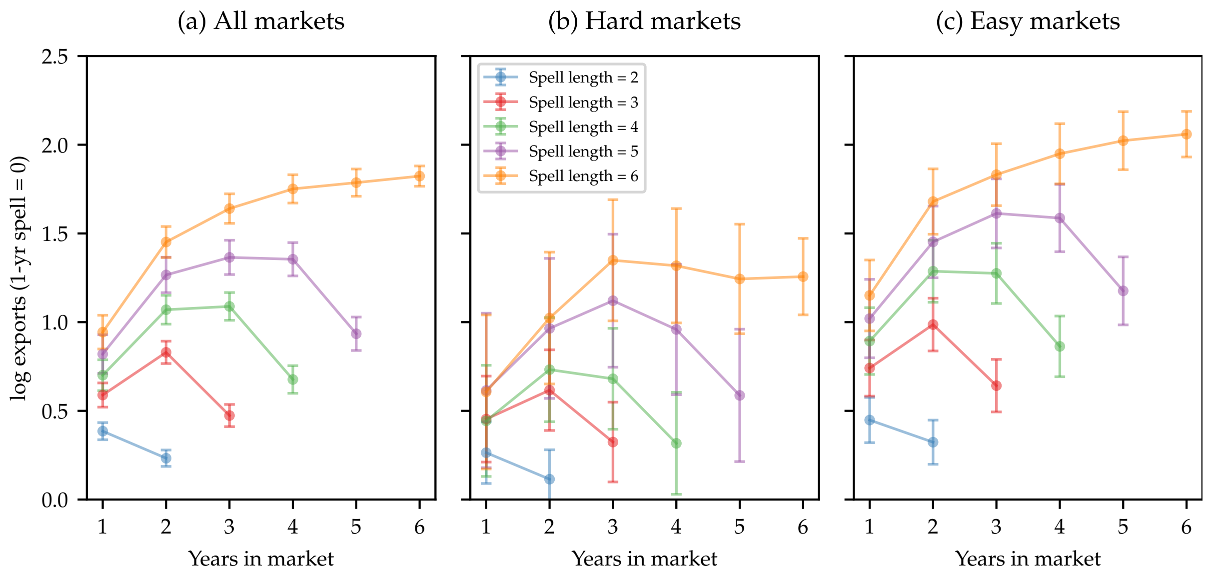

Notes: Panel (a) shows estimates of $\beta_{m,n}$ from (2). Panel (b) shows estimates of $\beta_{m,n,g}$ from (3) for markets in the bottom 50% of export participation, and panel (c) shows estimates for markets in the top 10%. Each line shows $\beta_{m,1}$, $\beta_{m,2}$, …, $\beta_{m,6}$ (or $\beta_{m,1,g}$, $\beta_{m,2,g}$, … in the second two panels) for a set value of m. Vertical bars with horizontal caps show 95% confidence intervals.

shorter spells are ultimately terminated by persistent negative shocks. Although these findings are not new, documenting that these patterns hold in an emerging economy such as Brazil as well as economies like Ireland provides additional evidence of their generality.

To study how these sales trajectories differ across markets, I first split the dataset into two groups: "hard" markets below the 50th percentile of export participation ($g = 1$) and "easy" markets above the 90th percentile ($g = 2$). I then estimate the following specification, where the spell length-tenure effect is interacted with an indicator for the market's group:[4]

$$\log ex_{i,j,t} = \sum_{g=1}^{2}\sum_{m=1}^{6}\sum_{n=1}^{m} \beta_{m,n,g}\mathbb{1}_{\{\text{spell}_{i,j,t}=m\}}\mathbb{1}_{\{\text{tenure}_{i,j,t}=n\}}\mathbb{1}_{\{\text{group}_j=g\}} + \gamma\mathbb{1}_{\{\text{censored}_{i,j,t}\}} + f_{i,t} + f_{j,t} + \epsilon_{i,j,t}. \tag{3}$$

The results for hard and easy destinations are shown in panels (b) and (c) of Figure 1, respectively. Comparing the two panels, we see that the Fitzgerald et al. (2023) findings reproduced in panel (a) are more pronounced in easy destinations than in hard ones. In panel (c), which shows the results for easy destinations, there is more variation in new entrants' sales conditional on spell duration and more growth in sales over the course of a spell than in panel (b), which shows the results

[4] In a preliminary version of the paper in which I did not include firm fixed effects, I split the data into two subsamples (one for hard markets and another for easy markets) and ran specification (2) for these two subsamples separately. Although this approach may be more straightforward, it cannot be used when including firm fixed effects, as the same firm's average sales will differ across groups, and these fixed effects would absorb some of the difference in that firm's sales across groups. In order to include firm fixed effects, a single specification must be estimated on the entire sample.

for hard destinations. The largest differences are in the results for six-year export spells.[5] In easy markets, sales in the first year of a six-year spell are more than 100% higher than a one-year spell's sales, versus about 60% for hard markets. Similarly, sales eventually grow by more than 100% over the course of a six-year spell in easy markets, versus about 60% in hard markets.

I continue my firm-level analysis by studying how the likelihood of continuing as an exporter depends on tenure in a market. Following Fitzgerald et al. (2023) and Ruhl and Willis (2017), I use a linear probability model of the form

$$\mathbb{1}_{\{\text{exit}_{i,j,t}=1\}} = \sum_{n=1}^{6} \beta_n \mathbb{1}_{\{\text{years in market}_{i,j,t}=n\}} + f_{i,t} + f_{j,t} + \epsilon_{i,j,t}. \tag{4}$$

As before, I include firm–year and market–year fixed effects, but here I restrict the sample to spells that are not censored by the end of the dataset. The coefficient β_n indicates how much more likely an exporter that has survived for n years is to exit than a firm that has just begun to export. Panel (a) of Figure 2 shows the results of this analysis, which are consistent with the aforementioned studies' findings: exit from a given market becomes less likely the longer a firm has been exporting to that market. To study whether the effects of tenure on survival differ across markets, I interact the tenure indicator with the group indicator, yielding the following specification:

$$\mathbb{1}_{\{\text{exit}_{i,j,t}=1\}} = \sum_{g=1}^{2} \sum_{n=1}^{6} \beta_{n,g} \mathbb{1}_{\{\text{tenure}_{i,j,t}=n\}} \mathbb{1}_{\{\text{group}_j=g\}} + f_{i,t} + f_{j,t} + \epsilon_{i,j,t}. \tag{5}$$

Panel (b) of Figure 2 shows the results. Conditional on tenure, firms are always more likely to exit from hard markets than easy markets. This is consistent with Table 1, which shows that the overall exit rate is higher in hard markets. However, exit rates decline with tenure at a similar rate in both groups of markets; the level effect of tenure on survival is higher in easy markets, but the effect of a change in tenure on survival is not.

As in section 2.1, these results confirm previous findings in the literature and show that they apply to both emerging and advanced economies, but they also reinforce my new findings that the dynamics of exporters' performance vary with market characteristics. The next list of facts summarizes these results.

Facts 3: Effects of tenure on sales and survival

1. Exporters that achieve longer spells sell more upon entry and grow faster (Fitzgerald et al., 2023).

2. Exporters with longer tenures exit less often (Ruhl and Willis, 2017; Gumpert et al., 2020).

[5] In the online appendix, I show that all the differences in $\beta_{m,n,g}$ between easy and hard markets are highly statistically significant. The largest p-value is 0.15%, and the vast majority are smaller than 0.05%.

3. The effects of tenure and spell length on sales are larger in markets with greater export participation (new finding).
4. Exporters exit less often from markets with greater export participation after conditioning on tenure (new finding).

Figure 2: Exit rates conditional on tenure

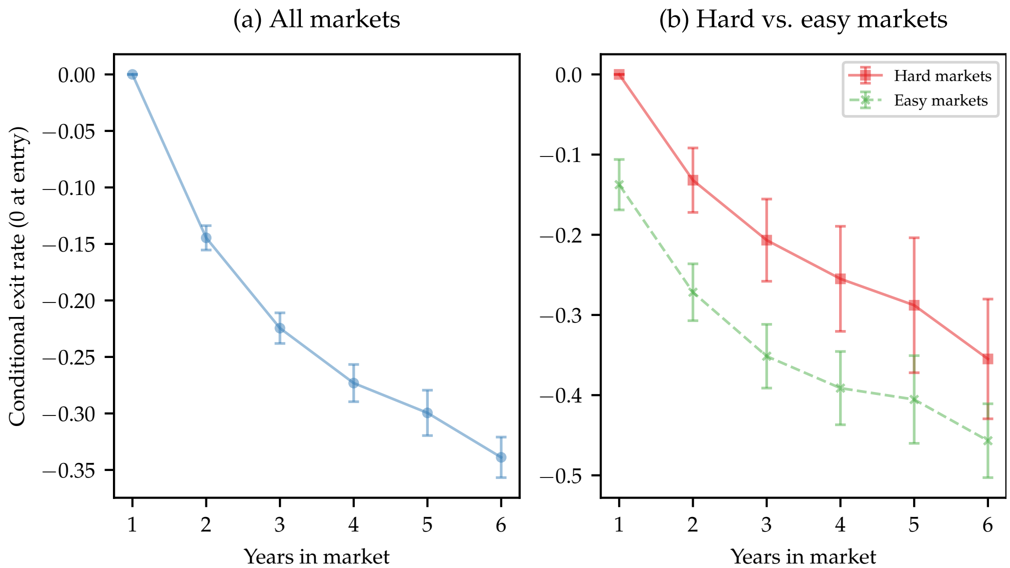

Notes: Panel (a) shows estimates of β_n from (4). Panel (b) shows estimates of $\beta_{n,g}$ from (5). Solid red line with square markers in panel (b) shows estimates for markets in the bottom 50% of export participation, and dashed green line with 'x' markers shows estimates for destinations in the top 10%. Vertical bars with horizontal caps show 95% confidence intervals.

3 Model of export market penetration dynamics

The model environment consists of one exporting country and J importing countries (i.e., markets). indexed by $j = 1, \ldots, J$. The exporting country is populated by a continuum of firms that produce differentiated goods using constant-returns-to-scale technologies. Each market is populated by a measure L_j of identical consumers with income per capita Y_j and constant-elasticity-of-substitution preferences. Trade barriers are captured by iceberg transportation costs, τ_j, which also vary across markets. As in Arkolakis (2010), firms are heterogeneous in their customer bases in each market, which they can increase endogenously by advertising. The costs of retaining old customers and acquiring new ones depend on a firm's current customer base, which leads firms to gradually accumulate foreign customers over time.

As in Arkolakis (2010) and Ruhl and Willis (2017), I assume that importing countries are large relative to the exporting country so that aggregate prices and quantities in the former are independent of outcomes in the latter. I also assume that export activities are small relative to the total size of the exporting country's economy so that the wage in the exporting country, normalized to one without loss of generality, is independent of export-sector outcomes. Finally, I assume for the moment that all aggregate variables, including trade barriers and other destination characteristics, are constant to economize on notation; this section restricts attention to the model's stationary equilibrium. In my quantitative analysis, however, I also analyze transition dynamics that follow changes in trade barriers.

3.1 Firm characteristics

There is a unit measure of firms in the exporting country that produce differentiated varieties according to constant-returns-to-scale technologies.[6] Firms are heterogeneous in productivity, $a \in \mathbb{R}_{++}$; demand in each market, $\mathbf{z} = (z_1, z_2, \ldots, z_J) \in \mathbb{R}^J_{++}$; and the fraction of consumers in each market to which they can sell, $\mathbf{m} = (m_1, m_2, \ldots, m_J) \in [0,1]^J$.

Productivity is common to all markets, and evolves according to a Markov process with transition function $G(a', a)$. Demand in each market j evolves independently according to a Markov process with transition function $H(z'_j, z_j)$. A firm's customer base in each market is chosen endogenously in a manner that I describe below. Each period, a firm has a chance $1 - \delta(a)$ of dying, which I allow to depend on its productivity to capture the fact that smaller firms shut down more frequently (Alessandria et al., 2021b). When a firm dies, it is replaced by a new firm with productivity and demand shocks drawn from their respective ergodic distributions, $\bar{G}(a)$ and $\bar{H}(z_j)$. Newborn firms have zero customers in all export markets.

3.2 Export demand, pricing, and profits

Firms compete monopolistically as in Melitz (2003) and Chaney (2008). Market j's demand for a firm's product depends on the market's characteristics, L_j and Y_j; the firm's price in that market, p; the firm's demand shock in that market, z; and the fraction of consumers in that market to which the firm can sell, m. Conditional on purchasing the firm's product, an individual consumer in market j has a standard downward-sloping demand function:

$$c_j(z, p) = L_j Y_j (p/z)^{-\theta}, \tag{6}$$

where the parameter θ is the elasticity of substitution between varieties.[7] Total demand for the firm's product in market j depends on the firm's price as well as the number of customers it can serve, as in Arkolakis (2010) and Eaton et al. (2011):

$$q_j(z, m, p) = m c_j(z, p). \tag{7}$$

[6] I abstract from firm creation in this paper. My focus is on exporter performance in bilateral trade relationships, and the interpretation of this assumption is that the prospect of exporting to a single destination, even a large one, is too small to affect firm creation incentives. Studying the relationship between bilateral trade and firm creation is a promising avenue of investigation that could yield important insights, but I leave this for future research.

[7] The price level in each market is normalized to one; Y_j can be interpreted as purchasing power-adjusted income per capita.

Conditional on its productivity, a, demand, z, and customer base, m, a firm chooses its price in each market j to maximize profits:

$$\pi_j(a,z,m) = \max_p \left\{ pq_j(z,m,p) - \frac{\tau_j q_j(z,m,p)}{a} \right\}. \tag{8}$$

The optimal price is given by the standard constant-markup solution,

$$p_j(a) = \frac{\theta}{\theta-1}\frac{\tau_j}{a}. \tag{9}$$

The firm's exports to market j and associated profits can be written as

$$ex_j(a,z,m) = \left(\frac{\theta}{\theta-1}\right)^{1-\theta} mL_jY_j\tau_j^{1-\theta}(az)^{\theta-1} \tag{10}$$

and

$$\pi_j(a,z,m) = \frac{1}{\theta}\left(\frac{\theta}{\theta-1}\right)^{1-\theta} mL_jY_j\tau_j^{1-\theta}(az)^{\theta-1} := \tilde{\pi}_j m(az)^{\theta-1}, \tag{11}$$

respectively.

3.3 Market penetration dynamics

A firm's customer base in each market evolves over time as it attracts new customers and loses some of its old ones. Consider a firm with current customer base m in a market j. Let $n \in [0, 1-m]$ denote the number of new customers the firm attracts as a share of that market's overall customer population, and let $o \in [0, m]$ denote the number of old customers the firm retains, measured in the same way. Then the firm's customer base in the next period is given simply by

$$m' = n + o. \tag{12}$$

Note, though, that as the firm's current customer base grows, the pool of potential new customers shrinks while the pool of old customers that can potentially be retained grows. I use the term "potential entrant" to refer to a firm with $m = 0$ and the term "incumbent" to refer to a firm with $m > 0$. The terms "entrant" and "new exporter" equivalently refer to a potential entrant that chooses $m' > 0$; a new exporter becomes an incumbent in the next period.

Customer attraction and retention both depend on a firm's advertising efforts. I use s (for search) to denote advertising targeted at new customers and r (for retention) to denote advertising targeted at old customers. Following Arkolakis (2010), the marginal effect of search effort on customer attraction is increasing in the total number of potential new customers, $(1-m)L_j$, and decreasing in the fraction of potential new customers a firm successfully attracts, $n/(1-m)$:

$$n_j'(s) = \psi_n L_j^{-\alpha_n} (1-m)^{-\beta_n} \left(\frac{1-m-n_j(s)}{1-m} \right)^{\gamma_n}. \tag{13}$$

The parameter α_n governs returns to population size in advertising to new customers. The smaller α_n, the easier it is to attract new customers in larger countries. Similarly, β_n governs the returns to scale with respect to the size of the pool of potential new customers in a particular market. The smaller β_n, the easier it is for a firm to attract new customers when its current market penetration is low. I refer to α_n and β_n as the macroeconomic and microeconomic returns to market size parameters, respectively. The parameter γ_n represents the degree of diminishing returns in advertising to new customers. The higher γ_n, the fewer additional new customers are reached by each additional unit of search advertising. Finally, ψ_n is the efficiency of advertising to new customers. The higher ψ_n, the lower the average cost of customer attraction.

Similarly, the marginal effect of retention effort is increasing in the total number of old customers, mL_j, and decreasing in the fraction of old customers the firm successfully retains, o/m:

$$o_j'(r) = \psi_o L_j^{-\alpha_o} m^{-\beta_o} \left(\frac{m-o_j(r)}{m} \right)^{\gamma_o}. \tag{14}$$

The parameters α_o, β_o, γ_o, and ψ_o have similar interpretations to their counterparts above. Note though, that microeconomic returns to market size increase as a firm's customer base, and thus the pool of old customers who can be retained, grows. Differences between the parameters in (13) and (14) allow for the possibility that advertising to old customers works differently than advertising to new customers. For example, it may be that the macroeconomic market size effect is less pronounced ($\alpha_o > \alpha_n$) because advertising to current customers is more analogous to contacting them individually one after another than to mass advertising on the radio or television. It might also be the case that returns to advertising to current customers diminish less rapidly ($\gamma_o < \gamma_n$). Indeed, when I calibrate the model's parameters so that it matches the facts described in section 2, I find precisely these differences.

3.4 Dynamic market penetration costs

As a firm builds its customer base over time, its search and retention costs change. Solving the differential equations (13) and (14) yields the costs of attracting n new customers and retaining o

old customers, respectively:

$$s_j(m,n) = \frac{L_j^{\alpha_n}(1-m)^{\beta_n}}{\psi_n(1-\gamma_n)}\left[1-\left(\frac{1-m-n}{1-m}\right)^{1-\gamma_n}\right], \tag{15}$$

$$r_j(m,o) = \frac{L_j^{\alpha_o} m^{\beta_o}}{\psi_o(1-\gamma_o)}\left[1-\left(\frac{m-o}{m}\right)^{1-\gamma_o}\right]. \tag{16}$$

In what follows, I use $s_{j,n}$ and $r_{j,o}$ to denote the partial derivatives of the advertising cost functions with respect to their second arguments. Although these expressions bear more than a passing resemblance to the market penetration cost function in Arkolakis (2010), they depend not only on the number of customers a firm attracts or retains, but also on the firm's current customer base.

