

NBER WORKING PAPER SERIES

MULTINATIONAL EXPANSION IN TIME AND SPACE

Stefania Garetto
Lindsay Oldenski
Natalia Ramondo

Working Paper 25804
<http://www.nber.org/papers/w25804>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2019

We thank our discussants Maggie Chen, Javier Cravino, Ana Maria Santacreu, Michael Sposi, and Stephen Yeaple for very helpful comments in more than one occasion. We have also benefited from comments from George Alessandria, Costas Arkolakis, Ariel Burstein, Jonathan Eaton, Oleg Itskhoki, Samuel Kortum, Andrei Levchenko, Eduardo Morales, Ezra Oberfield, and Andres Rodriguez-Clare, as well as seminar participants at various conferences and institutions. Xiao Ma provided outstanding research assistance. The statistical analysis of firm-level data on US multinational companies was conducted at the Bureau of Economic Analysis, US Department of Commerce, under arrangements that maintain legal confidentiality requirements. The views expressed are those of the authors and do not reflect official positions of the US Department of Commerce. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Stefania Garetto, Lindsay Oldenski, and Natalia Ramondo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Multinational Expansion in Time and Space
Stefania Garetto, Lindsay Oldenski, and Natalia Ramondo
NBER Working Paper No. 25804
May 2019
JEL No. F1,F23

ABSTRACT

This paper studies the expansion patterns of the multinational enterprise (MNE) in time and space. Using a long panel of US MNEs, we document that: MNE affiliates grow by exporting to new markets; the activities of MNE affiliates persist during the affiliate's life, usually starting with sales to their host market and eventually expanding to export markets; and MNE affiliates' entry into new locations does not depend on the location of preexisting affiliates. Informed by these facts, we develop a multi-country quantitative dynamic model of the MNE that features heterogeneity in firm-level productivity, persistent aggregate shocks, and a rich structure of costs that affect MNE expansion. Importantly, MNE affiliates can decouple their locations of production and sales, and endogenously choose to enter or exit the host and the export markets. We introduce a compound option formulation that allows us to capture in a tractable way the rich heterogeneity that is observed in the data and that is necessary for quantitative analysis. Using the calibrated model, our quantitative application to Brexit reveals that export platforms are important for understanding the reallocation of MNE activity in time and space, and that the nature of the frictions to MNE activities matters for aggregate firm dynamics.

Stefania Garetto
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
garettos@bu.edu

Natalia Ramondo
School of Int'l Relations and Pacific Studies
University of California at San Diego
9500 Gilman Drive, MC 0519
La Jolla, CA 92093
and NBER
nramondo@ucsd.edu

Lindsay Oldenski
Georgetown University
37th and O Streets, NW
Washington DC, 20057
Lindsay.Oldenski@georgetown.edu

1 Introduction

Many important questions in international economics involve the complex activities of multinational enterprises (MNEs) in time and space. Consider the recent rise in protectionism worldwide: the debate on the United Kingdom abandoning the European Union (EU), “Brexit”, is one example. Under Brexit, would MNEs pull out from the United Kingdom and reallocate toward other countries? Would MNEs located in the EU be affected? How would trade flows linked to MNEs change? Providing sound answers to these and other similar questions requires an understanding of the dynamic patterns of MNE expansion and the nature of the costs these firms face.

Despite their importance, the behavior of MNEs and their affiliates in time *and* space has received little attention in the literature.¹ On the empirical side, this is primarily due to data limitations. On the theoretical side, the nature of the costs of MNE activities—whether variable, fixed, or sunk, and whether host- or destination-country specific—poses challenges to tractability, particularly in a multi-country dynamic setting where MNEs can separate the locations of production and sales. This paper contributes to filling the gap in the literature by introducing a new multi-country dynamic model of the MNE which is informed by a new set of facts on the behavior of MNE affiliates. The quantitative model is aimed at answering counterfactual questions about the effects of MNE behavior in time and space.

Our analysis uses a long panel of US MNEs and their foreign affiliates from the Bureau of Economic Analysis (BEA). Studying the behavior of US MNEs and their foreign affiliates is a relevant setup not only because the United States is the main source of MNEs in the world, but also because MNE affiliates are the main channel through which US firms reach foreign consumers. In 2009, for instance, majority-owned affiliates of US MNEs abroad accounted for 75 percent of US sales to foreign customers; forty percent of those affiliates’ sales were exports, i.e., sales to customers outside the affiliate’s host market.

We start by documenting three facts about the dynamic behavior of US MNEs and their affiliates. First, MNE expansion happens mainly at the extensive, rather than intensive, margin. We observe that MNEs expand by entering new markets, either with a new affiliate or exporting from an existing one. We do not find evidence of growth at the intensive margin within a country: the ratio of affiliate-to-parent sales is flat over the affiliate’s life. Second, the activities of MNE affiliates persist over the affiliate’s life. The vast majority of affiliates are born specialized in sales to the host market, and local sales remain the affiliate’s main activity; affiliates may start exporting later in life. Third, the location of a new MNE affiliate does not depend on the location of preexisting

¹ See Antràs and Yeaple (2014) for a detailed survey on the main facts and theories about MNEs.

affiliates—in other words, the pattern of affiliates’ entry does not display “extended gravity”, in stark contrast with the facts that Morales et al. (2018) document for exporters.

Guided by these facts, we build a multi-country dynamic model of MNE expansion. Home-based firms decide whether, when, and where to open foreign affiliates. Affiliates, in turn, can sell both to their host market and to any other market. Affiliate operations, both in the host and in the export markets, are subject to sunk, fixed, and variable costs. The MNE decisions of whether to set up an affiliate in a market, and whether to export from it, are shaped by the interaction of firm-specific characteristics, persistent aggregate productivity and demand shocks, and the array of MNE costs.

Our model’s structure is based on two main assumptions. First, guided by the observation that almost all affiliates in the data have some horizontal sales at birth, we assume that firms that decide to do Foreign Direct Investment (FDI) must first set up an affiliate and sell to the local market, and only then can they consider exporting from that affiliate. Since the model is set up in continuous time, the decision of opening an affiliate and the decision of exporting from it can be made virtually simultaneously, generating affiliates with exclusively horizontal sales and affiliates with both horizontal and export sales at birth, as in the data. Second, consistent with the lack of extended gravity in the data, we assume that the decision of opening an affiliate—and eventually exporting from it—is independent across markets. For instance, a firm’s decision of opening an affiliate in Germany and exporting to France from it is independent from having an affiliate in France. We introduce interdependence in location choices into the model by allowing the decision to open an affiliate in a country to depend on the set of countries where the affiliate can export to. In this way, the problem of the MNE takes the form of a *compound option*: opening an affiliate in a country is an option which, when exercised, gives access to a set of additional options, such as exporting from the affiliate to any other location.