For a firm with current customer base m that wishes to expand (or perhaps shrink) its customer base to m', the total cost of customer attraction and retention—the market penetration cost in the terminology of Arkolakis (2010)—is given by the solution to the static problem

$$f_j(m,m') = \min_{n\in[0,1-m],\ o\in[0,m]} \left\{s_j(m,n)+r_j(m,o)\right\} \text{ subject to } m' = n+o. \tag{17}$$

I use $n_j(m,m')$ and $o_j(m,m')$ to denote the optimal policy functions for customer attraction and retention, respectively. The solution to this problem can be characterized as follows. For entrants, who have no old customers to retain, the market penetration cost is equal to the attraction cost: $f_j(0,m') = s_j(0,m')$, $n_j(0,m') = m'$, and $r_j(0,m') = 0$. For incumbents with growing customer bases such that the marginal cost of attracting the last new customer is lower than the marginal cost of retaining the first old customer, then no old customers should be retained: If $m' > m$ and $s_{j,n}(m,m') < r_{j,o}(m,0)$, then $n_j(m,m') = m'$ and $o_j(m,m') = 0$. For incumbents with shrinking customer bases such that the marginal cost of retaining the last old customer is lower than the marginal cost of attracting the first new customer, then no new customers should be attracted: If $m' < m$ and $r_{j,o}(m,m') < s_{j,n}(m,0)$, then $n_j(m,m') = 0$ and $o_j(m,m') = m'$. In all other cases, the marginal attraction and retention costs are equal at the optimum: $s_{j,n}(m,n) = r_{j,o}(m,o)$.

3.5 Key properties of the market penetration cost function

The market penetration cost function (17) has several key properties that allow the model to account for both cross-sectional and life-cycle facts about exporters documented in section 2. In what follows, I use $f_{j,m'}$ to denote the first partial derivative of (17) with respect to end-of-period market penetration, $f_{j,m'm'}$ to denote the second partial, and $f_{j,mm'}$ to denote the cross partial.

First, the marginal cost of accessing the first customer is always strictly positive regardless of a firm's current customer base:

$$f_{j,m'}(m,0) > 0,\ \forall m. \tag{P1}$$

This property generates the extensive margin of trade: firms with sufficiently low productivity and/or demand will find the cost of accessing the first customer prohibitive. In this dynamic context, this property delivers endogenous exit as well as endogenous entry: some incumbent exporters will opt not to retain any of their old customers or attract any new ones.[8,9]

Second, the marginal market penetration cost is increasing, and it is impossible to saturate the market:

$$f_{j,m'm'}(m,m') > 0\ \forall m,m',\ \lim_{m'\to 1} f_{j,m'}(m,m') = \infty\ \forall m. \tag{P2}$$

This property implies that higher-productivity and/or higher-demand firms will attract more new customers and retain more of their old customers, but that even the best firms will never fully penetrate an export market even after accumulating customers over many periods.[10]

Third, the market penetration cost is decreasing in a firm's current customer base, both overall and at the margin:

$$f_{j,m}(m,m') < 0,\ f_{j,mm'}(m,m') < 0\ \forall m,m'. \tag{P3}$$

This property mirrors a common result from sunk-cost models, in which the cost of entering a foreign market is higher than the cost of continuing to serve it. Here, it implies that firms derive two benefits from expanding their customer bases: increased sales in the present and reduced exporting costs in the future.

Fourth, measured relative to purchasing power, exporting is more expensive on the whole and at the margin in smaller, poorer destinations:

$$\frac{\partial f_j(m,m')/(L_jY_j)}{\partial X_j} < 0,\ \frac{\partial f_{j,m'}(m,m')/(L_jY_j)}{\partial X_j} < 0\ \forall m,m',\ X_j \in \{L_j, Y_j\}. \tag{P4}$$

3.6 Optimal market penetration dynamics

Once the firm has determined the most cost-effective way to increase (or decrease) its market penetration, it chooses how much it should do so in order to maximize the present discounted

[8] Moreover, if the marginal cost of attracting the first new customer for a potential entrant, $s_{j,n}(0,0)$, exceeds the marginal cost of retaining the first old customer for a typical incumbent, the model will generate exporter hysteresis as in Baldwin (1992): the average entrant will be more productive than the average incumbent.

[9] This is where modeling the distinction between retaining old customers and attracting new ones is especially crucial. It is possible to model market penetration dynamics without making this distinction and still match many of the facts documented in this paper, but doing so requires one to assume that exit from the export market is exogenous, or that there is an additional fixed cost of exporting on top of the market penetration costs. See Steinberg (2019) for an early version of this model with no distinction between old and new customers and exogenous exit.

[10] These properties are inherited from similar properties of the underlying attraction and retention costs, s_j and r_j. Note that because it is impossible to retain all old customers—$r_{j,o}$ goes to infinity as o approaches 1—all firms experience customer turnover, even firms with growing customer bases.

value of the profits from exporting:

$$V_j(a,z,m) = \max_{m' \in [0,1]} \left\{ \pi_j(a,z,m') - f_j(m,m') + \frac{\delta(a)}{1+R} \mathbb{E}\left[V_j(a',z',m') | a,z \right] \right\} \tag{18}$$

where the parameter R is the firm's discount rate. This formulation of the problem is virtually identical to the Bellman equations in sunk-cost models of exporting such as Das et al. (2007), Alessandria and Choi (2007), and Alessandria et al. (2021b) in which the cost of exporting depends on a firm's current status as an exporter; the only difference is that export status is a continuous variable in my model, rather than binary. I use $m'_j(a,z,m)$ to denote the optimal policy function at this stage.

Using the envelope theorem, the solution to this problem is characterized by the following inequality:

$$f_{j,m'}\left(m,m'\right) \geq \tilde{\pi}_j(az)^{\theta-1} - \frac{\delta(a)}{1+R} \mathbb{E}\left[f_{j,m}(m',m'') | a,z \right] \tag{19}$$

where m' and m'' are shorthand for $m'_j(a,z,m)$ and $m'_j(a',z',m'_j(a,z,m))$, respectively. The left-hand side of this expression is the marginal cost of exporting. The first term on the right-hand side is the marginal increase in flow profits the firm gains from increasing its market penetration. The second term on the right-hand side is the expected change in the cost of exporting in the next period. Note that property (P3) implies that this term is positive and increasing in m': increasing market penetration today reduces the cost of exporting tomorrow. If this condition holds with equality, the firm chooses m' to equate the marginal cost of exporting with the marginal benefit. This property also implies that the policy function is increasing in m: firms with higher market penetration rates at the beginning of the period choose higher market penetration rates at the end of the period. In turn, this implies that firms gradually build up their customer bases over time after entering a market.

If, on the other hand, the marginal cost of attracting or retaining the very first customer, $f_{j,m'}(m,0)$, exceeds the marginal benefit, the firm will exit (if $m > 0$) or not enter (if $m = 0$). Entry is characterized by a threshold $\underline{z}_j(a)$ such that firms with demand shocks below this threshold will choose not to enter:

$$f_{j,m'}\left(0,0\right) = \tilde{\pi}_j(a\underline{z}_j(a))^{\theta-1} - \frac{\delta(a)}{1+R} \mathbb{E}\left[f_{j,m}(0,m'') | a,z \right]. \tag{20}$$

The entry threshold is decreasing in a: high-productivity firms are more likely to enter than low-productivity ones. Exit is characterized by a threshold $\overline{z}_j(a,m)$ such that firms with demand below this threshold will choose to exit:

$$f_{j,m'}\left(\overline{z}_j(a,m),0\right) = \tilde{\pi}_j(az)^{\theta-1} - \frac{\delta(a)}{1+R} \mathbb{E}\left[f_{j,m}(0,m'') | a,z \right]. \tag{21}$$

The exit threshold is decreasing in both z and m. Importantly, the latter implies that larger negative demand shocks are required to induce firms with larger customer bases to exit.

Figure 3 illustrates how the features of the model work together to generate realistic exporter dynamics. Consider a potential entrant with zero customers and a high enough demand shock, z_{hi}, to warrant entering a particular market. Panel (a) shows how the firm's optimal market penetration choice as a new entrant, m_1, is determined. It is shown in the figure as the intersection of the firm's marginal benefit, the horizontal dotted line labelled $\tilde{\pi}_{hi} + Q\mathbb{E}[f_m(m_1, m'')|z_{hi}]$, and the entrant's marginal cost curve, the solid green curve labelled $f_{m'}(0, \cdot)$. Panel (b) shows the firm's policy function as the solid blue curve labelled $m'(a, z_{hi}, \cdot)$; the firm's choice in this period is the point $(0, m_1)$ located on this curve.

Figure 3: Entry, expansion, and exit in the model

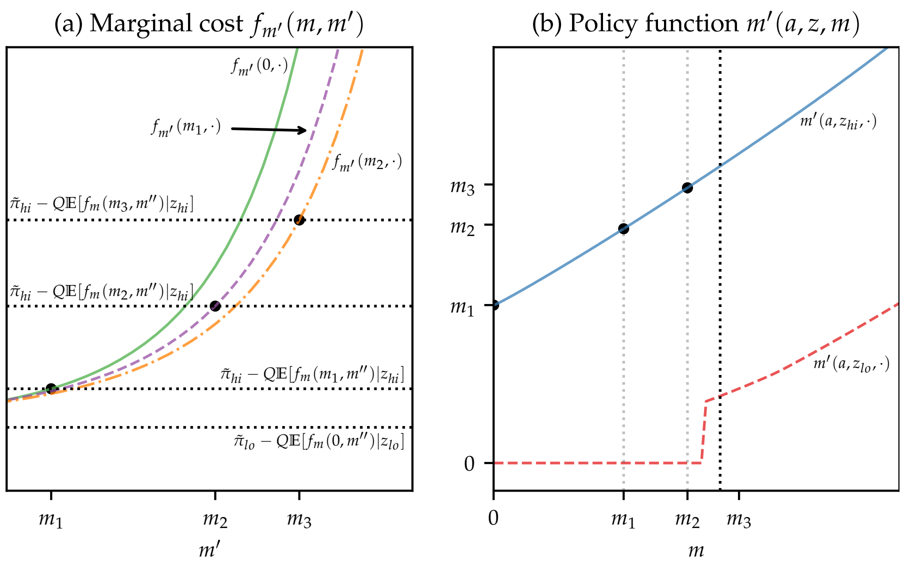

Notes: Panel (a) shows marginal costs and benefits of market penetration with high (z_{hi}) and low (z_{lo}) demand shocks. Solid green curve is the marginal cost for a new entrant ($m = 0$). Dashed purple (orange) curve shows marginal cost after 1 (2) period(s) of exporting. Dotted black lines are marginal benefits, with associated labels using the shorthand $\tilde{\pi}_{hi} = \tilde{\pi}_j(az_{hi})^\theta$, $\tilde{\pi}_{lo} = \tilde{\pi}_j(az_{lo})^\theta$, and $Q = \delta(a)/(1+R)$. Panel (b) shows the market penetration policy functions for high (solid blue) and low (dashed red) demand shocks. Market subscripts j omitted in all labels for brevity.

In period 1, the firm's marginal cost curve shifts outward to the purple curve in panel (a) labelled $f_{m'}(m_1, \cdot)$ due to property (P3). The firm's optimal market penetration choice in this period, m_2, is given by the intersection of this new marginal cost curve and the firm's marginal benefit. Note, though, that the marginal benefit has increased to $\tilde{\pi}_{hi} - Q\mathbb{E}[f_m(m_2, m'')|z_{hi}]$. This is also due to property (P3): the increase in market penetration from m_1 to m_2 has reduced its expected marginal market penetration cost in the next period. The firm's choice in period 1 is shown in panel (b) as the point (m_1, m_2) on the firm's policy function. In period 2, the firm's marginal cost curve shifts outward again, to the dashed orange curve labeled $f_{m'}(m_2, \cdot)$ in panel (a). Suppose, however, that the firm receives a bad demand shock, z_{lo}, such that the marginal benefit of exporting is now lower than the marginal cost, shown by the lowest horizontal dotted line in panel (a) labelled $\tilde{\pi}_{lo} + \beta\mathbb{E}[f_m(0, m'')|z_{lo}]$. Instead of continuing to expand, the firm decides to exit. The red dashed curve labelled $m'(a, z_{lo}, \cdot)$ in panel (b) shows the policy function associated with this lower level of demand; the firm's decision to exit is shown as the point $(m_2, 0)$ on this curve.

Now suppose instead that the firm keeps its higher demand shock instead of receiving the bad one. In this case, its optimal market penetration choice, m_3, is shown in panel (a) as the intersection of its current marginal cost curve, $f_{m'}(m_2, \cdot)$, and its marginal benefit, $\tilde{\pi}_{hi} - Q\mathbb{E}[f_m(m_3, m'')|z_{hi}]$. This choice is shown in panel (b) as the point (m_2, m_3) on the original policy function (the solid blue curve). In this case, the firm's policy function under the bad demand shock z_{lo} at m_3 is positive: this level of market penetration is high enough that the firm will no longer choose to exit if it receives the bad shock. This illustrates how the model generates higher exit rates among smaller, newer exporters.

3.7 Aggregation and equilibrium

The final piece of the model is a law of motion that describes how the distribution of exporters evolves over time. Let $\Psi_j(a, z, m)$ denote the joint distribution of productivities, demand shocks, and market penetration rates in market j. This distribution evolves according to the law of motion

$$\Psi_j'(\mathcal{A} \times \mathcal{Z} \times \mathcal{M}) = \int_{\mathbb{R}^2_{++} \times [0,1]} Q_j(a, z, m, \mathcal{A} \times \mathcal{Z} \times \mathcal{M}) \, \mathrm{d}\Psi_j(a, z, m) \tag{22}$$

where $\mathcal{A}$ and $\mathcal{Z}$ denote typical subsets of $\mathbb{R}_{++}$, $\mathcal{M}$ denotes a typical subset of $[0, 1]$, and $Q_{j,}(a, z, m, \mathcal{A} \times \mathcal{Z} \times \mathcal{M})$ is the probability that a firm with productivity a, demand shock z, and customer base m transitions to a state in the set $\mathcal{A} \times \mathcal{Z} \times \mathcal{M}$. This transition function is given by

$$\begin{aligned} Q_j(a, z, m, \mathcal{A} \times \mathcal{Z} \times \mathcal{M}) = \delta(a) \int_{\mathbb{R}^2_{++}} \mathbb{1}_{\left\{m_j'(a,z,n) \in \mathcal{M}\right\}} \, \mathrm{d}G(a', a) \mathrm{d}H(z', z) \\ + (1 - \delta(a)) \int_{\mathbb{R}^2_{++}} \mathbb{1}_{\{0 \in \mathcal{M}\}} \, \mathrm{d}\bar{G}(a') \mathrm{d}\bar{H}(z'). \end{aligned} \tag{23}$$

The first term on the right-hand side is the probability that a firm survives, chooses a new customer base in the set $\mathcal{M}$, draws a productivity shock in the set $\mathcal{A}$, and draws a demand shock in the set $\mathcal{Z}$. The second term is the probability that a firm dies and is replaced by a new firm with no customers, productivity in $\mathcal{A}$, and demand in $\mathcal{Z}$.

A stationary equilibrium consists of: (i) a collection of export cost policy functions $\left\{f_j(m, m')\right\}_{j=1}^J$ that solve the cost minimization problem (17); (ii) a collection of value functions and market penetration policy functions $\left\{V_j(a, z, m), m_j'(a, z, m)\right\}_{j=1}^J$ that solve the firm's dynamic problem (18); and (iii) a collection of distributions $\{\Psi_{j,t}\}_{j=1}^J$ that satisfy the law of motion (22). In my quantitative analysis, I solve for transition dynamics following permanent and temporary changes in trade costs as well as stationary equilibria. A transition equilibrium is an infinite sequence of the objects described above that satisfy the relevant conditions at each point in time.

3.8 Accounting for the facts

The model cannot be characterized in closed form, but it is easy to see how it accounts for the facts described in section 2. The cross-sectional Facts 1.1–1.3 are generated by the same mechanisms as in Arkolakis (2010). Properties (P1) and (P4) described in section 3.5 imply that the entry threshold, $\underline{z}_j(a)$, is higher in harder markets with smaller populations, smaller incomes per capita, or higher trade costs. A higher entry threshold means that fewer firms enter these markets, accounting for Fact 1.1. Those firms that do enter these markets are more productive on average, which means that they are more likely to enter other markets as well, accounting for Fact 1.3. Property (P2) implies that firms with lower productivities and/or demand shocks have fewer customers than other firms in a given market. As these firms only export to easy markets, exports in these markets are more concentrated among the firms with the highest sales, accounting for Fact 1.2.

Facts 2.1–2.5 regarding new-exporter dynamics are explained by combining property (P3) with the other properties. As Figure 3 illustrates, the fact that the marginal market penetration cost $f_{j,m'}(m, m')$ is decreasing in m implies that new entrants (firms with $m = 0$) face higher marginal costs than incumbents and therefore have fewer customers, accounting for Fact 2.1. Using the result that the exit threshold, $\bar{z}_j(a, m)$, is decreasing in m with Fact 2.1 implies that new entrants are more likely to exit than incumbents, accounting for Fact 2.2. Property (P4) implies that the exit threshold is higher in harder markets, accounting for Fact 2.3.[11] Facts 2.4 and 2.5 are established by combining all four properties. As mentioned above, these properties imply that easier markets attract more low-productivity exporters, and that these exporters have fewer customers than high-productivity exporters in these markets. The property that the exit threshold, $\bar{z}_j(a, m)$, is decreasing in both a and m implies that these small exporters exit frequently. Thus, the marginal exporter has fewer customers and exits more frequently compared to incumbents in easy markets than is the case in harder markets.

Property P3, which implies that firms accumulate customers gradually over time, is also crucial for generating Facts 3.1–3.4 regarding the effects of tenure on sales and survival. Fact 3.2, which states that longer tenures are associated with lower exit rates, is explained in the same way as Fact 2.2: the exit threshold is decreasing in m, which means that the likelihood of exit falls as a firm adds to its customer base each period. Fact 3.1, which states that longer spells have greater sales upon entry and a stronger effect of tenure on sales, is the product of several factors. In the model, longer spells are driven by higher demand shocks at entry: higher demand at the beginning of a spell means it takes longer on average for demand to fall below the exit threshold. Property (P2)

[11] Strictly speaking, establishing this is complicated by the fact that while firms that export to harder markets have fewer customers in these markets than they do in easy ones due to property (P2), these firms are also more productive on average. Controlling for firm fixed effects, as I do in section 2.2, sidesteps this issue. Note also that this does not require property (P3).

implies that higher demand on entry leads to more customers, accounting for the first part of Fact 3.1. Property P3 implies that more customers on entry means more customers in the next period, and the period after that, and so on. Note that Fact 3.2, which implies higher returns in the future from gaining more customers today (the second term on the right-hand side of (19)), reinforces this effect. This accounts for the second part of Fact 3.1. Fact 3.4, which states that exit is less likely in easy markets after conditioning for tenure, is explained by property P4. Fact 3.3, which states that Fact 3.1 is more pronounced in easy markets than in hard ones, is also generated by property P4. Firms penetrate easy markets more deeply, which allows them to take greater advantage of property P3 over time. Additionally, regarding Fact 3.4, the gains from lower market penetration costs in the future are larger in easy markets.

3.9 Relationship to other models

The model generalizes several existing export participation frameworks that are commonly used in quantitative studies.

Static market penetration model. If retaining old customers is impossible ($\psi_o = 0$), the overall market penetration cost is equal to the cost of attracting new customers, which means that the cost of exporting does not depend on a firm's current customer base: $f_j(m') \equiv f_j(0, m') = s_j(0, m')$. In this case, the cost of exporting is the same as in the static model of Arkolakis (2010). While the firm's decision about whether to export and how many customers to serve is now static, the idiosyncratic dynamics of productivity and demand still generate variation in sales and survival over firms' life cycles. Taking these persistent shocks into account, this specification of the model in its entirety is equivalent to Arkolakis (2016), which integrates Arkolakis (2010) with a theory of firm dynamics driven by exogenous productivity shocks. I refer to this model as the *static MP model*. The first-order condition that describes the firm's optimal market penetration rate in the static MP model is

$$f_j'\left(m'\right) \geq \tilde{\pi}_j(az)^{\theta-1}, \tag{24}$$

where $f_j'(m')$ is the marginal market penetration cost. This implies that the exit threshold depends only on the firm's productivity and is equal to the entry threshold: $\bar{z}_j(a) = \underline{z}_j(a)$.

Sunk cost model. If $\gamma_n = \gamma_o = 0$, the marginal attraction and retention costs are both constant. Because the marginal benefit of serving additional customers is also constant (equation (11) is linear in m), firms choose either to serve all customers in a market or none. In this case, the cost of exporting depends only on whether a firm has any current customers to retain; that is, whether a firm is currently an exporter. This specification is equivalent to the *sunk cost model* of Das et al. (2007); the constant marginal attraction and retention costs, which are equal to $f_{j,0} \equiv L_j^{\alpha_n}/\psi_n$ and $f_{j,1} \equiv L_j^{\alpha_o}/\psi_o$, can be interpreted as the sunk entry cost and fixed continuation cost, respectively.

The firm's problem now boils down to a discrete choice about whether to export:

$$V_j(a,z,m) = \max_{m' \in \{0,1\}} \left\{ \pi_j(a,z,m') - f_{j,m} + \frac{\delta(a)}{1+R} \mathbb{E}\left[V_j(a',z',m')|a,z\right]. \right\} \tag{25}$$

The sunk cost model is characterized by entry and exit thresholds that depend only on the firm's productivity, but may differ if the entry and continuation costs are different. Specifically, if $f_j(0) > f_j(1)$, then the entry threshold will be larger than the exit threshold: $\underline{z}_j(a) > \bar{z}_j(a)$. This is typically the case in calibrated sunk-cost models (see, e.g., Das et al., 2007; Ruhl and Willis, 2017). Note that if the marginal attraction cost is constant ($\gamma_n = 0$) and retaining old customers is impossible ($\psi_o = 0$), the market penetration cost can be interpreted as a fixed per-period cost of exporting and the model collapses to the static model of Chaney (2008).

Models with exogenous new-exporter dynamics. My theory is also similar to, but does not generalize, models of exporter dynamics driven by exogenous shocks. Ruhl and Willis (2017) extend the sunk cost model to allow export capacity to shift upward deterministically over time in order to capture the fact that new exporters sell less than incumbents. Alessandria et al. (2021b) build on Ruhl and Willis (2017) by making these shifts stochastic, which allows their model to capture the fact that new exporters are also less likely to survive. The firm's problem in Alessandria et al. (2021b), which I refer to as the *exogenous NED model*, takes the same form as in the sunk cost model, but with the addition of a new exogenous state variable that represents an idiosyncratic shock to the technological cost of exporting (i.e., an iceberg cost). The key difference between my model and Alessandria et al. (2021b) is that in the latter, export capacity is driven purely by luck and is therefore orthogonal to productivity and demand. In my model, on the other hand, export capacity is endogenously correlated with exogenous states because firms with good states choose to accumulate more customers.

Models with endogenous exporter dynamics. There are several other papers that use models of endogenous customer accumulation to account for new-exporter dynamics, most notably those developed by Fitzgerald et al. (2023) and Piveteau (2020). There are two key differences between these models and mine. First, they assume that all entrants start exogenously with the same number of customers in all destinations. My theory explains why new entrants have fewer customers than incumbents and generates dispersion in entrants' sales both within and across destinations, consistently with the data. Second, these models require sunk and fixed costs that vary exogenously across firms and destinations in addition to customer accumulation costs in order to generate realistic entry and exit patterns. In my model, extensive-margin dynamics are driven solely by the marginal cost of serving the first customer, $f_j(\cdot,0)$, which varies endogenously across firms, across destinations, and over time. This suggests that while the models of Fitzgerald et al. (2023) and Piveteau (2020) are successful in accounting for exporter dynamics in historical data (if anything, more successful than my model as they have more free parameters to play with), my model may be more suitable for counterfactual analysis.

4 Calibration

I calibrate the model's parameters using indirect inference. The targeted statistics are the facts about market-level variation in exporter performance documented in section 2.1. The calibrated model succeeds in matching these targets, but also in accounting for the non-targeted facts about firm-level variation documented in section 2.2. After calibrating and validating the model, I compare its performance to that of three other models discussed in section 3.9: the static MP model; the sunk cost model; and the exogenous NED model. I then explore how the costs of exporting chosen by firms in equilibrium in the baseline model vary with time in a market and across destinations, and discuss how these findings relate to other findings in the literature.

4.1 Procedure

The first step in my calibration procedure is to choose a set of destinations and assign values to their characteristics, L_j, Y_j, and τ_j. I use the same 63 destinations in the Brazilian microdata analyzed in section 2; as before, their characteristics are taken from the CEPII Gravity database. The second step is to calibrate the parameters that govern the distribution of firms' exogenous types and the cost of exporting. I assume that demand follows a standard first-order autoregressive process in logs with persistence ρ_z and innovation dispersion σ_z^2. I assume that productivity is unconditionally distributed log-normally with variance σ_a^2, and that each period firms retain their productivities with probability ρ_a and draw new ones with probability $1-\rho_a$.[12]. Following Alessandria et al. (2021b), I parameterize the death rate as $1-\delta(a) = \max(0, \min(e^{-\delta_0 a} + \delta_1, 1))$.

With these parameterizations, there are 16 parameters that must be calibrated: ρ_a, σ_a, ρ_z, and σ_z govern the distributions of the exogenous state variables; δ_0 and δ_1 govern survival; α_n, β_n, γ_n, and ψ_n govern the cost of attracting new customers; α_o, β_o, γ_o, and ψ_o govern the cost of retaining old customers; θ governs the elasticity of substitution between varieties; and R governs the rate at which firms discount future profits. Following Ruhl and Willis (2017) and Alessandria et al. (2021b), I set θ externally to 5, a common value in the literature that implies a trade elasticity of 4 in the absence of firm-level responses. I set the discount rate R externally to match the average Brazilian real interest rate of 10% during 2000-2005.[13]

This leaves me with 14 parameters whose values must be jointly determined. I use an indirect inference strategy to find values of these parameters that minimize the distance between statistics

[12]This approach is similar to a standard log-normal AR(1) process, but is more computationally tractable because an exporter's continuation value conditional on drawing a new productivity is independent of its current productivity. It is commonly used in models with Pareto-distributed productivities (see, e.g., Buera et al., 2011)

[13]Brazilian real interest rates were high and volatile during the 1980s and 1990s, and then declined after the Brazilian currency was allowed to float in 1999. The 10% figure I use is almost identical to the value used by Ruhl and Willis (2017) for Colombia.

computed using the Brazilian microdata and the same statistics computed using simulated data generated by the model. Specifically, for each of the six measures of exporter performance discussed in section 2.1, I target the cross-destination average shown in panel (a) of Table 1 and the coefficients on population, income per capita, and trade barriers shown in panel (b). Additionally, I target a multilateral export participation rate of 26%.[14] For each candidate parameter vector, I use the model to simulate a panel of firms and calculate these statistics by applying the same processing and analysis that I apply to the real data. I then search over the parameter space to find the vector of parameters that minimizes the mean squared difference between the simulated and actual moments, where each statistic is weighted by the inverse of its standard error. To ensure that the estimated parameter vector is a global minimum, I partition the parameter space into sets of increasingly small subspaces, use a stochastic population-based global optimization method in each subspace, and "polish off" each subspace's best candidate parameter vector using a simplex-based method. Essentially, my approach follows the Subplex method (Rowan, 1990) but adds a stochastic search in each subspace. It is similar to the TikTak algorithm described in Arnoud et al. (2019).

There are a total of 25 target statistics (the average and three slope coefficients for each of the six statistics from section 2.1 plus the overall export participation rate). The 14 estimated parameters are therefore over-identified, but several of the target moments are correlated. Looking at the average number of other destinations served, for example, the cross-destination mean is negatively correlated with the three slope coefficients because exporters in the most popular markets serve only a few other destinations; increasing the magnitude of the slope coefficients for this statistic also raises the overall mean. Consequently, an exactly identified estimation strategy would be problematic, whereas my strategy ensures there is sufficient independent variation in the target statistics to pin down all the parameters.

Each of the target statistics affects some parameters more than others. The cross-destination averages of the top-five share and the number of other destinations served pin down the variances of the productivity distribution and demand shock, σ_a and σ_z. The overall export participation rate determines the level of the new-customer attraction cost, ψ_n, while the average exit rate influences the level of the old-customer retention cost, ψ_o, and the minimum death probability, δ_1. The average exit rate of entrants relative to that of incumbents affects the persistence of the demand shock, ρ_z, and the sensitivity of the exit rate to productivity, δ_0. The slope coefficients of the exit rate, number of other destinations served, and relative exit rate of entrants play dominant roles in identifying the returns to market size in attracting new customers and retaining old ones, α_n and α_o. Finally, the average and slope coefficient of the relative entrant size and the slope coefficient of the top-five share jointly pin down the convexity parameters of the attraction and retention costs, γ_n and γ_o.

[14] I do not have data on non-exporting Brazilian firms, so I rely on the estimate of Ruhl and Willis (2017) for Colombia.

4.2 Parameter values

Table 2 lists the parameter estimates resulting from the procedure described above.[15] Panel (a) shows the parameters that govern the distribution and evolution of firms' exogenous characteristics. The variance of multilateral productivity shocks is somewhat lower than that estimated in models of multilateral exporter dynamics (see, e.g., Ruhl and Willis, 2017; Alessandria and Choi, 2014; Alessandria et al., 2021b), but the persistence of these shocks is higher. However, demand shocks are less persistent than productivity shocks, and the product of productivity and demand, $a \times z_j$, exhibits similar dispersion and persistence to other studies' productivity processes. The variance of the productivity distribution over and above the variance in the demand shock process is needed to capture the high concentration of exports among the largest exporters in the average destination and the fact that most firms have relatively small portfolios of export destinations, while the lower persistence of demand helps account for variation in exit rates across destinations. The survival function parameters, shown in panel (b), are similar to the estimates of Alessandria et al. (2021b), who use business dynamics data in which firm creation and death can be directly observed.

Table 2: Calibrated parameter values

Parameter	Meaning	Value
(a) Distribution of firm types		
σ_x	Variance of productivity	1.02
ρ_x	Persistence of productivity	0.98
σ_z	Variance of demand	0.44
ρ_z	Persistence of demand	0.60
δ_0	Correlation of survival with productivity	34.7
δ_1	Minimum death probability	0.03
(c) New customer attraction costs		
α_n	Macro return to market size	0.51
β_n	Micro return to market size	0.94
γ_n	Convexity	6.44
ψ_n	Level	0.10
(d) Old customer retention costs		
α_o	Macro return to market size	0.96
β_o	Micro return to market size	0.79
γ_o	Convexity	3.82
ψ_o	Level	0.06

[15] Reporting standard errors is not feasible in part due to the computational burden of solving and simulating the model repeatedly, but also due to the fact that the objective function exhibits discrete jumps and other non-differentiabilities.

The parameters of the advertising cost functions are shown in panels (c) and (d). The macroeconomic return to market size is significantly larger in attracting new customers than in retaining old ones: $\alpha_n < \alpha_o$. This captures the hypothesis laid out in section 3.3 that advertising to current customers is more akin to contacting them one by one, whereas advertising to new customers is more like advertising en masse through media channels. Conversely, the microeconomic return to market size is larger for advertising to old customers: $\beta_o < \beta_n$. The new-customer attraction cost function is more convex than the old-customer retention cost function: $\gamma_n > \gamma_o$. This indicates that it is harder to attract large blocks of new customers than to retain large blocks of old ones. Finally, the level parameter is higher for customer attraction than for retention: $\psi_n > \psi_o$. Taken at face value, this would seem to suggest that acquiring new customers is cheaper than retaining old ones, which would contrast sharply with the large startup costs in standard sunk-cost models like Das et al. (2007) and Alessandria and Choi (2007) required to generate realistic export participation and turnover. However, as I shall show in section 4.5, we see that this is actually not the case when we examine the exporting costs that firms choose to incur in equilibrium more closely.

4.3 Performance on targeted and non-targeted statistics

The calibrated model closely replicates the targeted statistics. The first row in panel (a) of Table 3 shows the means of the six measures of exporter performance discussed in section 2.1 in the simulated data. The mean top-five share, average number of other destinations served, exit rate, relative entrant size, and relative entrant exit rate are all close to the empirical means reported in Table 1. The mean number of exporters is two-thirds higher in the model than in the data, but this moment is weighted less than the other measures' means due to its relatively high standard error. Panel (b) of Table 3 shows these measures' associations with market characteristics in the model. All coefficients but one (the effect of population size on the exit rate) have the correct sign, and of these all but two (the effect of GDP per capita on the overall exit rate and the effect of population on the relative exit rate of entrants) have the correct magnitude.

The model also reproduces, at least qualitatively, the other facts documented in section 2 that were not targeted in the calibration. Panels (a)–(c) of Figure 4 show how sales vary with tenure and spell length in the model. In each figure, lines show estimates of specification (2) or (3) using simulated data from the model, and shaded areas show the confidence intervals for the actual empirical estimates. Just as in the data, the most successful exporters exhibit the strongest growth in sales over the course of their export spells, and the differences in sales growth between more- and less-successful exporters are more pronounced in easy destinations than in hard ones. The model estimates for all markets, shown in panel (a), and for easy markets, shown in panel (c), are very close to their empirical counterparts. However, the model estimates for hard markets are lower than the empirical estimates—the model actually overstates the differences in sales dynamics between hard and easy markets.

Table 3: Exporter performance across markets in baseline and alternative models

Statistic/coefficient	Num. exporters	Top-five share	Avg. num. dests.	Exit rate	Entrant rel. size	Entrant rel. exit rate
(a) Cross-market averages						
Baseline model	980	0.58	18.82	0.44	0.37	0.08
Static mkt. pen. model	811	0.57	14.80	0.55	0.33	0.11
Sunk-cost model	677	0.40	16.46	0.37	1.19	0.01
Exog. exporter dyn. model	520	0.50	13.41	0.38	0.35	0.27
(b) Associations with destination characteristics: Baseline model						
log GDPpc	0.761	0.071	−2.149	−0.074	−0.135	0.036
log population	0.234	0.015	−1.207	0.024	−0.099	0.038
log trade barrier	−0.744	−0.069	2.217	0.072	0.128	−0.039
(c) Associations with destination characteristics: Static mkt. pen. model						
log GDPpc	0.748	0.063	−1.685	−0.071	−0.099	0.028
log population	0.357	0.031	−0.836	−0.034	−0.047	0.013
log trade barrier	−0.737	−0.062	1.684	0.071	0.100	−0.028
(d) Associations with destination characteristics: Sunk-cost model						
log GDPpc	0.843	0.080	−2.051	−0.075	−0.302	0.033
log population	0.120	0.003	−1.174	0.070	−0.525	0.036
log trade barrier	−0.811	−0.078	2.152	0.071	0.341	−0.033
(e) Associations with destination characteristics: Exog. exporter dyn. model						
log GDPpc	0.781	0.086	−1.362	−0.124	−0.072	0.082
log population	0.123	0.014	−0.814	0.053	−0.132	0.048
log trade barrier	−0.747	−0.084	1.607	0.112	0.082	−0.070

Notes: Panel (a) shows average values across destinations of exporter performance measures in each model. Panels (b)–(e) show the coefficients β, γ, and δ from regression (2) estimated using simulated data from each model.

Panels (a)–(b) of Figure 5 show how the likelihood of exit depends on tenure in the model. As in the data, the exit rate falls with tenure, and exit rates conditional on tenure are higher in harder markets than easy ones. However, the exit rate does not decline as much with tenure in the model as in the data. This is consistent with Fitzgerald et al. (2023), who find that customer accumulation plays a key role in explaining the post-entry dynamics of export quantities, whereas slow learning about idiosyncratic demand shocks is important to matching post-entry survival dynamics; my model features the former but not the latter.