Our two assumptions, together with the compound-option structure, are key for achieving tractability of the model while, at the same time, preserving the rich heterogeneity necessary for quantitative analysis. Separating the decisions of performing affiliate activities by country significantly simplifies the dynamic problem of the MNE. Because of the continuous-time specification, value functions can be solved in closed form (as simple additive functions of the firm’s realized profit flow, the option value of further expansion, and the option value of exit). In turn, the independence assumption, when coupled with the compound-option structure, implies static independence, but dynamic interdependence, of location decisions. These features of our model avoid the permutational problem present in static models of export platforms; such problem would be extremely hard to solve in a multi-country dynamic setup. Furthermore, because of the tractability of the firm’s problem, we are able to aggregate individual firms’ outcomes and solve for the evolution of price indexes. In this

way, we can construct a measure of the welfare changes induced by changes in MNEs' activities.

While the static components of our model are standard and follow Melitz (2003), the way we formulate the dynamic problem is new to the international trade literature. We build on insights from the literature on real options to solve general models of investment under uncertainty (see Dixit and Pindyck, 1994). The application to the complex decisions of the MNE appears natural, since MNE location and export decisions happen over time and they are likely to be affected by uncertainty in demand and other market characteristics.²

We calibrate the model to static and dynamic moments related to the behavior of US MNE affiliates located in the top ten host countries for US FDI, over thirty years. Our calibration implies that opening—and operating—affiliates is more costly than exporting from them, for most host countries. Exports to the United States (the Home country) are generally associated with lower barriers than exports to other destinations. Heterogeneity, however, is large across host countries, sales type, and types of frictions.

The calibrated model is able to reproduce non-targeted observations related to the selection patterns of MNE affiliates within and across host markets. Concretely, the model matches fairly well the size advantage, in terms of horizontal sales, observed for: affiliates that export over affiliates that do not export; affiliates that export earlier in life over affiliates that export later; and affiliates that are opened first within the MNE over subsequent affiliates.

Armed with the calibrated model, we perform various counterfactual exercises with the goal of evaluating the importance of the different frictions to MNE expansion, of the endogenous responses of price indexes, and of the compound-option structure. Since our sample includes affiliates located in the United Kingdom, Ireland, Germany, and France, we use the potential withdrawal of the United Kingdom from the European Union (EU), Brexit, as our main counterfactual exercise.

Different potential implementations of Brexit are related to the increase of different types of export costs between the United Kingdom and other EU countries. Our model predicts that an increase in export costs between the United Kingdom and other EU countries would have a static effect, a dynamic effect, and an equilibrium price effect. First, export activities between the United Kingdom and the EU would become more costly, so that sales from UK-based affiliates to the EU, as well as sales from EU-based affiliates to the United Kingdom, would decline. This decline would drive the decrease in the incentive to open affiliates in the United Kingdom and in other EU countries, due to

² The compound option formulation can prove useful for problems related to global sourcing decisions of the MNE, which are likely to occur over time and under uncertainty. One can imagine a set up in which making an investment to source an input from a given country opens up the possibility of sourcing other, more upstream, inputs from a different location.

the smaller, and costlier, export market. Second, increases in trade costs would affect the affiliate export band of inaction, and hence, affiliate export entry and exit rates would change. Finally, increases in trade frictions would have the effect of raising prices not only in the United Kingdom, but also in the EU, encouraging more export entry from the United Kingdom into those markets. The strength of each effect on aggregate firm dynamics would vary depending on the nature of the shock to export costs. For instance, while increasing sunk export costs would increase both the sales and number of affiliates selling to the EU from the United Kingdom, increases in iceberg trade costs would drastically decrease both the number and sales of UK-based US affiliates to the EU.

This paper is related to the existing literature in several ways. First, most contributions in the literature have analyzed MNE expansion in space, but not in time. As it is evident in the static models in Tintelnot (2017), Fan (2017), Arkolakis et al. (2018), and Head and Mayer (2019), allowing firms to set up affiliates in countries that differ from the destinations of their sales can result in a complex problem when fixed costs of production are included. The sharp patterns that we document from observing affiliates over time help to simplify this problem by reducing the choice set of firms in a way that is consistent with the data.

Second, there is a small, but growing, literature that analyzes different aspects of the dynamic behavior of the MNE. Papers in this literature, however, limit the spatial dimension of the problem. In a model with firm-level shocks, Gumpert et al. (2018) focus on the life-cycle dynamics of the proximity-concentration tradeoff—i.e., sales of MNE affiliates are exclusively directed to their host market and act as an alternative to exports from the Home market—and assess the role of MNEs on new exporters’ dynamics, using rich firm-level data from various countries. Given the nature of their question, the analysis does not consider export platforms, and focuses on life-cycle, rather than aggregate, firm dynamics. Fillat and Garetto (2015) build a dynamic two-country model of the proximity-concentration tradeoff with aggregate shocks. Importantly, they introduce the idea that MNE activities can be treated as a real option that gets exercised once an affiliate is opened abroad.³ Fillat et al. (2015) extend this idea to a multi-country setup. Both papers focus on the link between the MNE expansion decisions and asset prices, and both assume that the activities of affiliates are restricted to their market of operation.⁴ Our model treats MNE activities as a compound, rather than a simple, option. In this way, we are able to preserve the tractability of the

³ Impullitti et al. (2013) also use a real option to model the entry and exit patterns of exporters.

⁴ Other related papers in the MNE literature are the following. Ramondo et al. (2013) study the implications of the proximity-concentration tradeoff under uncertainty and use BEA data to assess the predictions of the model. Conconi et al. (2016) couple a model of the proximity-concentration tradeoff with learning and test it using detailed data for Belgium. Focusing only on MNEs, Egger et al. (2014) propose a learning model to explain the pecking order of markets observed in the entry patterns of German MNEs. All these papers, however, limit the spatial dimension of their analysis by considering only horizontal FDI sales, as well as the dynamic dimension, by considering two-period models.

problem in a dynamic multi-country setup, and expand on the spatial dimension by separating the activities of MNE affiliates between the locations of production and sales.

Third, our paper is naturally related to the large literature on export dynamics, which has been primarily concerned with quantifying the various costs of export activities and their welfare implications.⁵ Complementing this literature, we quantify the frictions to MNE expansion, and analyze the implications in terms of aggregate firm dynamics and welfare. However, an important difference with the literature on export dynamics is that the nature of the MNE problem is more complex, and subject to more frictions, than the exporter problem. MNEs choose not only which markets to serve, as an exporter does, but also the location from which to serve each of those markets. Our compound-option structure, together with the assumptions on the timing and independence of entry decisions, allows us to solve the complex problem of the MNE in a dynamic multi-country setup.

Finally, our paper relates to the large literature that analyzes the dynamics of domestic firms, which goes back to Davis et al. (1996), and more recently Decker et al. (2014, 2016). Our facts suggest that the dynamics of MNE affiliates are starkly different from the dynamics of domestic firms. We interpret these differences as indicative of the fact that new US firms face a very different set of frictions in the domestic and foreign markets.

2 Establishing the Facts

We document three novel facts on the dynamic behavior of foreign affiliates of US multinational enterprises (MNEs). First, MNE expansion happens at the extensive, rather than the intensive, margin. Second, the vast majority of affiliates are born specialized in sales to the host market. This activity remains the main activity of the affiliate, while export activities may start later in life. Finally, the location of new affiliates does not depend on the location of preexisting affiliates.

⁵ Earlier contributions by Baldwin and Krugman (1989), Roberts and Tybout (1997), Das et al. (2007), and Alessandria and Choi (2007) find evidence of large sunk costs of exporting by focusing on observed patterns of export entry and exit. Subsequent analyses, such as Eaton et al. (2008) and Ruhl and Willis (2017), incorporate facts related to the life-cycle dynamics of new exporters and find that those costs are much lower. Alessandria et al. (2018) take a further step and also calculate the welfare gains from trade in a dynamic setting that matches well the life-cycle export facts. Arkolakis (2016) presents rich micro evidence on firm selection and export growth that supports dynamic theories of endogenous entry costs *vis-à-vis* standard export sunk costs. Finally, Fitzgerald et al. (2017), using detailed data on export prices and quantities, show that life-cycle growth of those two variables are quite different.

2.1 Data

Our empirical analysis uses firm-level data on the operations of US MNEs from the Bureau of Economic Analysis (BEA). The data include detailed information on the operations of MNEs in the United States and their affiliates abroad, for the period 1987-2011. We restrict the sample to majority-owned affiliates that do not operate in tax haven countries, have manufacturing as their primary activity, and belong to a US parent operating in any sector.⁶ We further consolidate affiliates belonging to the same parent and operating in the same country and 3-digit industry. Finally, for the facts presented in this section, we focus on affiliates that open during our sample period and that survive for at least ten consecutive years in the market. This restriction implies that we exclude affiliates that open in 2003 or later, as well as observations belonging to the affiliate's eleventh year of life, or greater. We also remove affiliates and parents with zero total sales.⁷

Crucially, the BEA data break down affiliate sales by destination: the host market of operation (horizontal sales), and other markets (exports). The data further distinguish between exports to the United States and to third markets. Every five years the BEA conducts a more detailed benchmark survey, which further distinguishes affiliate exports to Canada, the United Kingdom, and Japan.⁸

Table 1 shows the number of observations with positive horizontal and export sales in our sample. Almost 95 percent of our affiliate-year observations have some horizontal sales, while more than two thirds of them have some exports. More than one third of the observations correspond to affiliates with horizontal sales only, while the share of affiliates with only exports is around six percent. Since affiliates that only export are few and account for a small share of total affiliate sales, the model we present in Section 3 does not generate pure exporters.

Appendix A provides more details on the data coverage and sample construction.

2.2 The expansion of MNE affiliates

We start by documenting that MNE affiliates grow by entering new export markets.

Figure 1 shows the ratio of affiliate-to-parent sales, by affiliate age, for all, horizontal, and export

⁶ Our sample is primarily composed of affiliates that are majority owned during their whole life. Only about one percent of affiliates go from majority to minority owned and less than two percent go from minority to majority owned.

⁷ This restricted sample covers 23 percent of all affiliates in manufacturing as well as 38 percent of their total sales. Facts computed using a larger sample with a five-year survival threshold display the same patterns (not shown).

⁸ The distinction between the United States and other export markets of the affiliate does not make any substantial difference for the facts documented below. We do use the available country-specific affiliate export data in our calibration. It is also important to notice that the BEA data do not record parents' direct exports by destination.

Table 1: Number of observations, by sale type.

	Horizontal sales	Export sales
No. of observations	38,080	38,080
with positive sales	36,135 (95%)	25,958 (68%)
of pure type	14,035 (37%)	2,418 (6.3%)
Sales accounted by pure type	15.6%	7.7%


Note: Observations are at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. A pure-type affiliate is an affiliate for which at least 99 percent of sales are either only horizontal or only export sales.

sales. We plot the coefficients on age dummies from estimating, by Ordinary Least Squares (OLS),

$$\log Y_{iap} = \sum_{a=2}^{10} D_{iap} + \beta \log \text{global emp}_{iap} + u_{iap}. \quad (1)$$

The dependent variable Y_{iap} denotes the ratio of affiliate-to-parent sales for affiliate i belonging to parent p at age a , while D_{iap} is a set of age dummies. We include country-year and affiliate fixed effects in all the specifications as well as a control for the global size of the corporation in terms of employment, $\log \text{global emp}_{iap}$.

Figure 1: Affiliate-to-parent sales ratio, intensive margin.



Notes: OLS coefficients on age dummies (relative to the entry year) from estimating (1), with affiliate fixed effects and country-year fixed effects. Five-percent confidence intervals shown in dash lines. Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Each panel includes affiliates with positive sales in the given category.

The affiliate-to-parent sales ratio is flat: at all ages, this ratio is not significantly different from the ratio in the entry year. A similar lack of growth is observed not only for all affiliate sales, but also

for horizontal and export sales, separately.⁹

MNE affiliates do not grow at the intensive margin, but they do expand at the extensive margin, i.e., adding destinations other than the host market. Table 2 shows the results from estimating, by OLS,

$$\log Y_{iap} = \sum_{\delta \in \{-5, 1\} \cup \{1, 5\}} D_{i\delta p} + \beta \log \text{global emp}_{iap} + u_{iap}, \quad (2)$$

where $D_{i\delta p}$ equals one when affiliate i is δ years away from starting exporting. We also control for the global employment of the MNE, and include country-year and affiliate fixed effects.

Results in Table 2 are relative to sales in the year of entry into exports, the excluded category. In each of the five years preceding export entry (-5- to -1), the ratio of affiliate-to-parent sales is significantly lower than the ratio at entry. After export entry (1 to 5), this ratio is flat. The similarity of the coefficients for $\delta = \{-5, \dots, -1\}$ indicates that the affiliate-to-parent sales ratio increases only at the time of export entry. Expansion happens only at the extensive margin of sales destinations.¹⁰

Robustness. One could argue that the lack of growth in MNE affiliates’ sales may be due to the fact that the affiliate “inherits” the age of the parent so that, *de facto*, it is a much older firm, and hence, has lower growth rates. This may well be happening, as documented for multi- versus single-plant firms in the United States by Kueng et al. (2017).¹¹ Unfortunately, the BEA data do not record the age of the parent firm. However, we can look at the affiliate position in the opening sequence of the MNE—i.e., first affiliates versus subsequent affiliates. In this way, we can compare affiliates belonging to younger MNEs with affiliates belonging to older MNEs (or the same MNE at older ages). Columns 1 and 2 in Appendix Table B.1 show that first affiliates do not appear to grow faster than subsequent affiliates—even though the age-dummy coefficients for subsequent affiliates are more precisely estimated.

A second argument can be that, since global value chains (GVCs) have been growing very fast in the last decades, affiliates linked to them might be growing faster than non-GVC affiliates. In columns 3 and 4 in Appendix Table B.1, we show the results of estimating (1) separately for GVC

⁹ Notice that this finding is in stark contrast with the export dynamics literature, which documents small export shares at entry and intensive margin growth in exports over the life of surviving exporters. See, among others, Arkolakis (2016), Fitzgerald et al. (2017), and Ruhl and Willis (2017).