4.4 Comparison with other models' performance

As discussed in section 3.9 above, the theory of export costs developed in this paper generalizes several existing models and provides an explanation for key assumptions in others. Here, I analyze the extent to which these other models can account for the facts documented in section 2 and use these results to shed light on which of my theory's ingredients are most important in accounting for these facts. I analyze three alternative models: the static MP model of Arkolakis (2016), in which the market penetration cost does not depend on a firm's current customer base; the sunk

Figure 4: Effects of tenure and spell length on sales in baseline model vs. other models

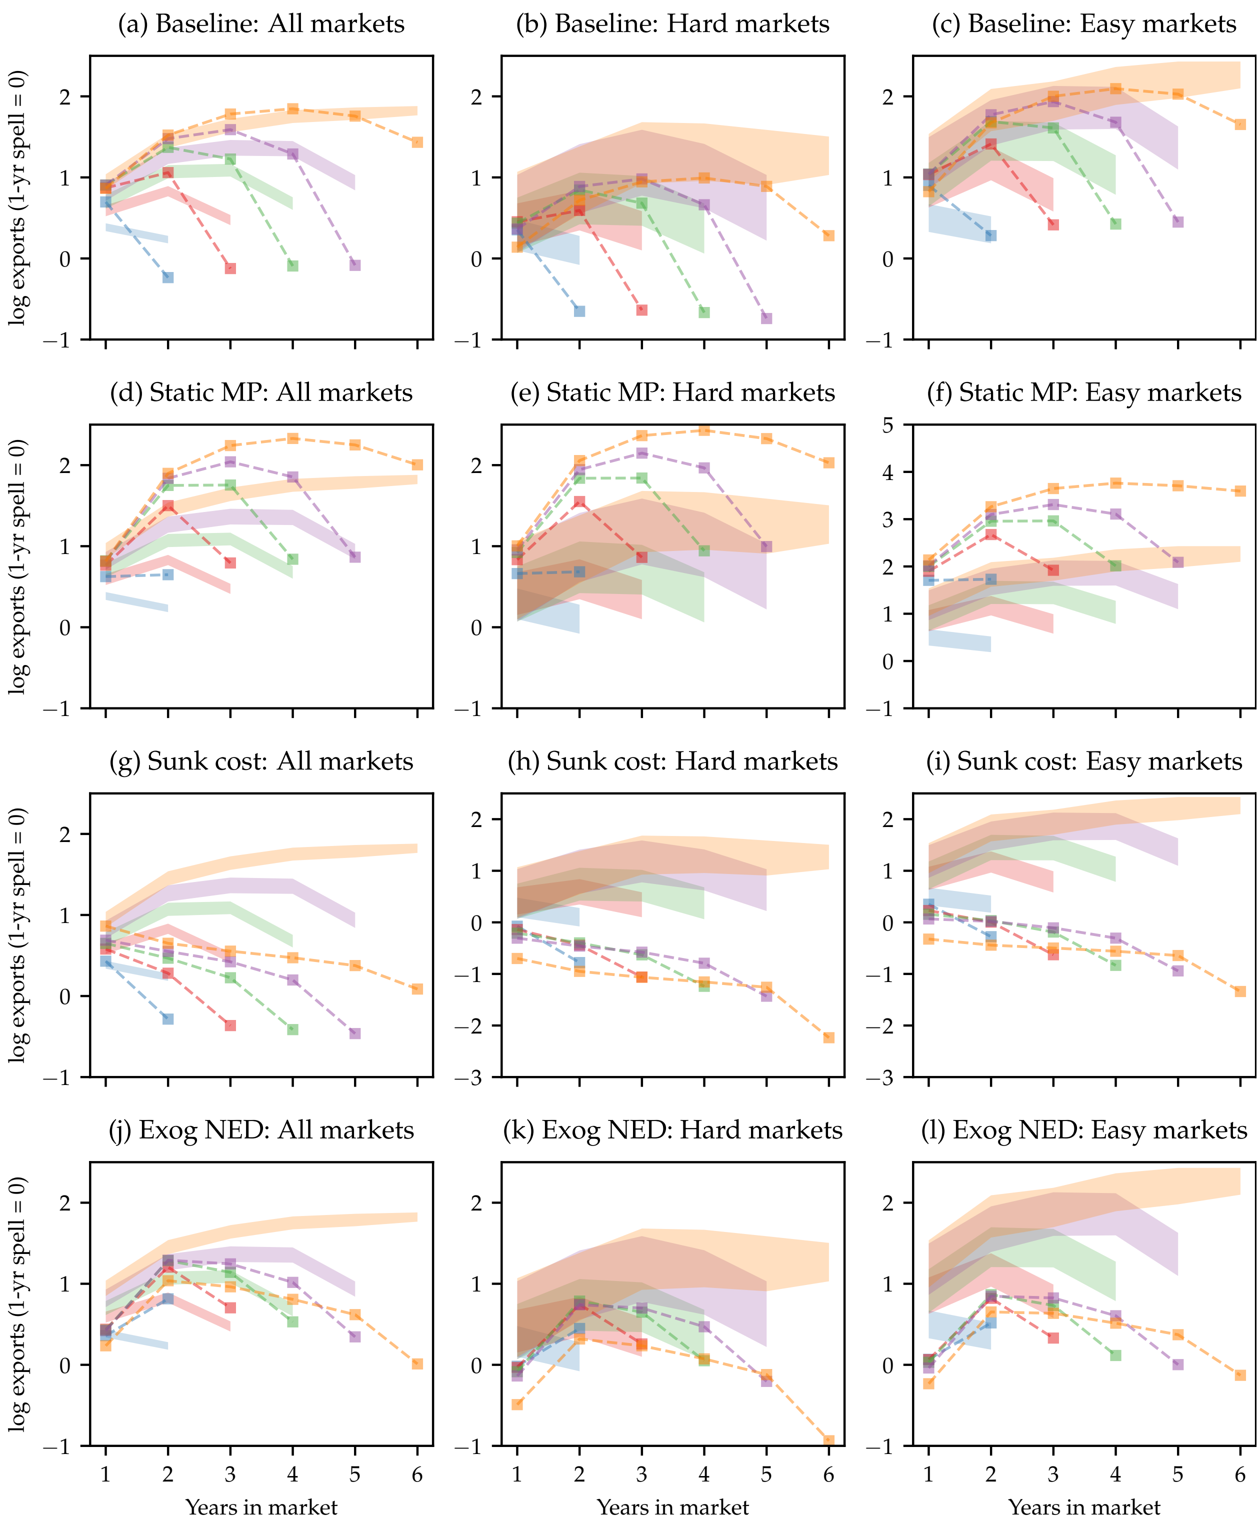

Notes: Each row shows simulated results from a different model specification. First column shows estimates of $\beta_{m,n}$ from (2). Second column shows estimates of $\beta_{m,n,g}$ from (3) for markets in the bottom 50% of export participation, and third column shows estimates for markets in the top 10%. Each line shows $\beta_{m,1}, \beta_{m,2}, \ldots, \beta_{m,6}$ (or $\beta_{m,1,g}, \beta_{m,2,g}, \ldots$ in the second two panels) for a set value of m. In both columns, shaded areas show 95-percent confidence intervals for estimates using actual data discussed in section 2.2.

Figure 5: Exit rates conditional on tenure in baseline model vs. other models

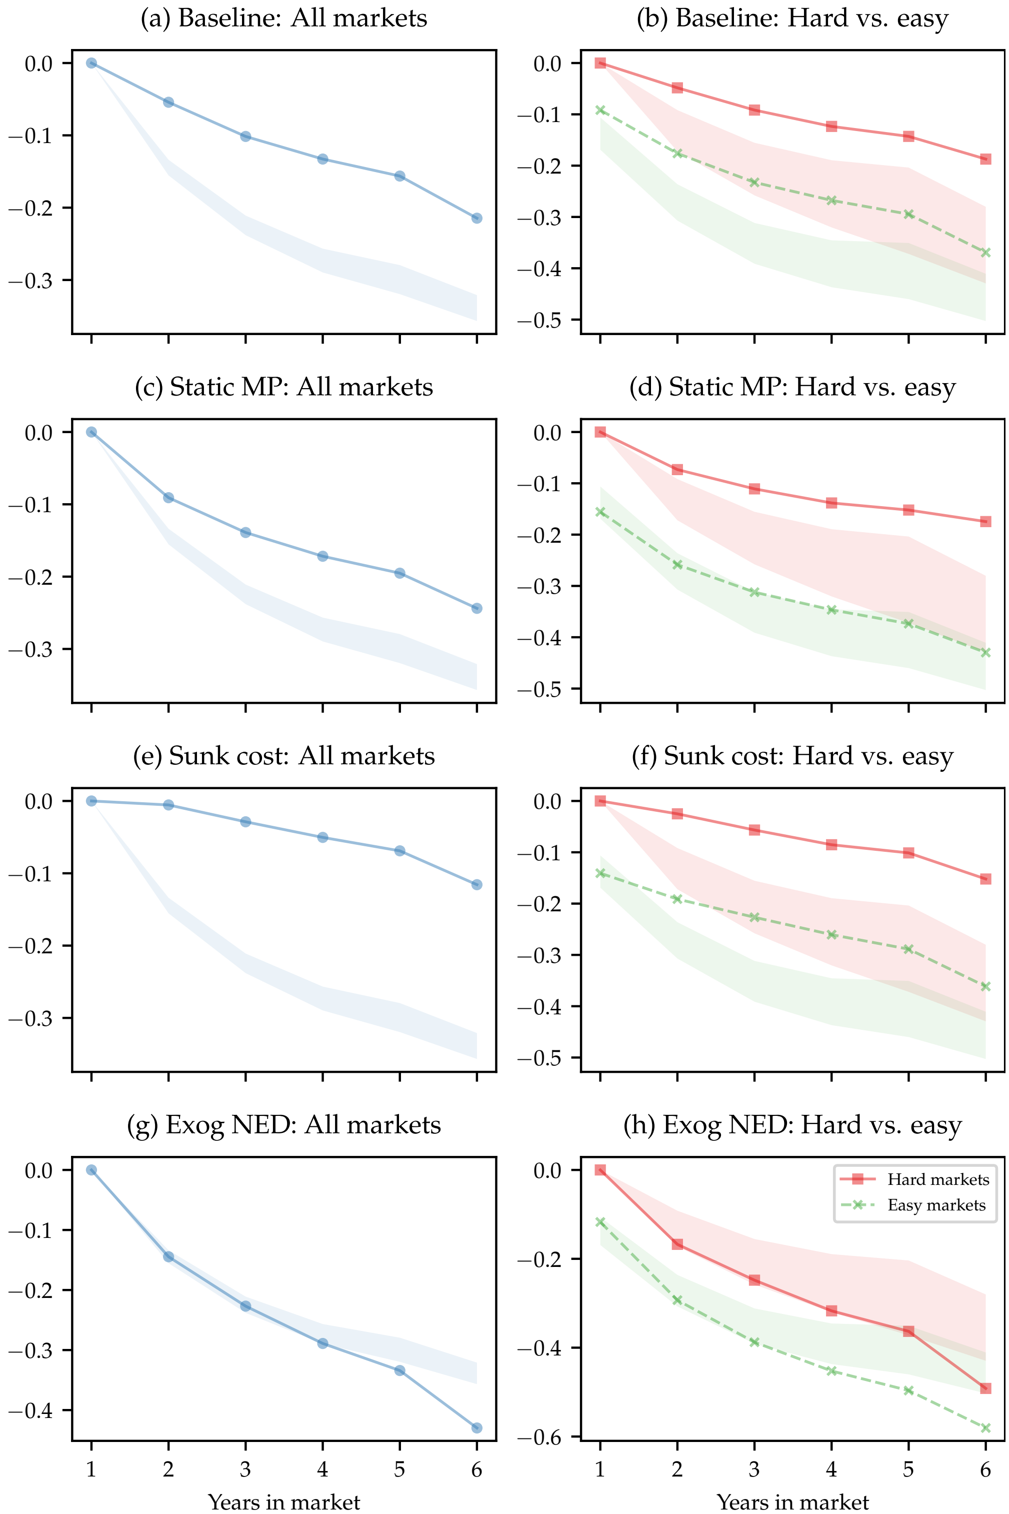

Notes: Each row shows simulated results from a different model specification. First column shows estimates of β_n from (4). Second column shows estimates of $\beta_{n,g}$ from (5). In second column, solid red lines with square markers show estimates for markets in bottom 50% of export participation, and dashed green lines with 'x' markers show estimates for destinations in top 10%. In both columns, shaded areas show 95-percent confidence intervals for estimates using actual data discussed in section 2.2.

cost model of Das et al. (2007), which features a dynamic export participation decision but no market penetration decision; and the exogenous NED model of Alessandria et al. (2021b), which adds exogenous shocks to export capacity to the sunk cost model. In each alternative model, I hold fixed the stochastic processes for multilateral productivity, a, bilateral demand, z, and survival, $\delta(a)$, at their calibrated values listed in Table 2. This allows me to weigh the relative contributions of endogenous market penetration dynamics and exogenous shocks in accounting for the facts described in section 2 above.

Static MP model. In my static MP model, I set $\psi_o = 0$ so that retaining old customers is impossible, and hold fixed the relevant customer attraction parameters (the macroeconomic return to scale, α_n, and the degree of convexity, γ_n; the microeconomic return to scale, β_n, is irrelevant).[16] The results of the calibration exercise in this model are shown in the second row of panel (a) and panel (c) of Table 3; panels (d)–(f) of Figure 4; and panels (c)–(d) of Figure 5.

The static MP model is successful in generating a high level of cross-sectional sales concentration and new exporters that sell substantially less than incumbents, but it generates too much turnover on average and too little turnover among new entrants. This indicates that endogenous market penetration is an important driver of export participation dynamics. The reason is that incumbent exporters face the same export costs as entrants in this model, whereas in the baseline model, export costs fall as firms accumulate customers over time. Consequently, the exit threshold in the static MP model—which is the same as the entry threshold in that model—is a function solely of exogenous characteristics, whereas in the baseline model it also depends on a firm's customer base. However, this version of the model is about as successful as the baseline in capturing the associations between destination characteristics and exporter performance. This indicates that exogenous shocks to productivity, demand, and survival play important roles in generating these correlations.

This version of the model also captures qualitatively the differences in sales trajectories over long vs. short export spells, but sales grow faster over long spells in particular than in the data or the baseline model, especially for easy markets. This indicates that market penetration dynamics play an important role in quantitatively reproducing these patterns. The static market penetration model performs similarly to the baseline model in accounting for the effects of tenure on exit rates. As discussed above, this is consistent with Fitzgerald et al. (2023), who find that customer accumulation is a more important driver of sales growth than of survival dynamics.

[16] An alternative approach is to recalibrate α_n and γ_n, which are identified most strongly by the average number of other destinations served and the average share of exports accounted for by the top 5 percent of exporters. Choosing new values of these parameters to match these two moments does not materially alter the other results.

Sunk cost model. In my sunk cost model, I set $\gamma_o = \gamma_1 = 0$ so that the marginal attraction and retention costs are constant, and recalibrate the efficiency parameters ψ_n and ψ_o to match the overall multilateral export participation rate of 26% and the average bilateral exit rate of 40% observed in the data.[17] The results are shown in the third row of panel (a) and panel (d) of Table 3; in panels (g)–(i) of Figure 4; and in panels (e)–(f) of Figure 5.

The sunk cost model generates too little concentration of exports among top exporters. Without the market penetration margin, low-productivity/low-demand exporters have the same number of customers as high-productivity/high-demand exporters, and so the former sell too much in this model relative to the latter. This version of the model also fails to generate new-exporter dynamics: entrants are too large and too likely to survive compared to incumbents. This is consistent with the findings of Ruhl and Willis (2017), who show that customer accumulation and other sources of firm-level intensive margin growth are crucial to capturing these dynamics. However, like the static market penetration model, this model succeeds in capturing the associations between destination characteristics and exporter performance, confirming that exogenous shocks are important to accounting for these patterns.

This version of the model also cannot account for post-entry sales trajectories. In fact, in this model, sales tend to fall after a firm enters a new export destination. This is driven by regression to the mean in productivity and demand. Firms begin exporting after they receive good shocks and do not exit until they receive sufficiently bad shocks, so new exporters tend to be more productive and have greater demand for their products than incumbents. In the absence of the customer accumulation margin, sales tend to fall post-entry as productivity and demand converge to the mean. The sunk cost model also fares worse than the baseline in capturing the effect of tenure on survival, although it fares well in generating a difference in survival trajectories between hard and easy markets.

Exogenous NED model. In my exogenous NED model, I assume that a firm enters a new market with m_0 customers and each period it continues to export there is a chance ρ_m its customer base will grow to $m_1 > m_0$. Once it has reached this greater level of market penetration, its customer base can fall back to m_0 with the same probability. I also normalize $m_1 = 1$ and choose m_o so that this version of the model also matches the observed ratio of the average entrant's sales to the average incumbent's. As in the sunk cost model, I recalibrate ψ_o and ψ_n to match the multilateral export participation rate and average bilateral exit rate. The results are shown in the last row of panel (a) and panel (e) of Table 3; in panels (j)–(l) of Figure 4; and in panels (g)–(h) of Figure 5.

[17] Firms' incentives to enter and exit are different in this model due to the absence of the market penetration margin, so leaving ψ_n and ψ_o unchanged would lead to a different export participation rate than the baseline model. This approach allows me to analyze the role of the market penetration margin in accounting for the facts at hand while holding fixed the measure of exporting firms. Overall, the results are similar when leaving ψ_n and ψ_o fixed at their baseline values, except that the export participation rate is higher and turnover is less frequent.

The exogenous NED model does better than the sunk cost model—but not as well as the baseline and static MP models—in generating concentration of sales among top exporters. It also succeeds in capturing the propensity of entrants to exit more often than incumbents. This is consistent with the findings of Alessandria et al. (2021b), who show that gradual growth in export capacity helps account for entrants' low likelihood of survival. Like the previous two alternative models, this model, too, succeeds in capturing the associations between destination characteristics and exporter performance. Once again, this confirms the importance of exogenous shocks in explaining these patterns.

Despite this model's ability to capture many of the facts about exporter performance across markets, it cannot account for post-entry sales trajectories. Sales grow for one period after entry, but then they decline for the remainder of an export spell, as in the sunk cost model. This indicates that endogenous customer accumulation dynamics—which allow different firms to choose different market penetration paths in a given destination, and the same firm to choose different paths in different destinations—are crucial to accounting for these trajectories. However, the exogenous new NED model fares better than the baseline model in accounting for the effect of tenure on exit. This suggests that exogenous shocks to export capacity may be an important driver of this pattern.

4.5 Equilibrium exporting costs

To illustrate how market penetration costs in the baseline model vary across markets and across firms within each market, I compute four variables at the market level in the simulated data from the calibrated model: average cost in levels, normalized to one in the average market; average cost for entrants relative to incumbents; the average cost/profit ratio; and average cost/profits for entrants relative to incumbents. Panel (a) of Table 4 reports summary statistics for these variables, and panel (b) reports associations with destination characteristics. In levels, export costs vary by three orders of magnitude across markets; the lowest value is barely 1% of the average and the highest value is more than ten times greater. There is also a fair amount of variation in the ratio of entrants' export costs to incumbents' costs, which ranges from 0.17 to 1.7. When measured relative to profits, however, there is much less variation, both in the overall average cost and the cost paid by entrants relative to incumbents.

To illustrate how exporting costs vary within firms over different export spells, I estimate how spell length and tenure affect exporting costs following the Fitzgerald et al. (2023) approach described in section 2.2. Here, though, I exclude the last year of a firm's export spell (i.e., when $m = n$) because firms that choose to exit endogenously pay zero export costs. Panels (a)–(c) of Figure 6 report the estimated effects on the level of exporting costs. In levels, exporting costs increase over the course of an export spell, and are higher in easy destinations than in hard ones. Panels (d)–(f), however, which report the estimated effects on the ratio of exporting costs to profits, tell a different story. Measured relative to profits, exporting costs are highest at the beginning of an

Table 4: Calibrated export costs across markets

Statistic/coefficient	Avg. cost	Entrant rel. cost	Avg. cost/profits	Entrant rel. cost/profits
(a) Summary statistics				
Mean	1.000	0.538	0.830	1.080
Min	0.018	0.170	0.560	0.778
Max	12.148	1.657	1.298	1.245
Std. dev.	1.767	0.310	0.167	0.109
(b) Associations with market characteristics				
log GDPpc	0.478	−0.233	0.023	0.006
log population	0.836	−0.234	−0.087	0.051
log trade barrier	−0.346	0.241	−0.030	0.001

Notes: Panel (a) shows averages across destinations. Panel (b) shows coefficients β, γ, and δ from regression (2) estimated using simulated data from baseline model. Variable definitions are as follows. Avg. cost: average market penetration cost across all firms, normalized by average across all markets. Entrant rel. cost: Average market penetration cost of new entrants divided by average cost of incumbents. Avg. cost/profits: Average ratio of market penetration costs to profits for all firms. Entrant rel. cost/profits: Average ratio of market penetration costs to profits for entrants, divided by same ratio for incumbents.

export spell, and decline more sharply in easy destinations than in hard ones.

Broadly, these results are consistent with models of exogenous new-exporter dynamics like Ruhl and Willis (2017) and Alessandria et al. (2021b), in which startup costs are similar to continuation costs, especially when measured relative to profits, and with richer models of exporter dynamics like Piveteau (2020) and Fitzgerald et al. (2023), which require exporting costs that vary exogenously across firms and markets to fit the data. The key difference is that variation in exporting costs in my model is endogenous, providing a theoretical foundation for these studies' assumptions. Moreover, exporting costs in my model vary within firms, across markets, and even within markets across export spells.

5 Aggregate implications

A common theme in the trade dynamics literature is that micro matters for macro: firm-level responses drive the dynamics of aggregate trade flows in response to trade reforms and other aggregate shocks. Here, I study the aggregate implications of the theory developed in this paper by analyzing how trade responds in the short and long run to a permanent reduction in trade costs, both in the baseline model and in the alternative models discussed above in section 4.4.

For each market in the Brazilian customs data, I solve for the transition dynamics of aggregate bilateral trade flows following a permanent, unanticipated 10% reduction in iceberg trade costs. I then break the resulting time series into two groups, as in the analyses in the previous sections: markets in the top 10% of export participation (easy destinations) and destinations in the bottom

Figure 6: Effects of tenure and duration on exporting costs

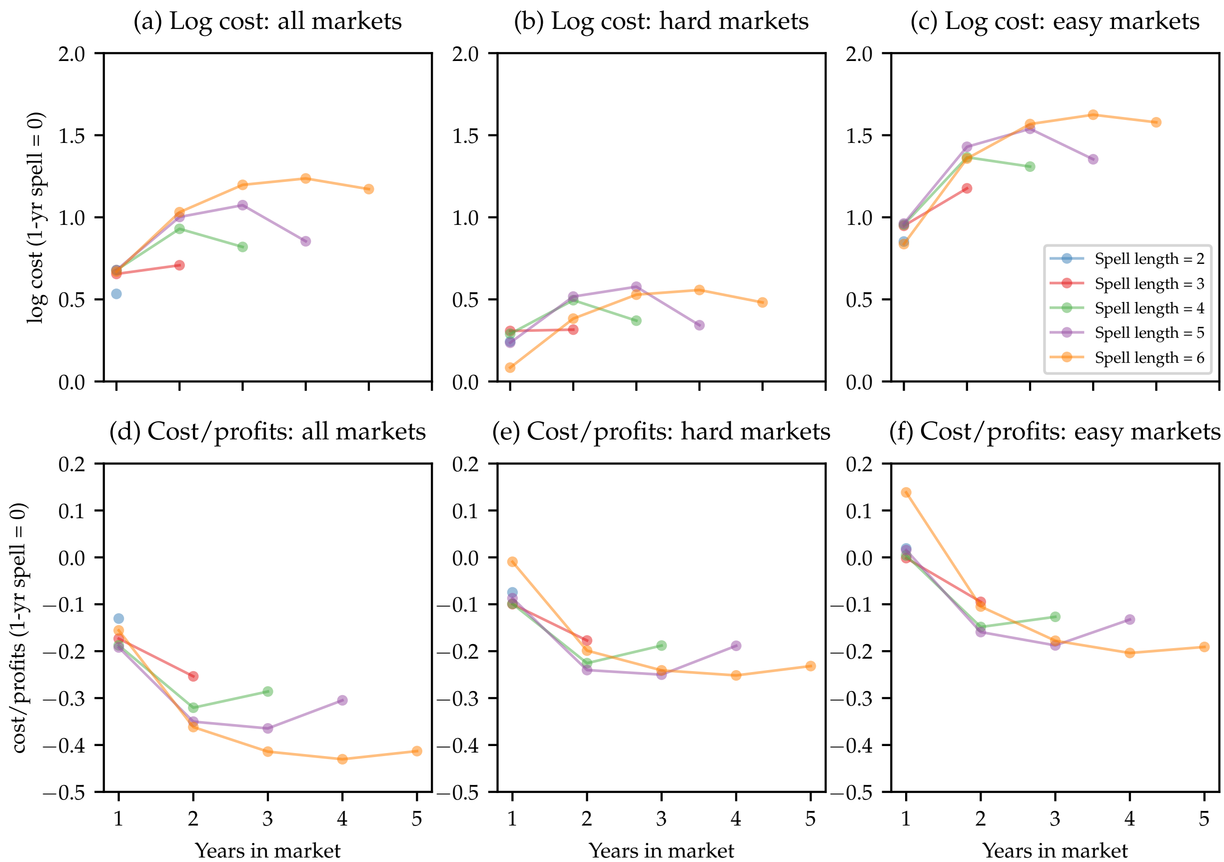

Notes: Figure shows estimates of $\beta_{m,n}$ from versions of (2) with measures of exporting costs instead of export volumes on the left-hand side. In the upper panels, the dependent variable is the log export cost: $\log f_j(m_{i,j,t}, m_{i,j,t+1})$. In the lower panels, it is the export cost normalized by profits: $f_j(m_{i,j,t}, m_{i,j,t+1})/\pi_j(a_{i,j,t}, z_{i,j,t}, m_{i,j,t})$. The center (right) panels show estimates for destinations in the bottom 50% (top 10%) of export participation. Each line shows $\beta_{m,1}, \beta_{m,2}, \ldots, \beta_{m,6}$ for a set value of m.

50% (hard destinations). I repeat this process for the baseline model, the static MP model, the sunk cost model, and the exogenous NED model. I adopt the following timing to make the different forces at work as transparent as possible. In period 0, the model is in its initial steady state. In period 1, trade costs fall after firms have made their market penetration decisions, so trade rises only because of the price elasticity of demand. Thus, the period-1 trade elasticity in all destinations is $\theta - 1 = 4$, the elasticity that would obtain in a model without any firm-level adjustments at all. In period 2, firms begin to adjust their market penetration rates, entering and expanding due to the increase in demand, and the trade elasticity begins to rise. Figure 7 reports the results, with the trade elasticity shown in panel (a), the export participation rate in panel (b), and the average market penetration rate in panel (c).

Figure 7: Transition dynamics after permanent trade reform in baseline model vs. other models

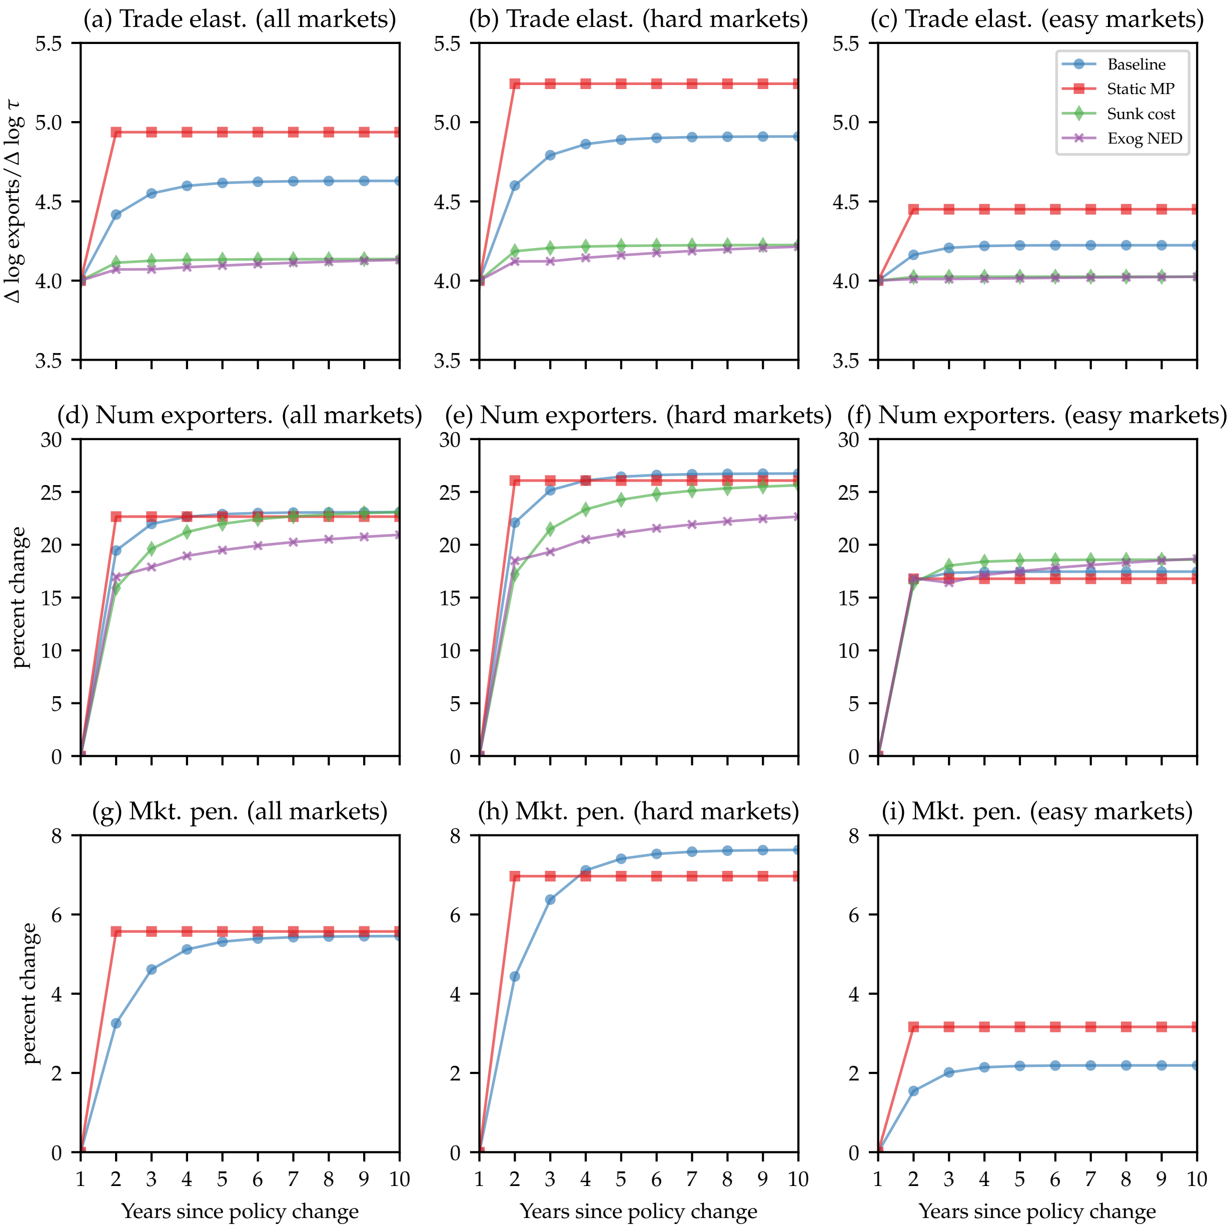

Notes: Figure shows transition dynamics following a permanent, unanticipated 10% drop in trade costs. Top row shows trade elasticity, measured as the cumulative log change in aggregate trade divided by the log change in trade costs. Middle row shows the change in the export participation rate (the fraction of firms with strictly positive market penetration). Bottom row shows change in the average value of m' for firms with $m' > 0$. First column shows the average transition for all markets. Second and third columns show average transitions for markets in the bottom 50% and top 10% of export participation, respectively. In each panel, blue line with circles is the baseline model; red line with squares is the static MP model; green line with diamonds is the sunk cost model; and purple line with 'x' markers is the exogenous NED model.

Baseline model. In the long run, export participation and market penetration both respond substantially more in hard destinations than in easy destinations. This is related to observations by Arkolakis (2010), Eaton et al. (2011), Kehoe and Ruhl (2013), and others that the least-traded products respond the most to trade reforms. Consequently, harder destinations have long-run trade elasticities more than 20 percent greater than easy destinations, consistent with the findings of Adão et al. (2020). The long-run trade elasticity in easy destinations is only slightly above $\theta - 1$. As Arkolakis (2010) points out, this is because exports are highly concentrated among the largest firms in these destinations, and "the largest firms in a market grow at a positive rate that (asymptotically) depends only on the price elasticity of demand." In the short run, trade takes several years to converge to its new higher level as new firms enter and incumbent exporters build up their customer bases. The export participation rate converges within about six years, while the average market penetration rate takes about ten years. In hard destinations, where exports are more evenly distributed across firms and these firm-level adjustments are more pronounced, the trade elasticity also takes about ten years to converge. In easy destinations, by contrast, trade flows converge almost immediately because firm-level adjustments are quantitatively less important.

Static MP model. The static market penetration model has similar long-run implications for trade flows as the baseline model. In both models, trade grows more in response to a permanent trade liberalization. This is because trade is more concentrated among large exporters in easy destinations than hard destinations, coupled with the fact that the convex market penetration costs that are present in both models imply that large exporters respond less to changes in trade costs than small exporters. The static market penetration model predicts slightly larger long-run responses to trade liberalizations in all destinations (in both easy and hard destinations, the long-run trade elasticity in this model is slightly higher than in the baseline) because there is more convexity in the cost of acquiring new customers than retaining old ones. Figure 7 also shows, however, that there is no gradual adjustment in trade in the static market penetration model: trade converges immediately to its long-run level. This is because the market penetration decision in this model is static, which implies that there is no persistence at the firm level in export participation or market penetration (after controlling for productivity and demand; persistence in firms' exogenous characteristics creates some persistence in export participation). This prediction of the static market penetration model is clearly counterfactual, given the widely documented evidence that trade adjusts gradually (see, e.g. Ruhl, 2008; Boehm et al., 2020).

Sunk cost model. Trade reforms in the sunk cost model have similar effects across destinations on the extensive margin of trade as in the baseline model. In both models, the number of exporters grows more in hard destinations than in easy ones following a permanent reduction in trade costs. However, there is little difference across destinations in the dynamics of aggregate bilateral trade: the long-run trade elasticity in the sunk cost model is only slightly higher in hard destinations than in easy ones. In the baseline model, by contrast, trade grows substantially more in the former than in the latter, consistent with the empirical findings of Kehoe and Ruhl (2013) and Arkolakis (2010).

Exogenous NED model. The exogenous NED model exhibits similar transition dynamics as the sunk cost model, except that it takes longer for the number of exporters—and thus aggregate trade flows—to converge. As before, the fact that this model predicts negligible differences across destinations in long-run trade elasticities is at odds with the data, which indicate that trade responds more to policy changes in hard destinations than in easy ones.

6 Conclusion

In this paper, I study how and why exporting firms' performance dynamics vary across destinations and explore the aggregate implications of these patterns. I first use microdata from Brazil to document that in smaller, poorer markets, overall turnover is higher, new entrants are larger and less likely to exit relative to incumbents, and successful exporters' sales grow less dramatically over the duration of their export spells as compared to larger, richer markets. To account for these facts, I develop a model of export market penetration dynamics that synthesizes static frameworks such as Arkolakis (2010) and Eaton et al. (2011) in which firms choose how many customers to serve in each destination with dynamic frameworks such as Das et al. (2007) and Alessandria and Choi (2007) in which sunk entry costs lead firms to make forward-looking export participation decisions.

The key feature of the model is that the cost of exporting is increasing in the number of customers a firm wants to serve but decreasing in the number of customers a firm already has, which means that firms choose endogenously to grow their customer bases gradually over time. As in Arkolakis (2010), the marginal cost of serving a single customer is strictly positive regardless of the size of a firm's current customer base, which generates endogenous entry but also exit in my dynamic setting. Acquiring and retaining customers is more expensive in smaller, poorer markets relative to these markets' purchasing power, which leads exporters to exit more frequently from these markets and accumulate fewer customers over the course of their export spells.