¹⁰ It is worth noting that MNE expansion does not happen by opening more than one affiliate in a given host market. Only 6.98 percent of US MNEs have more than one affiliate in the same host country; at the country-MNE level the share is 4.6 percent; and at the country-firm-year level the share is only 3.38 percent.

¹¹ Kueng et al. (2017) document a stark difference in the life-cycle employment profiles of establishments belonging to single- versus multi-unit firms in manufacturing: while establishments in single-unit firms grow steeply, the ones in multi-unit firms do not grow.

Table 2: Affiliate-to-parent sales ratio, extensive margin.

D(years to export entry = -5)	0.001 (0.010)
D(years to export entry = -4)	-0.023** (0.010)
D(years to export entry = -3)	-0.024* (0.013)
D(years to export entry = -2)	-0.018** (0.009)
D(years to export entry = -1)	-0.019*** (0.007)
D(years to export entry = 1)	-0.018* (0.009)
D(years to export entry = 2)	-0.014 (0.009)
D(years to export entry = 3)	-0.005 (0.010)
D(years to export entry = 4)	-0.003 (0.009)
D(years to export entry = 5)	-0.005 (0.008)
log global employment	-0.012 (0.009)
Obs	38,080
R^2	0.011

Note: Results from estimating (2) by OLS. Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. The dependent variable *affiliate-to-parent sales ratio* refers to affiliate sales relative to the domestic sales of the US parent. We include only affiliates that start exporting during our sample period. All specifications include affiliate and country-year fixed effects. Standard errors, clustered at the parent level, are in parenthesis. Levels of significance are denoted *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

affiliates (defined as affiliates with positive intra-firm exports), and for non-GVC affiliates (defined as affiliates with zero intra-firm exports). Both groups of affiliates have flat affiliate-to-parent sales ratios—even though estimates are more precise among GVC affiliates..

A third argument for the observed lack of growth in MNE affiliates’ sales is related to the mode of FDI entry. If MNEs establish foreign affiliates mostly through a merger with—or an acquisition of—an existing firm (M&A), one could argue that “new” foreign affiliates are in reality preexisting firms that likely grew before their acquisition. The BEA asks a subset of affiliates whether they were created through an M&A or a greenfield project.¹² Columns 5 and 6 in Appendix Table B.1 show that, relative to age two (the first year for which sales are recorded for the whole year), the sales ratio grows very little, regardless of the mode of FDI entry.

A final concern is that the flat sales ratio observed for affiliates may be due to the fact that firms grow in a foreign market first through exports, and only subsequently through opening an affiliate. Since the BEA data do not include information about parent exports by destination market, we are not able to address this question directly. Gumpert et al. (2018), however, report that, for Norway and France, the difference in growth profiles for MNEs with previous export experience into a market and those without it is not significant, except for the first year of the affiliate’s life.

The fact presented in this section motivates an important feature of the model we present below, namely that MNE affiliates grow at the extensive margin by entering new destination countries.


2.3 The activities of MNE affiliates over time

We now present evidence on the specialization patterns of affiliates in terms of horizontal and export activities over time. We show that affiliates are born specialized in horizontal sales, and eventually they incorporate exports as a secondary activity later in life.

Figure 2 shows the evolution of the intensive and extensive margins of horizontal and export sales as shares of total sales of the affiliate. Figure 2a shows the evolution of the horizontal and export sales share, computed as an average over affiliates reporting, respectively, positive horizontal and positive export sales. On average, horizontal sales account for about 80 percent of affiliate sales at birth and decrease by ten percentage points over the first ten years of life of the affiliate, while the export share is flat at 40 percent. To capture the extensive margin of horizontal and export sales, Figure 2b plots the percentage of affiliates with non-zero horizontal sales and non-zero export sales. While the share of affiliates with horizontal sales is stable at more than 95 percent, the share of

¹² This question applies only to firms that opened mid-year, and thus, the reported information about sales covers only part of the entry year. For this reason, we compute the ratio of affiliate-to-parent sales starting at age two.

Figure 2: Intensive and extensive margins of affiliate sales, by type.



Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Horizontal and export sales refer, respectively, to sales to the market where the affiliate is located, and to sales to markets outside the local market. (2a): average sales, as a share of total affiliate sales, include affiliates with positive horizontal and export sales, respectively. (2b): number of affiliates, as a share of the total number of affiliates, include affiliates with positive horizontal and export sales, respectively.

exporting affiliates increases from 50 to 70 percent during the first ten years of life of the affiliate. In other words, for horizontal activities, changes in sales shares are due to the intensive margin, while export shares increase only because of affiliates that start exporting. Hence, the data suggest that, over time, affiliates incorporate export sales into their activities, but they never stop selling to their host market.

The patterns in Figure 2 are confirmed by OLS regressions that include a battery of fixed effects, as shown in Appendix Table B.2. Estimates that include affiliate fixed effects suggest that, on average, horizontal (export) sales shares decrease (increase) during the life of an affiliate, and the share of affiliates with exports is higher among older affiliates.

For robustness, Appendix Figure B.1 and Table B.3 report analogous results for the subset of affiliates that are pure type at birth—i.e., firms that in their first year of life either sell exclusively to the host market or only export. The results on pure-type affiliates reinforce the patterns shown in Figure 2: pure-type affiliates diversify activities over their life, moving from exclusively horizontal sales to also exporting. Relatedly, Appendix Figure B.2 shows the relationship between the intensity at which an activity is performed and the time at which the affiliate first starts that activity: the older the affiliate is when it starts exporting, the lower its export intensity.

The fact documented in Figure 2 motivates an important assumption of our model: all affiliates

start operations with some horizontal sales and may endogenously expand into export markets.

2.4 Geography and the entry pattern of MNE affiliates

We now present a fact related to the sequence of location decisions of the MNE. In particular, we document the lack of “extended gravity”: the location of a new affiliate does not depend on the location of preexisting affiliates, either in terms of geographic distance or of other measures of similarity between markets.

Table 3 shows that, for a given US parent, the unconditional probability of opening an affiliate in a country is very similar to—and in several cases not significantly different from—the probability of opening an affiliate conditional on already having an affiliate in a “similar” country. Following Morales et al. (2018), we define “similarity” in a variety of ways: similar countries may be located in the same continent, share a border, share a language, have similar income per capita, or all of the above.¹³ Of course, this comparison is only possible for US parents with at least two foreign affiliates. We further restrict the sample to affiliates located in the ten most popular host countries for US MNEs, which cover about 60 percent of all sales of majority-owned US affiliates abroad. The table shows that the largest differences in conditional and unconditional probabilities are observed for countries, such as China, that are typically part of global supply chains.¹⁴ Differences between unconditional and conditional probabilities become smaller when we restrict the sample to MNEs with more than five, and more than ten, affiliates (see Appendix Table B.5).

Our findings on the lack of extended gravity for MNE affiliate entry are in stark contrast to the findings for exporter entry. Morales et al. (2018) find that the probability of exporting to a given country is 12.7, 5.3, 3.0, and 2.9 *times* higher if the firm already exports to an adjacent country, to a country in the same continent, to a country sharing the same language, and to a similar country in terms of per capita income. In contrast, we find that the probability of opening an affiliate in the United Kingdom, for example, is 17.5, 8.7, 2.8, and 2.4 *percent* higher if the firm already has an affiliate in an adjacent country, a country in the same continent, a country sharing the same language, and a country in the same income group.

The observed lack of extended gravity in affiliate entry motivates a key assumption of our model: the choice of opening an affiliate in a country is independent from the location of preexisting

¹³ We use the World Bank classification and divide countries into four groups according to their GDP per capita. We refer to two countries belonging to the same income group as similar in terms of income per capita.

¹⁴ In fact, Appendix Table B.4 shows that the lack of extended gravity is much more pronounced among non-GVC affiliates (i.e., affiliates with zero intra-firm exports).

Table 3: Unconditional and conditional probability of affiliate entry.

	Unconditional	Continent	Border	Language	Income	All
	(1)	(2)	(3)	(4)	(5)	(6)
Canada	0.021	0.021	–	0.023	0.021	–
		(0.525)	–	(0.000)	(0.553)	–
United Kingdom	0.025	0.027	0.030	0.026	0.026	0.030
		(0.000)	(0.143)	(0.292)	(0.008)	(0.143)
Germany	0.023	0.026	0.029	0.028	0.024	0.028
		(0.000)	(0.000)	(0.010)	(0.000)	(0.010)
Ireland	0.010	0.010	0.011	0.010	0.010	0.011
		(0.001)	(0.010)	(0.000)	(0.005)	(0.011)
China	0.027	0.037	0.050	0.048	0.051	0.057
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
France	0.021	0.024	0.028	0.023	0.022	0.029
		(0.000)	(0.000)	(0.018)	(0.000)	(0.000)
Brazil	0.016	0.022	0.027	0.025	0.023	0.019
		(0.000)	(0.000)	(0.063)	(0.000)	(0.614)
Singapore	0.016	0.023	0.044	0.017	0.016	0.045
		(0.000)	(0.000)	(0.000)	(0.300)	(0.000)
Mexico	0.024	0.029	0.028	0.034	0.031	0.024
		(0.000)	(0.620)	(0.000)	(0.000)	(0.961)
Japan	0.016	0.021	–	–	0.016	–
		(0.000)	(0.000)	(0.000)	(0.224)	(0.000)

Note: Probabilities of affiliates' entry into the top-ten most popular destinations of US MNEs. Conditional probabilities refer to the probability of observing a MNE opening an affiliate in country i given that the parent already has an affiliate in a "similar" country. Column 6 refers to similarity in all the dimensions in columns 2-5. The sample is restricted to parents with at least two affiliates worldwide. P-values from tests of equality of the conditional and unconditional probabilities are in parentheses. Conditional probabilities in **bold** are not significantly different from the relevant unconditional probability.

affiliates of the same parent.

Relatedly, at the affiliate-export level, we find suggestive evidence that an affiliate’s decision to enter an export market is independent from whether its parent already has an affiliate in that country. With the BEA data, for the benchmark year 2004, we are able to examine the coexistence of affiliates’ exports to three countries (Canada, the United Kingdom, and Japan) with the presence of affiliates owned by the same parent in those countries. Of the 20,359 affiliates that export to Canada, 64 percent belong to a US parent that also has affiliates located in Canada. Similarly, of the 5,017 affiliates that export to the United Kingdom, 70 percent belong to a US parent that also has affiliates located in that country. Finally, of the 5,224 affiliates that export to Japan, 47 percent belong to a US parent that also has affiliates located in Japan.

This finding motivates an independence assumption on export entry in the model: an affiliate’s decision to export to a destination does not depend on the location of other affiliates of the same parent.

3 A Dynamic Model of MNE Expansion

We build a quantitative dynamic model in which MNEs open affiliates across countries over time. Affiliates sell in their host markets, and they choose whether to export to other markets from there. We impose assumptions that are guided by the facts documented in Section 2 and are key for the tractability of the model.

While the static components of our model are standard and follow Melitz (2003), the dynamic part of the model is based on a compound option formulation that allows us to capture the richness of the decisions of the MNE in time and space. This formulation is novel to the international trade literature, and it is the key element that makes the model amenable to quantitative analysis.

3.1 Preferences and technology

The economy consists of $N + 1$ countries: the Home country (the United States in our data) and N foreign countries. Time is continuous. In each country k , consumers have preferences over a composite good,

$$U_k = \int_0^\infty e^{-\rho t} Q_k(t) dt, \quad (3)$$

with ρ denoting the subjective time discount rate. The quantity $Q_k(t)$ aggregates a continuum of varieties, indexed by v , with a constant elasticity of substitution (CES) $\eta > 1$,

$$Q_k(t) = \left[\sum_i \sum_j \int_{\Omega_{ijk}(t)} \lambda_{ijk}^{\frac{1}{\eta}} q_{ijk}(v, t)^{\frac{\eta-1}{\eta}} dv \right]^{\frac{\eta}{\eta-1}}. \quad (4)$$

The quantity $q_{ijk}(v, t)$ denotes consumption of variety $v \in \Omega_{ijk}(t)$, and $\Omega_{ijk}(t)$ denotes the set of varieties sold to country k and produced by affiliates located in j belonging to firms from i , at time t . The term λ_{ijk} denotes a preference shifter.

Assumption 1. *Varieties consumed and produced are origin-location-destination specific.*

Assumption 1 implies that consumers perceive differently varieties produced in different locations by the same firm. For example, consumers in a given destination perceive Mœt Chandon champagne produced in France as different from Chandon sparkling wine produced by the same firm in California. Importantly, this assumption implies that there is no cannibalization of sales when a MNE serves a country by opening an affiliate there *and* also by exporting to it from an affiliate in a different location.¹⁵

Each country is populated by a continuum of firms. The Home country is the only source of MNEs: Home firms decide whether to operate only in their home market or to also establish affiliates abroad. For this reason, to simplify notation, we remove the index i that denotes a variety's origin country and use the subscript d to refer to the parent's operations at Home.

Labor is the only factor of production. Each firm produces with a linear technology and operates under monopolistic competition. As in Melitz (2003), each firm is characterized by a productivity parameter φ that determines the unit labor cost of the good produced. Each firm sets prices to maximize profits from sales to each destination, $p_{jk}(\varphi) = \tilde{\eta} c_{jk}(\varphi)$, with $\tilde{\eta} \equiv \eta/(\eta - 1)$ and $c_{jk}(\varphi) \equiv w_j \tau_{jk}/\varphi$. The term w_j denotes the wage in country j where production takes place, and τ_{jk} denotes the iceberg cost of shipping goods from production location j to destination k , with $\tau_{jk} \geq 1$, $\forall j \neq k$, and $\tau_{jj} = 1$, $\forall j$.