I calibrate the model so that it reproduces a subset of the facts I document in the empirical part of the paper, and validate it by demonstrating that it reproduces the remaining facts. To explore the role of market penetration dynamics in accounting for the facts at hand, I compare the model with several conventional alternatives that lack this feature. These alternative models can account for some of the variation across destinations in exporter performance dynamics, but market penetration dynamics are needed to capture the full range of this variation, particularly in the growth in successful exporters' sales over the duration of their export spells. I then use the calibrated model to explore how aggregate trade dynamics differ across destinations. As previous studies such as Ruhl and Willis (2017) and Alessandria and Choi (2014) have found, trade grows gradually in response to permanent trade reforms. In my model, though, exports to smaller, poorer markets grow more in the long run than exports to larger, richer ones, and these larger adjustments take longer to materialize. None of the alternative models are able to fully capture these patterns.

The analysis in this paper is limited to a partial equilibrium setting in which one firm's market penetration does not hinder or facilitate the entry and growth of other firms, and in which a firm's performance in one market does not affect its incentives to export to other markets. Studying the interactions between firms and between markets in equilibrium would allow one to answer questions such as: What is the role of congestion (or agglomeration) externalities in shaping export participation dynamics? How does a trade reform in one market affect export participation in other markets? The model developed in this paper is tractable enough to make incorporating these kinds of interactions feasible, making it a suitable starting point for a wide range of additional research.

References

Adão, Rodrigo, Costas Arkolakis, and Sharat Ganapati (2020) 'Aggregate implications of firm heterogeneity: A nonparametric analysis of monopolistic competition trade models.' NBER Working Paper 28081

Alessandria, George, and Horag Choi (2007) 'Do sunk costs of exporting matter for net export dynamics?' *The Quarterly Journal of Economics* 122(1), 289–336

_ (2014) 'Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization.' *Journal of International Economics*

Alessandria, George, and Horag Choi (2016) 'The Dynamics of the U.S. Trade Balance and Real Exchange Rate: The J Curve and Trade Costs?' Working Paper, February

Alessandria, George, Costas Arkolakis, and Kim Ruhl (2021a) 'Firm dynamics and trade.' *Annual Review of Economics*

Alessandria, George, Horag Choi, and Kim Ruhl (2021b) 'Trade Adjustment Dynamics and the Welfare Gains from Trade.' *Journal of International Economics*

Arkolakis, Costas (2010) 'Market Penetration Costs and the New Consumers Margin in International Trade.' *Journal of Political Economy* 118(6), 1151–1199

Arkolakis, Costas (2016) ' A Unified Theory of Firm Selection and Growth *.' *The Quarterly Journal of Economics* 131(1), 89–155

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg (2019) 'Benchmarking global optimizers.' Working Paper 26340, National Bureau of Economic Research, October

Baldwin, Richard E (1992) 'Measurable Dynamic Gains from Trade.' *Journal of Political Economy* 100, 162–74

Bernard, Andrew B., and J. Bradford Jensen (2004) 'Why Some Firms Export.' *The Review of Economics and Statistics* 86(2), 561–569

Bernard, Andrew B., J. Bradford Jensesn, Stephen J. Redding, and Peter K. Schott (2012) 'The Empirics of Firm Heterogeneity and International Trade.' *Annual Review of Economics* 4(1), 283–313

Boehm, Christoph, Andrei A. Levchenko, and Nitya Pandalai-Nayar (2020) 'The Long and Short (Run) of Trade Elasticities.' CEPR Discussion Papers 14645

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin (2011) 'Finance and development: A tale of two sectors.' *American Economic Review* 101(5), 1964–2002

Chaney, Thomas (2008) 'Distorted gravity: The intensive and extensive margins of international trade.' *American Economic Review* 98(4), 1707–21

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout (2007) 'Market Entry Costs, Producer Heterogeneity, and Export Dynamics.' *Econometrica* 75(3), 837–873

Eaton, Jonathan, Marcela Eslava, Maurice Kugler, and James Tybout (2007) 'Export Dynamics in Colombia: Firm-Level Evidence.' NBER Working Paper 13531

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2011) 'An Anatomy of International Trade: Evidence From French Firms.' *Econometrica* 79(5), 1453–1498

Fernandes, Ana M., Caroline Freund, and Martha Denisse Pierola (2016) 'Exporter behavior, country size, and stage of development: Evidence from the exporter dynamics database.' *Journal of Development Economics* 119, 121 – 137

Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi (2023) 'How Exporters Grow.' *The Review of Economic Studies* p. rdad070

Gumpert, Anna, Haishi Li, Andreas Moxnes, Natalia Ramondo, and Felix Tintelnot (2020) 'The life-cycle dynamics of exporters and multinational firms.' *Journal of International Economics* 126, 103343

Kehoe, Timothy J., and Kim J. Ruhl (2013) 'How important is the new goods margin in international trade?' *Journal of Political Economy* 121(2), 358–392

Melitz, Marc J. (2003) 'The impact of trade on intra-industry reallocations and aggregate industry productivity.' *Econometrica* 71(6), 1695–1725

Ottaviano, Gianmarco, and Thierry Mayer (2007) *The happy few: the internationalisation of European firms* (Bruegel)

Piveteau, Paul (2020) 'An empirical dynamic model of trade with consumer accumulation.' *American Economic Journal: Microeconomics*

Rowan, Tom (1990) 'Functional stability analysis of numerical algorithms.' Ph.D. Thesis, Department of Computer Sciences, University of Texas at Austin

Ruhl, Kim J. (2008) 'The international elasticity puzzle.' Working paper

Ruhl, Kim J., and Jonathan L. Willis (2017) 'New Exporter Dynamics.' *International Economic Review* 58, 703–726

Steinberg, Joseph B. (2019) 'Brexit and the macroeconomic impact of trade policy uncertainty.' *Journal of International Economics* 117(C), 175–195

Appendix for "Export Market Penetration Dynamics"

Joseph B. Steinberg

Appendix A provides robustness checks for the main empirical and quantitative results reported in the paper. Appendix B provides additional results about variation in exporters' performance within their own individual portfolios of export markets. Appendix C repeats the empirical analysis using publicly-available data from Mexico and Peru. The relevant figures and tables are at the end of each section. For details about the computer code required to solve the model and replication the analyses contained in this paper, please see my GitHub repository https://github.com/joesteinberg/dyn-mkt-pen.

A Robustness and sensitivity analysis

This section contains additional robustness checks and sensitivity analyses related to the main results reported in the text of the paper.

A.1 Robustness: statistical significance of differences in sales trajectories

Section 2.2 in the paper shows that the effect of time in a market on sales differs systematically across markets: exporters sell more initially and grow more over the duration of their export spells in easy markets than in hard markets. I take two approaches to formally testing the statistical significance of the differences in the regression coefficients $\beta_{m,n,g}$ between hard and easy markets. The first approach tests the null hypothesis that $\beta_{m,n,1} = \beta_{m,n,0}$. Since we are interested primarily in whether $\beta_{m,n,1} > \beta_{m,n,0}$, I compute one-sided p-values. All of these p-values are tiny; the highest is 0.00002% for spell length = 5 and tenure = 1. The second approach looks at the significance of the linear combination of the coefficients $\beta_{m,n,1} - \beta_{m,n,0}$. All of the 95% confidence intervals for these combinations are well above zero. Table A.1 reports these results.

A.2 Sensitivity analysis: differences between micro and macro returns to market size

In an earlier version of this paper I did not distinguish between micro- and macroeconomic returns to scale in market penetration, i.e., I set $\alpha_n = \beta_n$ and $\alpha_o = \beta_o$. The distinction made in the current version arose from a very helpful comment by Jonathan Eaton. The motivation is as follows. The size of the overall market (in a macro sense) dictates the effectiveness of an additional dollar of advertising; a single TV or radio ad will likely reach more eyes/ears in a larger country. But in a dynamic context, not all eyes/ears are the same. It may be more or less difficult to target advertisements towards a specific sub-population of the overall market. If it is easy to target people (a firm can direct its ads specifically towards potential new/old customers), then β_n should be lower than α_n and β_o should be higher than α_o. But if this is not possible, i.e., exposure to advertising is essentially randomized across the entire population, then this should not be true.

Note that in a static context (i.e., in the static market pentration model) this distinction is irrelevant, because the entire market consists of potential new customers. But in the dynamic model I study in this paper, the distinction matters because the number of potential new customers shrinks and the number of old customers grows as a firm penetrates a market.

To study how much this distinction matters, I look at four alternative calibrations:

- ($\alpha_n = \beta_n$) Increase α_n to the calibrated value of β_n of 0.94, which makes it easier to attract new customers in harder markets;
- ($\beta_n = \alpha_n$) Reduce β_n to the calibrated value of α_n of 0.56, which lowers the marginal cost of attracting new customers for entrants relative to incumbents;
- ($\alpha_o = \beta_o$) Reduce α_o to the calibrated value of β_o of 0.79, which makes it harder to retain old customers in harder markets; and
- ($\beta_o = \alpha_o$) Increase β_o to the calibrated value of α_o of 0.96, which lowers the marginal cost of retaining old customers for entrants relative to incumbents.

The results in all four alternative calibrations, which are shown in Table A.2 and figures A.1–A.2, are similar to the baseline results. The effects on turnover and the relative size of entrants follow directly from the explanations above. For example, in the first alternative, more firms export to harder markets and the exit rate is lower. The largest differences are seen in the effects of tenure and spell length on sales. In the first and third alternatives, sales grow more quickly over the course of longer export spells, similar to the static market penetration model. This indicates that luck (getting good demand shocks) plays a larger role in these alternatives than in the baseline. In the second and fourth alternatives, sales grow more slowly because the marginal market penetration cost changes less over the course of an export spell. Nevertheless, the main results are all still there, including the differences in sales trajectories between hard and easy markets.

A.3 Sensitivity analysis: no macro returns to market size

In addition to exploring the sensitivity of my results to differences between micro and macro returns to market size, I have also looked at what happens when there are no macro returns to market size at all. When $\alpha_n = \alpha_o = 1$, the per-customer marginal acquisition/retention costs are unaffected by population size. The results of this analysis are are also shown in Table A.2 and figures A.1–A.2. Not surprisingly, they are most similar to the results for the calibration where $\alpha_n = \beta_n = 0.94$, which is essentially the same.

Table A.1: Effects of tenure-spell length in hard vs. easy markets

		Coeff		(1) H0 test: Hard = Easy		(2) Linear combo: Easy - Hard					
Spell length	Tenure	Hard	Easy	F-stat	One-sided p-value	Coeff.	Standard error	t-stat	p-value	CI (lower)	CI (upper)
2	1	0.264	0.448	17.6	1.38e-05	0.184	0.0439	4.19	2.77e-05	0.0979	0.27
2	2	0.115	0.323	24.2	4.36e-07	0.209	0.0424	4.92	8.74e-07	0.125	0.292
3	1	0.454	0.741	19.8	4.29e-06	0.288	0.0647	4.45	8.59e-06	0.161	0.414
3	2	0.617	0.986	37.6	4.35e-10	0.37	0.0603	6.13	8.74e-10	0.251	0.488
3	3	0.324	0.642	27.7	6.96e-08	0.318	0.0604	5.27	1.4e-07	0.2	0.436
4	1	0.443	0.893	28.2	5.53e-08	0.45	0.0847	5.31	1.11e-07	0.284	0.616
4	2	0.732	1.287	49.5	9.77e-13	0.555	0.0788	7.04	1.97e-12	0.4	0.709
4	3	0.681	1.275	60	4.7e-15	0.594	0.0767	7.75	9.52e-15	0.444	0.745
4	4	0.317	0.863	49	1.31e-12	0.546	0.0781	7	2.65e-12	0.393	0.699
5	1	0.615	1.020	12.4	0.000218	0.405	0.115	3.52	0.000436	0.179	0.631
5	2	0.965	1.452	21.7	1.61e-06	0.488	0.105	4.66	3.23e-06	0.282	0.693
5	3	1.121	1.613	24.3	4.11e-07	0.492	0.0998	4.93	8.24e-07	0.296	0.688
5	4	0.959	1.587	40.9	7.86e-11	0.628	0.0982	6.4	1.58e-10	0.436	0.82
5	5	0.587	1.176	34.5	2.12e-09	0.589	0.1	5.87	4.27e-09	0.393	0.786
6	1	0.607	1.151	23.4	6.44e-07	0.544	0.112	4.84	1.29e-06	0.324	0.764
6	2	1.023	1.680	47	3.46e-12	0.657	0.0957	6.86	6.98e-12	0.469	0.844
6	3	1.349	1.831	30.1	2.01e-08	0.483	0.0879	5.49	4.04e-08	0.31	0.655
6	4	1.318	1.949	57.8	1.45e-14	0.631	0.083	7.6	2.93e-14	0.469	0.794
6	5	1.244	2.023	94.8	1.04e-22	0.779	0.08	9.74	2.15e-22	0.622	0.936
6	6	1.257	2.059	236	1.67e-53	0.803	0.0523	15.4	4.07e-53	0.7	0.905

Notes: Table reports results from two approaches to measuring the statistical significance of the differences between the effects of tenure-spell length on sales in hard versus easy markets (i.e., differences between $\beta_{m,n,1}$ and $\beta_{m,n,0}$ from specification (3) in the main text). Section (1) reports the F-statistic and one-sided p-value from the test of the null hypothesis that $\beta_{m,n,1} = \beta_{m,n,0}$. This approach uses Stata code from `https://www.stata.com/support/faqs/statistics/one-sided-tests-for-coefficients`. Section (2) reports results from using Stata's `lincom` command with the argument $\beta_{m,n,1} - \beta_{m,n,0}$.

Table A.2: Exporter performance across markets: sensitivity analysis

Statistic/coefficient	Num. exporters	Top-5 share	Avg. num. dests.	Exit rate	Entrant rel. size	Entrant rel. exit rate
(a) Cross-market averages						
$\alpha_n = \beta_n$	1,491	0.63	18.71	0.48	0.22	0.14
$\beta_n = \alpha_n$	937	0.58	17.53	0.45	0.38	0.07
$\alpha_o = \beta_o$	822	0.56	15.21	0.52	0.32	0.11
$\beta_o = \alpha_o$	1,229	0.62	23.23	0.36	0.39	0.07
$\alpha_n = \alpha_o = 1$	1,590	0.63	19.79	0.48	0.21	0.14
(b) Associations with destination characteristics: $\alpha_n = \beta_n$						
log GDPpc	0.761	0.071	-2.149	-0.074	-0.135	0.036
log population	0.234	0.015	-1.207	0.024	-0.099	0.038
log trade barrier	-0.744	-0.069	2.217	0.072	0.128	-0.039
(c) Associations with destination characteristics: $\beta_n = \alpha_n$						
log GDPpc	0.748	0.063	-1.685	-0.071	-0.099	0.028
log population	0.357	0.031	-0.836	-0.034	-0.047	0.013
log trade barrier	-0.737	-0.062	1.684	0.071	0.100	-0.028
(d) Associations with destination characteristics: $\alpha_o = \beta_o$						
log GDPpc	0.843	0.080	-2.051	-0.075	-0.302	0.033
log population	0.120	0.003	-1.174	0.070	-0.525	0.036
log trade barrier	-0.811	-0.078	2.152	0.071	0.341	-0.033
(e) Associations with destination characteristics: $\beta_o = \alpha_o$						
log GDPpc	0.781	0.086	-1.362	-0.124	-0.072	0.082
log population	0.123	0.014	-0.814	0.053	-0.132	0.048
log trade barrier	-0.747	-0.084	1.607	0.112	0.082	-0.070
(f) Associations with destination characteristics: $\alpha_n = \alpha_o = 1$						
log GDPpc	0.622	0.062	-2.155	-0.064	-0.069	0.025
log population	0.009	0.001	-0.153	-0.002	0.003	0.001
log trade barrier	-0.622	-0.062	2.158	0.064	0.071	-0.024