A firm's domestic profits are then given by $\pi_d(\varphi) = H(w_d/\varphi)^{1-\eta} P_d^\eta \lambda_d Q_d$, and variable profits from sales to k of an affiliate in j are given by $\pi_{jk}(\varphi) = H(\tau_{jk} w_j/\varphi)^{1-\eta} \lambda_{jk} P_k^\eta Q_k$, where $H \equiv \eta^{-\eta} (\eta-1)^{\eta-1}$ and P_k is the corresponding CES price index. Note that, for $j = k$, the variable π_{jk} denotes profits from horizontal sales, while for $j \neq k$, it denotes profits from affiliate export sales.

¹⁵ Assumption 1, together with the preference shifter λ_{ijk} , allow us to identify—and match—location-destination specific sales shares as observed in the data.

When a firm establishes an affiliate in a foreign country j , it has to pay a sunk entry cost $F_j^h > 0$. The affiliate starts by selling locally and incurs a per-period fixed cost $f_j^h > 0$. Once the affiliate is in place, it can expand its operations to export to other markets. An affiliate located in country j has to pay a sunk cost $F_{jk}^e > 0$ to start exporting to country k , and a per-period fixed export cost $f_{jk}^e > 0$. For simplicity, we assume that there are no per-period fixed costs associated with domestic production, so that all firms produce at Home.

3.2 The MNE dynamic problem: the compound option

We now present the MNE dynamic problem. At each point in time, a firm endogenously decides whether to open an affiliate in a foreign country, and whether—and where—to export from its existing affiliates, including exporting to the Home market. A firm may also decide to shut down affiliates, or to exit any of its affiliate export markets.

We use the notion of a *compound option* to model the dynamic problem of the MNE. Opening an affiliate in a country is an option that, when exercised, gives access to another set of options, namely the possibility of expanding to each export destination. Hence, the decision to open an affiliate in country j depends on the set of countries where the affiliate can export to. The compound-option structure allows us to easily solve the firm's problem backwards, as suggested by Dixit and Pindyck (1994, chap. 10). Conditional on the MNE having an affiliate in country j , one can solve for the value of exports to each destination and for the policy functions that induce the affiliate to start, or stop, exporting to each country $k \neq j$. Together with the value of horizontal sales, this determines the value of an affiliate in country j . One can then solve for the policy functions that induce the firm to open, or shut down, the affiliate.

The two assumptions we present next, together with the compound option formulation, lend tractability to the model and allow us to make headway toward a rich quantitative analysis. Both assumptions are guided by the empirical observations in Section 2.

Assumption 2. *The decision to open an affiliate in a country, and export from there, is independent from the decision to open affiliates in any other country.*

This assumption is equivalent to imposing the condition that the profits of an affiliate in a country are independent from the number of affiliates within the MNE. Together with Assumption 1, it implies that the same MNE can serve a market through affiliates operating there and through exports from affiliates located elsewhere.

Assumption 3. *A new affiliate starts selling in its host market before eventually starting to export.*

Assuming sequential decisions for the affiliate is a mere artifact to gain tractability: because the model is specified in continuous time, opening an affiliate and exporting from it can happen almost simultaneously. In this way, the model can generate affiliates that export from birth, as observed in the data.

Shock structure. Following Ghironi and Melitz (2005), we define the firm-level productivity φ as the product of a time-invariant firm-specific component, z , and a time-varying Home-country specific component, Z , so that $\varphi \equiv z \cdot Z$. The term z is a firm-specific draw from a distribution $G(z)$ which is constant over time, as in Melitz (2003). We assume that $Z = e^X$, where X is a Brownian motion with drift,

$$dX = \mu dt + \sigma dW, \quad (5)$$

for $\mu \in \mathfrak{R}$, $\sigma > 0$, and dW denoting a standard Wiener process. The specification in (5) is equivalent to assuming that aggregate Home productivity behaves like a random walk and that productivity growth is independently and identically distributed. This is a convenient functional form assumption that guarantees the tractability of the model's solution.

We assume that when a firm operates an affiliate in a foreign country, it transfers both the aggregate and the idiosyncratic components of the productivity shock to the host market. In this way, MNE operations contribute to the transmission of productivity shocks across countries, in the spirit of Cravino and Levchenko (2017).

We also introduce host-country aggregate demand shocks by assuming that aggregate demand in destination country k evolves according to a geometric Brownian motion,

$$dQ_k = \mu_k Q_k dt + \sigma_k Q_k dW_k, \quad (6)$$

where $\mu_k \in \mathfrak{R}$, $\sigma_k > 0$, and dW_k denotes a standard Wiener process, possibly correlated with the Home aggregate productivity shock.

Affiliate profits from sales to country k are linear in the term $e^{(\eta-1)X} Q_k$. Thus, it is convenient to define the “composite” shock $Y_k \equiv e^{(\eta-1)X} Q_k$, which captures the effect of both source- and destination-country aggregate shocks on affiliates' profits. The composite shock Y_k is also a geometric Brownian motion with drift $\tilde{\mu}_k$ and variance $\tilde{\sigma}_k^2$ given by

$$\tilde{\mu}_k = \mu_k + \mu(\eta - 1) + \frac{\sigma^2}{2}(\eta - 1)^2 + \gamma_k \sigma_k \sigma, \quad (7)$$

$$\tilde{\sigma}_k^2 = \sigma_k^2 + \sigma^2(\eta - 1)^2 + 2\gamma_k(\eta - 1)\sigma\sigma_k, \quad (8)$$

where γ_k denotes the correlation between e^X and Q_k .

We show below that the model can be solved in terms of realizations of the composite shock. Together with the independence assumption, this shock structure ensures that only Home- and destination-country shocks are important for the firm's dynamic decision to serve a market.

The shock structure we assume in the model is based on both observations in the data and analytical convenience. Empirically, Appendix Table B.6 shows that most of the variation in MNE affiliates' sales is explained by country-specific time-varying shocks and parent fixed effects, rather than parent- and affiliate-level time-varying shocks. In the calibration, the introduction of country-specific demand shocks allows us to match the evolution of affiliate shares in different host countries. Computationally, relying only on aggregate shocks makes feasible the aggregation of individual firms' decisions and the computation of equilibrium price indexes for many countries.¹⁶

Finally, the assumption that the MNE transfers its productivity to the affiliates abroad, together with the presence of country-level shocks, implies that, conditional on entry, a firm's Home and foreign sales perfectly co-move, as appears to be the case in the data (see Figure 1). The specification of the aggregate shock as a unit root process drives persistence in the affiliate's activities. Together with aggregate productivity growing over time ($\mu \geq 0$), this persistence gives rise to the dynamic patterns documented in Figure 2: affiliates start serving their host market, and later on, they start expanding internationally. The inclusion of iceberg trade costs further implies that export activities are performed at lower intensity.

Bellman equations. The state of the economy is described by the $(N + 1)$ -tuple (X, \mathbf{Q}) , where $\mathbf{Q} = [Q_1, \dots, Q_N]$. Let $\mathcal{V}(z, X, \mathbf{Q})$ denote the expected net present value of a Home-country firm with productivity z that follows an optimal policy when the state of the economy is (X, \mathbf{Q}) . Thanks to the independence assumption, we can write such value as

$$\mathcal{V}(z, X, \mathbf{Q}) = V_d(z, X) + \sum_{j=1}^N \max \{V_j^o(z, X, \mathbf{Q}), V_j^a(z, X, \mathbf{Q})\}. \quad (9)$$

The function $V_d(z, X)$ is the value of domestic operations, $V_j^o(z, X, \mathbf{Q})$ is the option value of opening an affiliate in country j , and $V_j^a(z, X, \mathbf{Q})$ is the value of an affiliate in country j , regardless of the destination of its sales. In turn, the value of an affiliate in country j is given by

$$V_j^a(z, X, \mathbf{Q}) = V_j^h(z, X, \mathbf{Q}) + \sum_{k \neq j} \max \{V_{jk}^o(z, X, \mathbf{Q}), V_{jk}^e(z, X, \mathbf{Q})\}. \quad (10)$$

The function $V_j^h(z, X, \mathbf{Q})$ is the value of horizontal sales in country j , $V_{jk}^o(z, X, \mathbf{Q})$ is the option

¹⁶ By relying on aggregate shocks only, we do not need to keep track of changes in the firms' productivity distribution over time, significantly reducing the dimensionality of the state space.

value of exporting to country k for an affiliate located in j , and $V_{jk}^e(z, X, \mathbf{Q})$ is the value of exports to country k for an affiliate located in j . In this way, the problem is formulated as a compound option because opening an affiliate in a country is equivalent to exercising an option that gives access to another set of options: the options to export to any other country.

Since all firms operate in the domestic market, the value of domestic operations is simply given by the evolution of domestic profits over time, and depends only on the domestic shock X . Over a generic time interval Δt ,

$$V_d(z, X) = \frac{1}{1 + \rho\Delta t} [\pi_d(z, X)\Delta t + E[V_d(z, X')|X]], \quad (11)$$

where X' denotes the realization of Home aggregate productivity next period.

If a domestic firm has not yet opened an affiliate in country j , all the value from its operations in j is option value—i.e., the value of the possibility of entering j in the future,

$$V_j^o(z, X, \mathbf{Q}) = \max \left\{ \frac{1}{1 + \rho\Delta t} E[V_j^o(z, X', \mathbf{Q}')|(X, \mathbf{Q})]; V_j^a(z, X, \mathbf{Q}) - F_j^h \right\}, \quad (12)$$

where \mathbf{Q}' denotes the vector of realizations of demand shocks next period. This equation captures the fact that a firm may keep the option of entering market j , or may enter country j by opening an affiliate there, in which case it pays the entry cost F_j^h and gets the value of having an affiliate in country j , $V_j^a(z, X, \mathbf{Q})$.

Since we assume that all affiliates sell in the market where they are located, the value of horizontal sales for an affiliate in country j is given by

$$V_j^h(z, X, \mathbf{Q}) = \max \left\{ \frac{1}{1 + \rho\Delta t} [(\pi_{jj}(z, X, \mathbf{Q}) - f_j^h)\Delta t + E[V_j^h(z, X', \mathbf{Q}')|X, \mathbf{Q}]]; V_j^o(z, X, \mathbf{Q}) \right\}. \quad (13)$$

This equation captures the fact that the affiliate may survive and make profits from horizontal sales in j , or may shut down, in which case it gets the value of the option of opening an affiliate in j , $V_j^o(z, X, \mathbf{Q})$.

As indicated in (10), the value of an affiliate is given by the value of its horizontal plus its export sales. The Bellman equation describing the value of the option to export to country k for a firm with an affiliate in country j is given by

$$V_{jk}^o(z, X, \mathbf{Q}) = \max \left\{ \frac{1}{1 + \rho\Delta t} E[V_{jk}^o(z, X', \mathbf{Q}')|(X, \mathbf{Q})]; V_{jk}^e(z, X, \mathbf{Q}) - F_{jk}^e \right\}. \quad (14)$$

This equation captures the fact that the affiliate may keep the option of exporting to country k —and get the continuation value of that option—or may start exporting to country k , in which case it pays the entry cost F_{jk}^e and gets the value of exporting to k from j , $V_{jk}^e(z, X, \mathbf{Q})$. In turn, this value is given by

$$V_{jk}^e(z, X, \mathbf{Q}) = \max \left\{ \frac{1}{1 + \rho \Delta t} [(\pi_{jk}(z, X, \mathbf{Q}) - f_{jk}^e) \Delta t + E[V_{jk}^e(z, X', \mathbf{Q}') | (X, \mathbf{Q})]] ; V_{jk}^o(z, X, \mathbf{Q}) \right\}. \quad (15)$$

This expression captures the fact that the affiliate may keep exporting to country k —and get the continuation value of that option—or may stop exporting to country k , in which case it gets the value of the option of exporting to k from j , $V_{jk}^o(z, X, \mathbf{Q})$.

Value functions. The problem can be solved backwards by first solving for $V_{jk}^o(z, X, \mathbf{Q})$ and $V_{jk}^e(z, X, \mathbf{Q})$, conditional on the firm having an affiliate in country j . Once the location of the affiliate is determined, the only country j -specific variables in the firm's problem are wages and iceberg trade costs, which are time invariant. For this reason, thanks to the independence assumption, the value functions only depend on the Home productivity shock and on the demand shock in destination country k . Since these shocks enter the profit functions linearly, we can replace them with the composite shock $Y_j \equiv e^{(\eta-1)X} Q_j$.

Solving for the value of exports conditional on the affiliate's location is a simple case of interlinked options (see Dixit and Pindyck 1994, ch. 7), with solution given by

$$V_{jk}^o(z, Y_k) = B_{jk}^o(z) Y_k^{\beta_k}, \quad (16)$$

$$V_{jk}^e(z, Y_k) = \frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f_{jk}^e}{\rho} + A_{jk}^e(z) Y_k^{\alpha_k}. \quad (17)$$

The variables $B_{jk}^o(z) > 0$ and $A_{jk}^e(z) > 0$ are firm-specific parameters, while $\alpha_k < 0$ and $\beta_k > 1$ are the roots of $\tilde{\sigma}_k^2 \xi^2 / 2 + (\tilde{\mu}_k - \tilde{\sigma}_k^2 / 2) \xi - \rho = 0$. The term $B_{jk}^o(z) Y_k^{\beta_k}$ in (16) represents the option value of exporting to country k and is increasing in the realization of the composite shock. Similarly, $A_{jk}^e(z) Y_k^{\alpha_k}$ in (17) is the option value of quitting export market k and is decreasing in the realization of the composite shock—i.e., the option of exiting an export market has a larger value in “bad times”. For each country pair (j, k) and for each firm with productivity z , the parameters $B_{jk}^o(z) > 0$, $A_{jk}^e(z) > 0$, and the thresholds for the realizations of the composite shock that induce the affiliate to start and stop exporting—i.e., the policy functions—can be recovered from the appropriate system of value-matching and smooth-pasting conditions.