Figure A.1: Effects of tenure and spell length on sales: sensitivity analysis

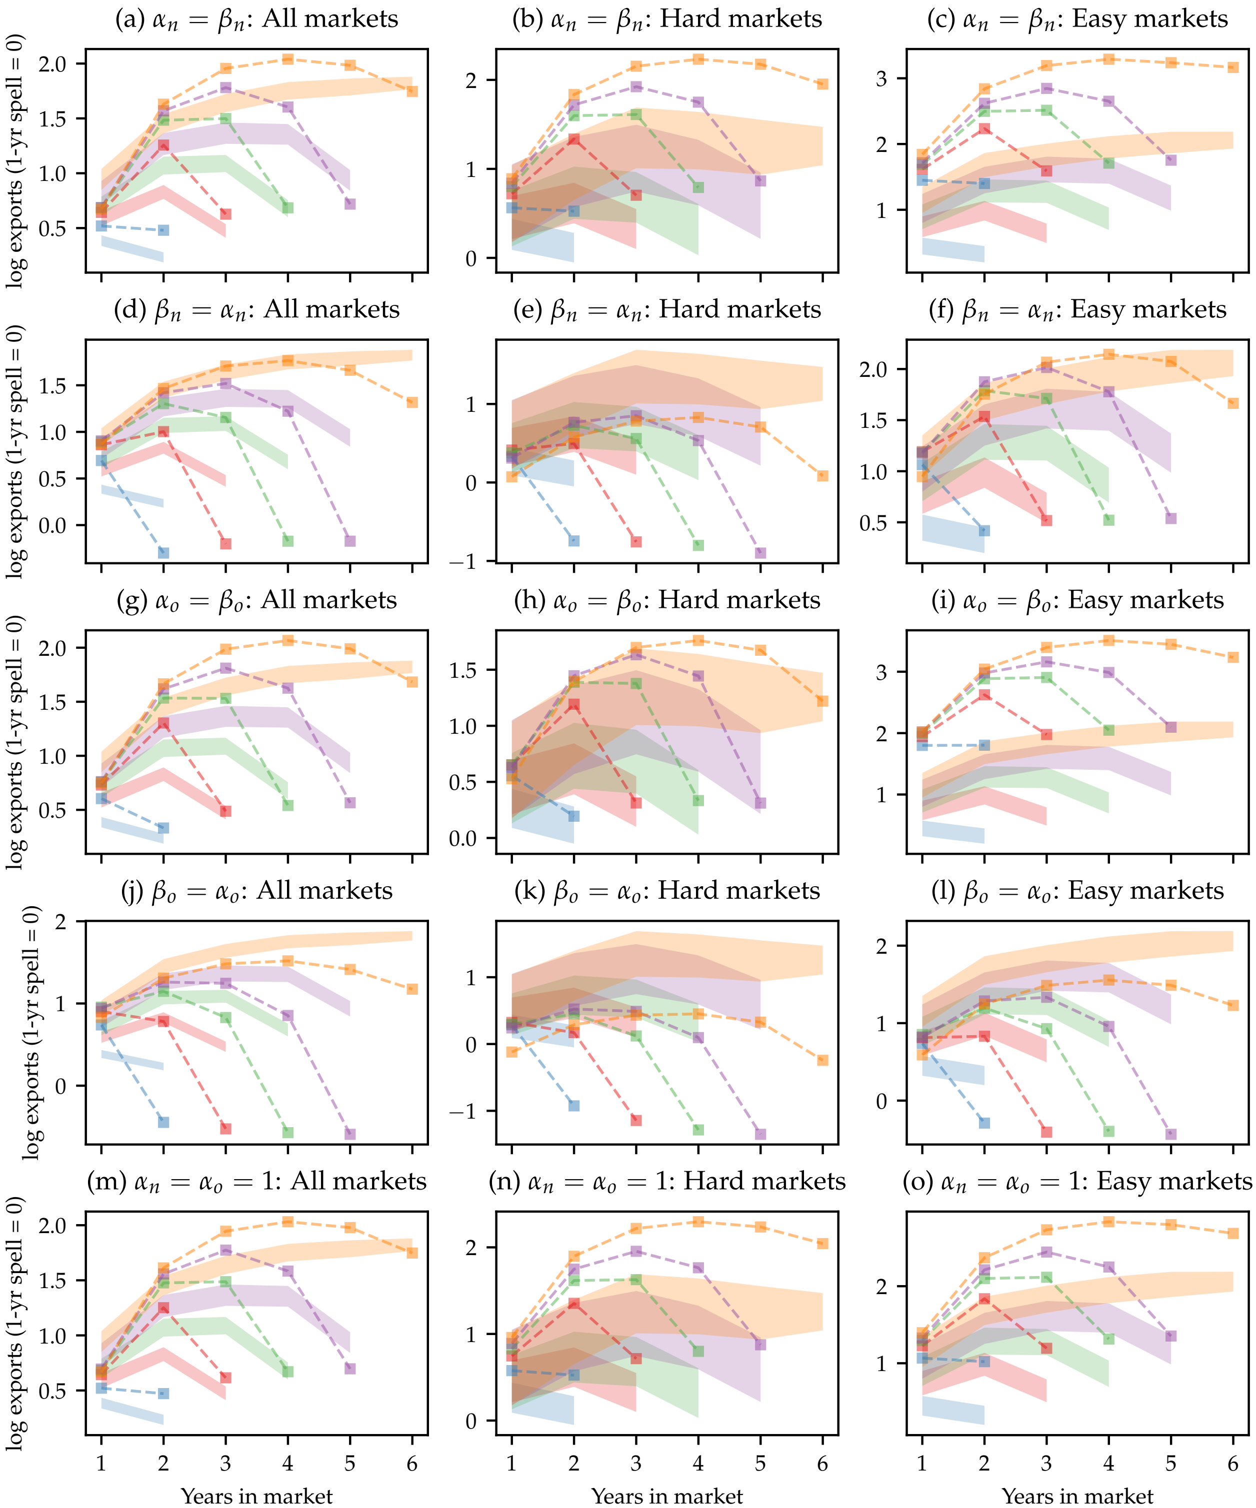

Notes: Each row shows simulated results from a different model specification. First column shows estimates of $\beta_{m,n}$ from specification (2) in the main text. Second column shows estimates of $\beta_{m,n,g}$ from specification (3) for markets in the bottom 50% of export participation, and third column shows estimates for markets in the top 90%. Each line shows $\beta_{m,1}, \beta_{m,2}, \ldots, \beta_{m,6}$ (or $\beta_{m,1,g}, \beta_{m,2,g}, \ldots$ in the second two panels) for a set value of m. In both columns, shaded areas show 95-percent confidence intervals for estimates using actual data discussed in section 2.2.

Figure A.2: Exit rates conditional on tenure: sensitivity analysis

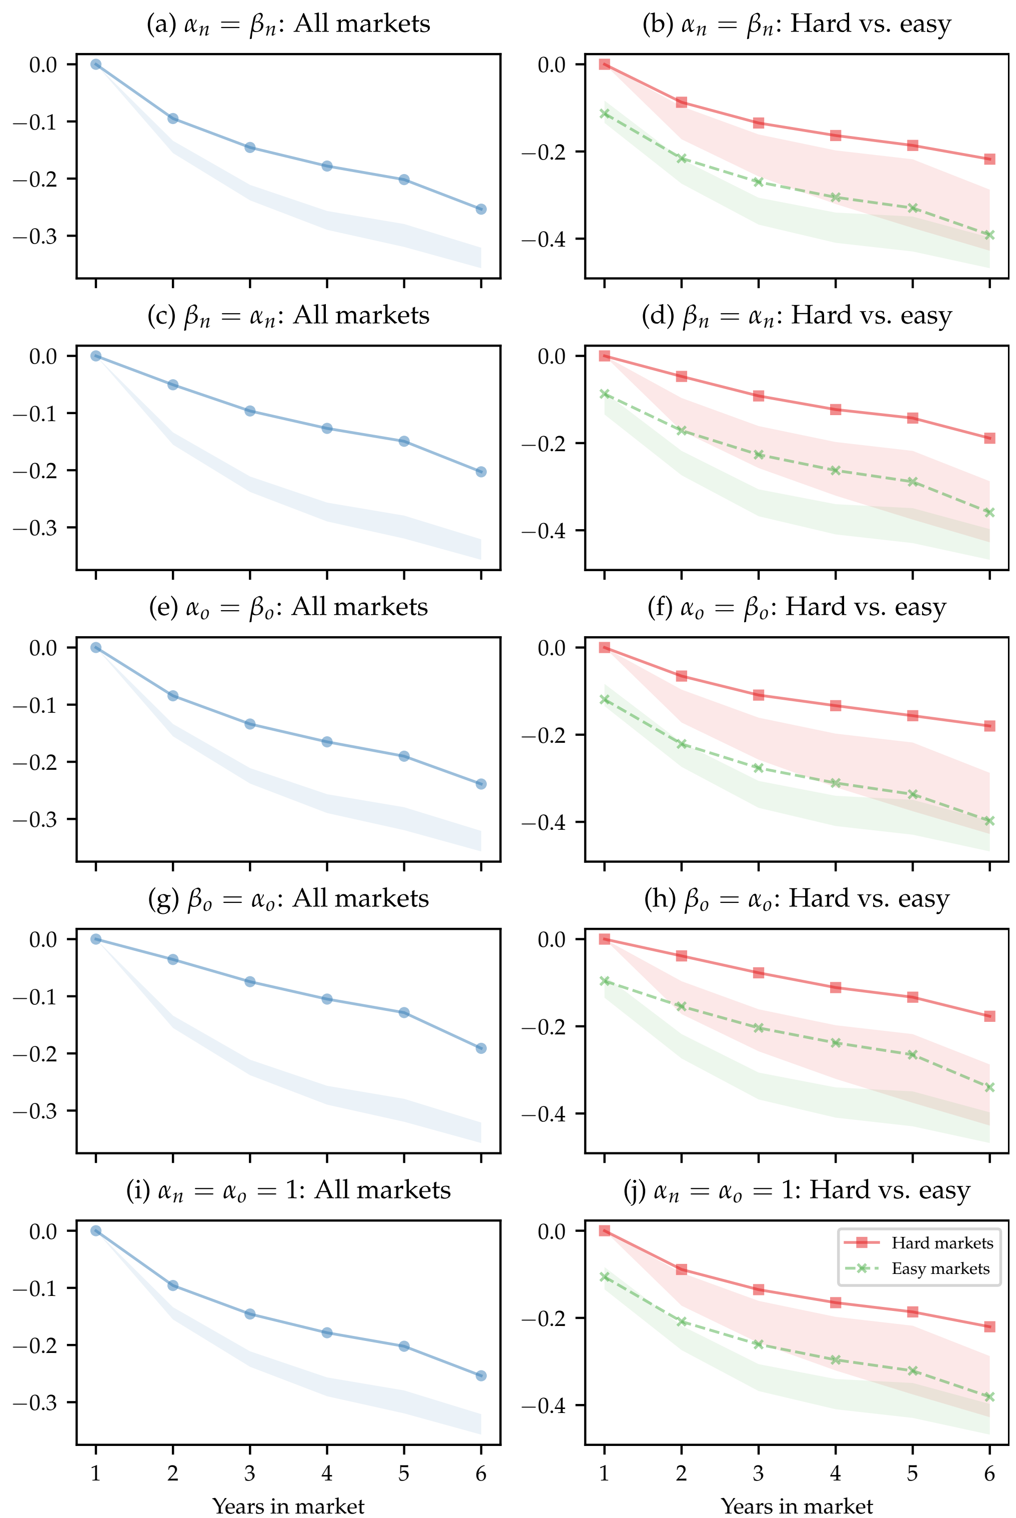

Notes: Each row shows simulated results from a different model specification. First column shows estimates of β_n from specification (4) in the main text. Second column shows estimates of $\beta_{n,g}$ from specification (5). In second column, solid red lines with square markers show estimates for markets in bottom 50% of export participation, and dashed green lines with 'x' markers show estimates for destinations in top 90%. In both columns, shaded areas show 95-percent confidence intervals for estimates using actual data discussed in section 2.2.

B Additional results

This section presents additional facts about exporter performance not discussed in the main text of the paper.

B.1 Additional cross-sectional statistics

It is well known that there are more small exporters and exports are more concentrated among top exporters in easier markets. In section 2.1 in the main text of the paper, I confirm that this is true in Brazil by showing that the share of exports accounted for the top 5% of exporters is higher in easier markets. Here, I dig further into how the cross-sectional distribution of exports differs across markets.

I compute the 25th, 75th, and 95th percentiles of exports for each market, normalized by the average exports to that market. Panel (a) of Table B.1 shows that there is a great deal of variation in these percentiles across markets, whether one looks at all firms, or incumbents or entrants only. Panel (b) shows that this variation is systematic by estimating

$$\log P_{j,t} = \alpha + \beta \log L_{j,t} + \gamma \log Y_{j,t} + \delta \log \tau_{j,t} + f_t + \epsilon_{j,t}, \tag{B.1}$$

where $P_{j,t}$ is one of the normalized percentiles in questions. All three percentiles are decreasing in output per capita and population and increasing in trade costs, but the 25th percentile is the most responsive to market characteristics while the 95th is the least responsive. This holds whether one looks at all firms, or incumbents of entrants only. This is consistent with previous findings in the literature that easier markets have more small exporters and that exports in these markets are more concentrated among the largest exporters. Panels (c) and (d) show that the model is consistent with all of these results.

B.2 Variation in performance within exporters' destination portfolios

This section of the appendix contains additional results about how how individual exporters' performance varies across the destinations to which they sell. I first group firms by the number of destinations in their export "portfolios." Figure B.1 shows the distribution of exporters and their total sales across all destinations by the sizes of their portfolios. 40 percent of exporters only sell to one destination. Of the remaining 60 percent, most sell to between 2–4 destinations. Only 12 percent of firms export to 10 or more destinations, but these firms account for about 75 percent of total exports in any given year, whereas firms that export to 4 or fewer destinations account for barely 10 percent. This finding is consistent with the "superstar" phenomenon documented elsewhere in the literature. The distribution of exporters by destinations served in the calibrated model is close to the empirical distribution, as is the distribution of exports. This confirms that the model captures how the cross sections of exporters varies across destinations.

I then rank the destinations within each firm's portfolio by sales and analyze how firms perform in high- vs. low-ranked destinations. Harder destinations have smaller populations with lower purchasing power and higher trade barriers, and so one would expect these destinations to be ranked lower on average in exporters' portfolios as well as having lower export participation overall. Table B.2, which reports associations between destinations' characteristics and their average ranks within exporters' portfolios, confirms that this is the case. The average rank of a destination within an exporter's portfolio is decreasing in population and income per capita, and increasing in trade barriers: firms export less to harder destinations than easy ones.[1] The calibrated model captures these patterns.

B.2.1 Turnover within exporters' destination portfolios

Since harder destinations have higher overall exit rates, one would therefore expect that firms are more likely to exit lower-ranked destinations within their portfolios. Table B.3, which lists exporters' average exit rates broken down by portfolio size (vertical axis) and destination rank (horizontal axis), confirms that this is the case as well: exporters that serve several markets have the highest exit rates in their least important destinations. However, many-destination exporters are less likely to exit from their least important destinations than single-destination exporters are from their sole destination. The calibrated model matches the propensity of multi-destination firms to exit more frequently from their least-important destinations.

B.2.2 Variation in sales and survival within exporters' destination portfolios

To analyze how the rank of a destination within a firm's portfolio affect its sales and likelihood of exit relative to other firms that export to that destination, I estimate regressions of the form

$$\log ex_{i,j,t} = \alpha + \sum_{m=1}^{10}\sum_{n=1}^{m} \beta_{m,n} \mathbb{1}_{\{\text{num. dests.}_{i,j}=m\}} \mathbb{1}_{\{\text{dest. rank}_{i,j,t}=n\}} + f_{j,t} + \epsilon_{i,j,t}, \tag{B.2}$$

$$\mathbb{1}_{\{exit_{i,j,t}=1\}} = \alpha + \sum_{m=1}^{10}\sum_{n=1}^{m} \beta_{m,n} \mathbb{1}_{\{\text{num. dests.}_{i,j}=m\}} \mathbb{1}_{\{\text{dest. rank}_{i,j,t}=n\}} + f_{j,t} + \epsilon_{i,j,t}. \tag{B.3}$$

Here, I top-code portfolio size and destination rank at 10. The coefficient $\beta_{m,n}$ in each regression measures how much more a firm sells (in the first specification) or how likely it is to exit (the second specification) in a given destination relative to a firm for whom this destination is its only market. Note that unlike the replication of Fitzgerald et al. (2023) in section 2.2, I do not include firm-year fixed effects or even firm fixed effects at all. It is not possible to do so while also including a control for the number of destinations a firm serves. Figure B.2 shows the results. In panel (a), we see that firms that export to at least 2 destinations sell more in their highest-ranked destination than firms that export to that destination only. The larger a firm's portfolio, the greater

[1] Estimating a Poisson or negative binomial regression on the raw firm-level data yields similar results.

the difference: firms that sell to 2 destinations sell about twice as much in their highest-ranked market as single-destination exporters; while firms with portfolios of 10 or more destinations sell about 5 times as much. We also see, however, that sales relative to single-destination exporters fall with a destination's rank. In fact, all firms except those with 10 or more destinations in their portfolios sell less in their lowest-ranked destinations than single-destination firms. In panel (b), we see these patterns reversed for exit rates. Multi-destination firms are less likely to exit from their most important destinations than single-destination firms—as much as 45 p.p. less likely for firms with 10 or more destinations. However, the gap shrinks as a destination's rank with an exporter's portfolio rises; firms that sell to 4 or fewer destinations are actually more likely to exit from their least important destinations than single-destination exporters.

The model predicts similar patterns of sales and exit rates across destinations within individual exporters' portfolios as observed in the data. In the model, as in the data, exporters with larger portfolios sell more in their top-ranked destinations relative to single-destination exporters, but sell less in their least important destinations. Similarly, exporters with larger portfolios are less likely to exit from their most important destinations as compared to single-destination exporters, and exporters with 4 or fewer destinations have about the same exit rate in their lowest-ranked destinations as single-destination exporters. Together, these tables and figures show that the model accurately captures the facts about how export performance varies within exporters' destination portfolios discussed above.

B.2.3 Variation in exporting costs within firms' destination portfolios

To dig more deeply into cross-firm variation in exporting costs and to analyze how individual firms' exporting costs vary across destinations, I follow the approach from section B.2.2 and estimate the effect of a destination's rank within an exporter's portfolio on the cost that exporter pays to access that destination, both in levels and relative to profits, using similar specifications to (B.2).

Panel (a) of Figure B.3 reports the estimated effects of destination rank on the level of exporting costs, using the same kind of specification as in (B.2) Exporters with the largest destination portfolios pay the highest exporting costs, especially in higher-ranked destinations. Firms that serve 10 or more destinations, pay about 4 times more to export to their top destinations than firms that serve those destinations only, and even firms that sell to only 2 destinations pay twice as much as single-destination exporters. Export costs fall with destination rank, however, and firms with 9 or fewer destinations in their portfolios (the vast majority) actually pay less to export to their least important destinations than firms that serve those destinations alone. Panel (b) of Figure B.3, which reports the results from estimating specification (B.2) with the ratio of export costs to profits as the dependent variable, shows that these results reverse when export costs are measured relative to profits. Firms with larger destination portfolios have lower export cost/profit ratios, especially in their highest-ranked destinations, and these ratios rise as destination rank falls. In brief, these results show that high-productivity and/or high-demand firms pay higher costs to export, but these costs are low relative to the large profits these firms earn from exporting.

Table B.1: Percentiles of normalized exports across markets

	All firms			Incumbents			Entrants		
Statistic/coefficient	25th	75th	95th	25th	75th	95th	25th	75th	95th
(a) Summary statistics (data)									
Mean	0.061	0.619	3.558	0.072	0.698	3.621	0.104	0.735	3.297
Min	0.004	0.160	2.435	0.006	0.203	2.182	0.017	0.284	2.178
Max	0.241	1.290	4.857	0.304	1.291	4.914	0.354	1.316	4.681
Std. dev.	0.061	0.297	0.581	0.071	0.283	0.525	0.084	0.307	0.478
(b) Associations with market characteristics (data)									
log GDPpc	-0.337	-0.215	-0.030	-0.319	-0.179	-0.023	-0.193	-0.154	-0.027
	(0.047)§	(0.031)§	(0.015)†	(0.046)§	(0.029)§	(0.013)*	(0.048)§	(0.031)§	(0.016)*
log population	-0.439	-0.164	-0.001	-0.403	-0.134	0.022	-0.356	-0.144	0.005
	(0.038)§	(0.025)§	(0.011)	(0.037)§	(0.024)§	(0.011)†	(0.034)§	(0.026)§	(0.012)
log trade barrier	0.487	0.221	0.010	0.452	0.178	-0.002	0.323	0.192	0.031
	(0.050)§	(0.033)§	(0.016)	(0.048)§	(0.032)§	(0.016)	(0.055)§	(0.034)§	(0.016)*
Num. observations	627	627	627	627	627	627	627	627	627
R^2	0.64	0.44	0.05	0.59	0.35	0.05	0.44	0.26	0.05
(c) Summary statistics (baseline model)									
Mean	0.031	0.604	4.276	0.026	0.671	4.518	0.078	0.788	3.831
Min	0.007	0.287	3.448	0.003	0.351	2.790	0.041	0.481	3.082
Max	0.112	1.171	5.051	0.173	1.423	5.608	0.131	1.156	4.262
Std. dev.	0.022	0.216	0.374	0.029	0.241	0.516	0.020	0.158	0.256
(d) Associations with market characteristics (baseline model)									
log GDPpc	-0.349	-0.235	-0.049	-0.265	-0.200	-0.018	-0.139	-0.131	-0.039
log population	0.123	-0.058	-0.025	0.351	0.007	-0.026	-0.081	-0.061	-0.010
log trade barrier	0.327	0.234	0.051	0.228	0.199	0.015	0.134	0.130	0.046

Notes: Panel (a) reports summary statistics. Panel (b) reports associations with destination characteristics. All specifications include year fixed effects. Standard errors are clustered at the market level. §, ‡, and † denote significance at the 0.1%, 1%, and 5% levels, respectively. The first three columns show results for the 25th, 75th, and 95th percentiles of the export distribution for all firms, normalized by the average export volume. The second three columns show results for incumbent exporters, and the last three columns show results for new entrants.

Table B.2: Associations between destination characteristics and average rank

	Data	Model
log GDPpc	-1.260	-1.987
	(0.216)§	
log population	-1.148	-2.530
	(0.138)§	
log trade barrier	2.254	2.079
	(0.283)§	
Num. observations	627	
R^2	0.55	

All specifications control for year fixed effects. Standard errors are clustered at the market level. §, ‡, and † denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table B.3: Exit rates by num. dest. and rank

	Destination rank					
Num. dest.	1	2	3	4	5-9	10+
(a) Data						
1	0.56	-	-	-	-	-
2	0.41	0.60	-	-	-	-
3	0.31	0.47	0.61	-	-	-
4	0.23	0.36	0.49	0.60	-	-
5-9	0.16	0.24	0.32	0.40	0.52	-
10+	0.06	0.08	0.10	0.13	0.21	0.31
(b) Model						
1	0.53	-	-	-	-	-
2	0.39	0.53	-	-	-	-
3	0.30	0.42	0.52	-	-	-
4	0.24	0.35	0.43	0.52	-	-
5-9	0.16	0.24	0.30	0.35	0.42	-
10+	0.05	0.08	0.11	0.13	0.18	0.27

Table reports exit rates by the number of markets to which a firm sells (rows) and the rank of a destination within a firm's portfolio (columns). Panel (a) reports results for the data and panel (b) shows results for simulated data from the baseline model.

Figure B.1: Distribution of exporters and exports by number of destinations

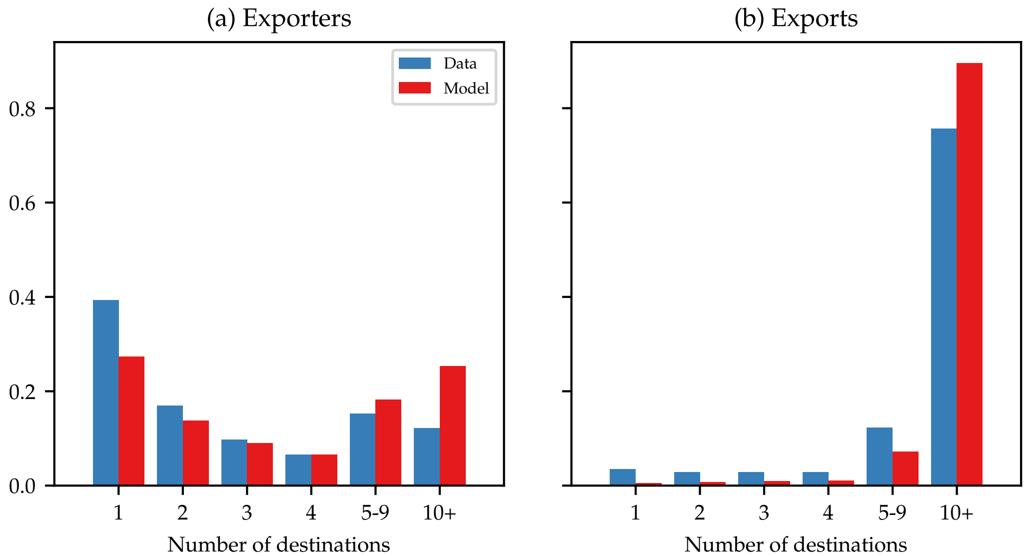

Notes: Panel (a) shows the fractions of firms that serve 1, 2, 3, 4, 5–9, and 10 or more destinations, respectively. Panel (b) shows the fraction of total exports that are accounted for by firms in these groups. In each panel, blue bars show the actual data and red bars show simulated data from the baseline model.

Figure B.2: Sales and exit rates by number of destinations served and destination rank

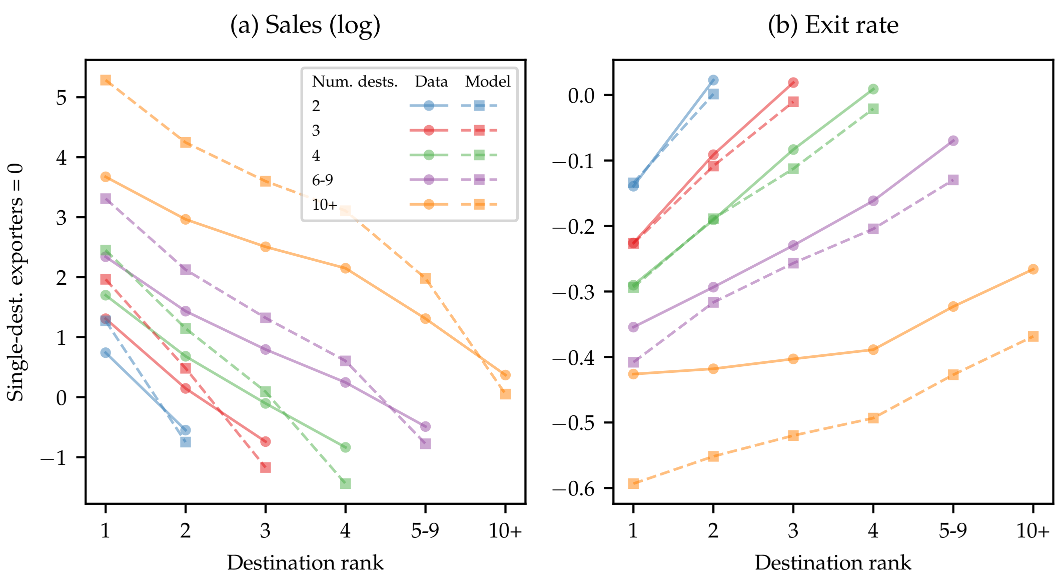

Notes: Panel (a) shows estimates of $\beta_{m,n}$ from (B.2), and panel (b) shows estimates (B.3) In each panel, solid lines show estimates from the actual data and dashed lines show estimates from the simulated data from the baseline model.

Figure B.3: Exporting costs by number of destinations served and destination rank

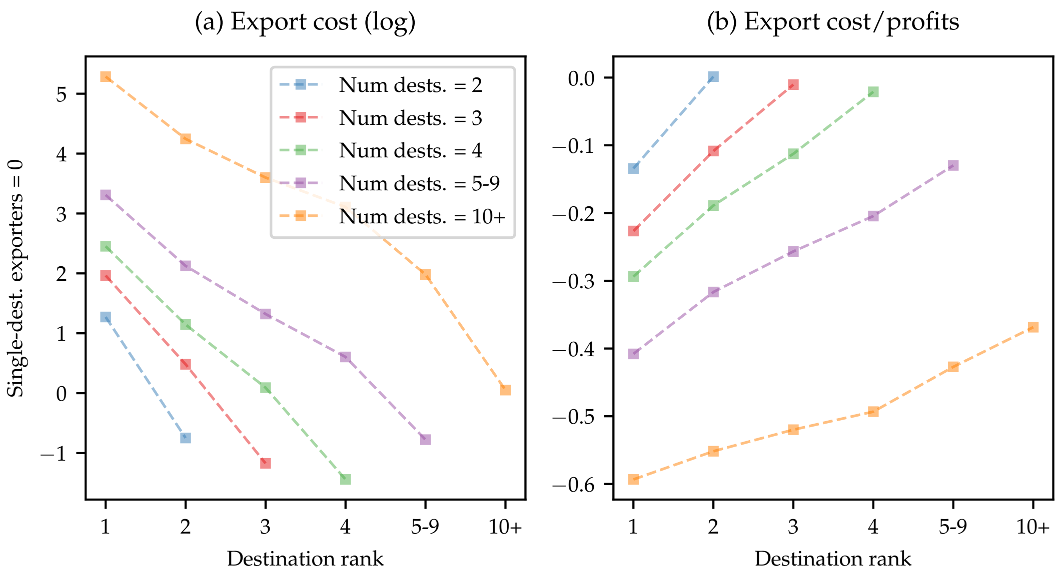

Notes: Figure shows estimates of $\beta_{m,n}$ from a version of (B.2) where measures of export costs are the dependent variables. Panel (a) shows estimates from the specification where log export costs are the dependent variable, and panel (b) shows estimates from the specification where the the ratio of export costs to profits the dependent variable. In each panel, solid lines show estimates from the actual data and dashed lines show estimates from the simulated data from the baseline model.

C Empirical results for Mexico and Peru

In this appendix, I report results from empirical analysis of two additional datasets on Mexican and Peruvian exporters from the World Bank's Exporter Dynamics Database. These transaction-level customs datasets have the same structure as the Brazilian data. The Mexican dataset covers the period 2001–2006 and contains about 23,000 firms per year.[2] The Peruvian dataset covers a longer time period, 1994–2008 but contains fewer firms, ranging from 2000 at the beginning of the sample to 5000 at the end. I apply exactly the same processing and analysis procedures described in section **??** to these datasets. Table C.1 and figures C.1–C.2 show the results.

[2]The dataset contains information on transactions through 2008, but there is a break in the coding of firm identifiers in 2007.

Table C.1: Destination-level measures of exporter performance (Mexico and Peru)

Statistic	Num. exporters	Top-5 share	Avg. num. dests.	Exit rate	Entrant rel. size	Entrant rel. exit rate
(a) Summary statistics (Mexico)						
Mean	682	0.67	14.61	0.46	0.37	0.33
Min	24	0.44	2.22	0.36	0.06	0.16
Max	16,908	0.92	24.28	0.60	1.13	0.45
Std. dev.	2,196	0.12	5.23	0.06	0.27	0.05
(b) Summary statistics (Peru)						
Mean	133	0.62	10.36	0.48	0.36	0.32
Min	21	0.31	3.63	0.35	0.07	0.16
Max	1,003	0.91	17.61	0.67	0.96	0.43
Std. dev.	178	0.15	3.68	0.07	0.21	0.06
(c) Associations with destination characteristics (Mexico)						
log GDPpc	0.299	0.065	-1.594	0.014	-0.096	0.003
	(0.042)§	(0.009)§	(0.264)§	(0.004)§	(0.023)§	(0.004)
log population	0.372	0.040	-1.244	-0.002	-0.047	0.006
	(0.031)§	(0.007)§	(0.241)§	(0.003)	(0.015)‡	(0.003)†
log trade barrier	-0.713	-0.031	2.699	0.020	0.084	-0.012
	(0.029)§	(0.007)§	(0.230)§	(0.003)§	(0.017)§	(0.003)§
Num. observations	288	288	288	288	288	288
R^2	0.84	0.51	0.70	0.45	0.13	0.05
(d) Associations with destination characteristics (Peru)						
log GDPpc	0.325	0.077	-0.956	-0.001	-0.062	0.008
	(0.092)§	(0.012)§	(0.296)‡	(0.008)	(0.020)‡	(0.008)
log population	0.236	0.042	-0.358	-0.010	-0.058	0.002
	(0.065)§	(0.014)‡	(0.198)*	(0.006)*	(0.013)§	(0.007)
log trade barrier	-0.575	-0.059	1.807	0.014	0.099	-0.013
	(0.091)§	(0.009)§	(0.268)§	(0.006)†	(0.022)§	(0.007)*
Num. observations	490	490	490	490	490	490
R^2	0.64	0.48	0.48	0.12	0.14	0.09

Notes: Panels (a) and (b) reports summary statistics for Mexico and Peru, respectively. Panels (c) and (d) report associations with destination characteristics for Mexico and Peru, respectively. Please see the note to table **??** for more details.

Figure C.1: Effects of tenure and duration on exporters' sales (Mexico and Peru)

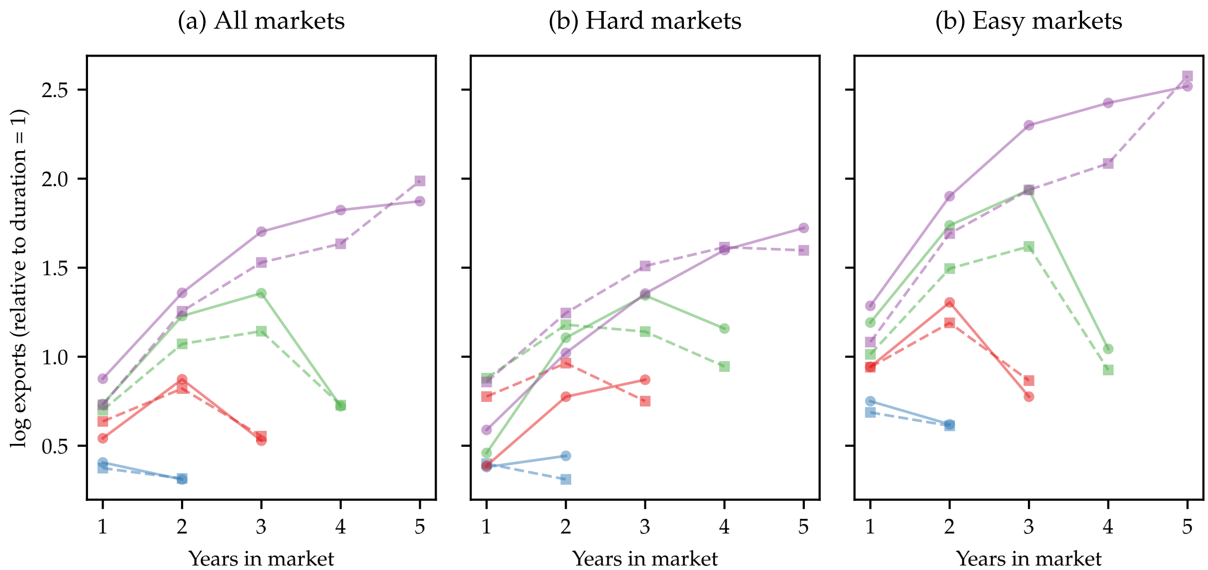

Notes: Panel (a) shows estimates of $\beta_{m,n}$ from (**??**). Panel (b) shows estimates of $\beta_{m,n,g}$ from (**??**) for markets in the bottom 50% of export participation, and panel (c) shows estimates for markets in the top 90%. Each line shows $\beta_{m,1}$, $\beta_{m,2}$, …, $\beta_{m,6}$ (or $\beta_{m,1,g}$, $\beta_{m,2,g}$,… in the second two panels) for a set value of m. Solid lines with round markers: Mexico. Dashes lines with square markers: Peru.

Figure C.2: Exit rates conditional on tenure (Mexico and Peru)

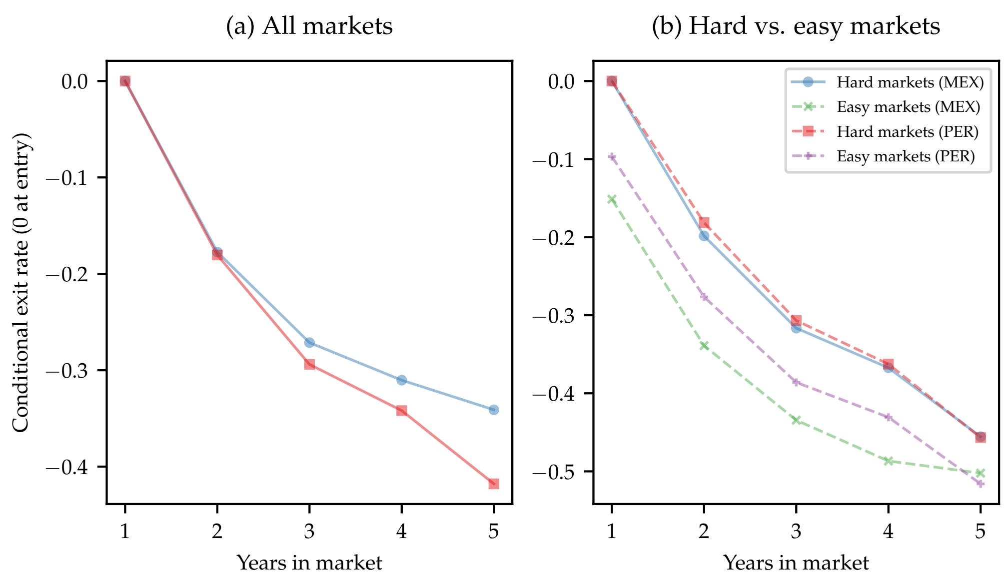

Notes: Panel (a) shows estimates of β_n from (**??**). Blue line with round markers: Mexico. Red line with square markers: Peru. Panel (b) shows estimates of $\beta_{n,g}$ from (**??**). Blue (red) line with round (square) markers: Markets in the bottom 50% of export participation for Mexico (Peru). Green (purple) line with 'x' ('+') markers: Markets in the top 90% of export participation for Mexico (Peru).