Following a similar procedure, one can show that the value of horizontal sales, conditional on having

an affiliate in country j , is given by the present discounted value of profits from horizontal sales plus the option value of shutting down the affiliate,

$$V_j^h(z, Y_j) = \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f_j^h}{\rho} + A_j^h(z)Y_j^{\alpha_j}, \quad (18)$$

where $A_j^h(z) > 0$ is a firm-specific parameter. As a result, the value of an affiliate in country j can be written as

$$\begin{aligned} V_j^a(z, \mathbf{Y}) &= A_j^h(z)Y_j^{\alpha_j} + \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f_j^h}{\rho} + \dots \\ &\quad \sum_{k \in \mathcal{A}_j(z)} \left[\frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f_{jk}^e}{\rho} + A_{jk}^e(z)Y_k^{\alpha_k} \right] + \dots \\ &\quad \sum_{k \notin \mathcal{A}_j(z)} \left[B_{jk}^o(z)Y_k^{\beta_k} \right], \end{aligned} \quad (19)$$

where $\mathcal{A}_j(z)$ is the subset of countries where an affiliate of firm z located in j exports to, and $\mathbf{Y} = [Y_1, \dots, Y_N]$. The implications of the independence assumption are clearly captured by (19): the value of an affiliate does not depend on the sales, or on the value of the firm's affiliates in other countries; it does depend, however, on the set of export destinations available from the affiliate's host country.

It remains to solve for the decision of a firm to set up an affiliate in country j . The option value of opening an affiliate in j is

$$V_j^o(z, Y_j) = B_j^o(z)Y_j^{\beta_j}. \quad (20)$$

Hence, for each host country j and for each firm with productivity z , the parameters $B_j^o(z) > 0$, $A_j^h(z) > 0$, and the thresholds of the realizations of the composite shock that induce the firm to open and shut down an affiliate can be recovered from the appropriate system of value-matching and smooth-pasting conditions.

Lastly, the value of domestic sales is simply given by the present discounted value of profits from domestic sales,

$$V_d(z, X) = \frac{\pi_d(z, X)}{\rho - \hat{\mu}}. \quad (21)$$

Details on the solution of the dynamic problem of the firm are shown in Appendix C.

3.3 Equilibrium price indexes

Thanks to the tractability of our multi-country model, we are able to solve for the price index in each country. This calculation entails keeping track of the evolution of the mass of affiliates located in each host country j and serving each destination country k . Appendix C shows the price indexes for each country j , and the law of motion for the mass of MNEs in each country j as well as the mass of firms in j that export to a destination k .

The ability to solve for equilibrium price indexes derives from the choices we made in terms of the setup of the model and shock structure. Traditionally, general equilibrium models of trade dynamics feature firm-level shocks but do not feature sunk costs (see, for example, Luttmer, 2007 and Arkolakis, 2016). Existing dynamic models with sunk costs characterize the equilibrium dynamics for a single firm, as in Das et al. (2007) and Morales et al. (2018), or focus on stationary equilibria where aggregate variables do not change over time, as in Alessandria and Choi (2007). These models are usually formulated in discrete time settings where the firm's value function itself needs to be solved numerically. Our continuous time formulation, coupled with unit root shocks, allows us to solve for the value functions in closed form (up to some constants). By including only aggregate shocks, we can easily solve for the price indexes as we do not need to keep track of the evolution of the firm's productivity distribution.

3.4 Model's implications

In this section, we derive some analytical properties of the model regarding the relationship between firm-level productivity, host market characteristics, and affiliate entry and export entry thresholds. In order to show these results analytically, we assume that the fixed costs of affiliate operations are “small”, so that there is no endogenous exit of affiliates, either from export markets or from their production locations. With this assumption, the option-value terms $A_h^e(z)$ and $A_{jk}^e(z)$, in (17), (18), and (19), are zero. Hence, we can obtain closed-form solutions for the affiliate entry and export entry thresholds,

$$Y_j^{OH}(z) = \left(\frac{\beta_j}{\beta_j - 1} \right) \cdot \left(\frac{f_j^h + \rho F_j^h}{\rho} - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) \right) \cdot \left(\frac{\rho - \tilde{\mu}_j}{\kappa_{jj}(z)} \right), \quad (22)$$

$$Y_{jk}^{OE}(z) = \left(\frac{\beta_k}{\beta_k - 1} \right) \cdot \left(\frac{f_{jk}^e + \rho F_{jk}^e}{\rho} \right) \cdot \left(\frac{\rho - \tilde{\mu}_k}{\kappa_{jk}(z)} \right). \quad (23)$$

The term $\kappa_{jk}(z) \equiv H(\tau_{jk}w_j/z)^{1-\eta}P_{jk}\lambda_{jk}$ is a firm-specific revenue term, and $\mathbf{V}_j^E(z, \mathbf{Y}_{-j})$ denotes the total value of exporting from an affiliate in j for a firm with productivity z .¹⁷ Details on the derivation of (22) and (23) are in Appendix C.

Proposition 1. *For a given host-destination country pair, more productive firms have lower affiliate entry thresholds and lower affiliate export entry thresholds: $\partial Y_j^{OH}(z)/\partial z \leq 0$ and $\partial Y_{jk}^{OE}(z)/\partial z \leq 0$.*

Proof. See Appendix C.

Under the assumption that, $\forall k \neq j$, $Y_j^{OH}(z) < Y_{jk}^{OE}(z)$ (the threshold for the shock realization that induces a firm to open an affiliate is lower than the one that induces the affiliate to export), Proposition 1 implies that: 1) affiliates that are exporters from birth have larger horizontal sales than affiliates born with exclusively horizontal sales; and 2) conditional on Home aggregate productivity—or host-country aggregate demand—increasing over time ($\tilde{\mu} \geq 0$), affiliates that start exporting later in life have lower horizontal sales than affiliates that start exporting earlier in life.

The upper panels of Figure 3 illustrate these predictions. The red and blue lines denote, respectively, the threshold for opening an affiliate in j , $Y_j^{OH}(z)$, and the threshold for starting exports from j to k , $Y_{jk}^{OE}(z)$. They are decreasing functions of the firm's productivity z , and hence, they are invertible functions. In Figure 3a, we assume that the realization of the aggregate shock is Y' and that we observe two firms having affiliates in the same host country j . Firm 1 with productivity z_1 has an affiliate in j with only horizontal sales, while firm 2 with productivity z_2 has an affiliate in j that also exports, so that $z_2 \geq z_j^e(Y') \geq z_1$. Since $z_2 \geq z_1$, the horizontal sales of the affiliate belonging to firm 2 must be larger than the horizontal sales of the affiliate belonging to firm 1. Now, suppose that the realization of the composite shock increases to $Y'' > Y'$. As illustrated in Figure 3b, $z_1 \geq z_j^e(Y'')$ and firm 1 will start exporting from its foreign affiliate in j . Hence, within a host country, affiliates that export earlier in life are more productive and exhibit larger horizontal sales than affiliates that start exporting later.

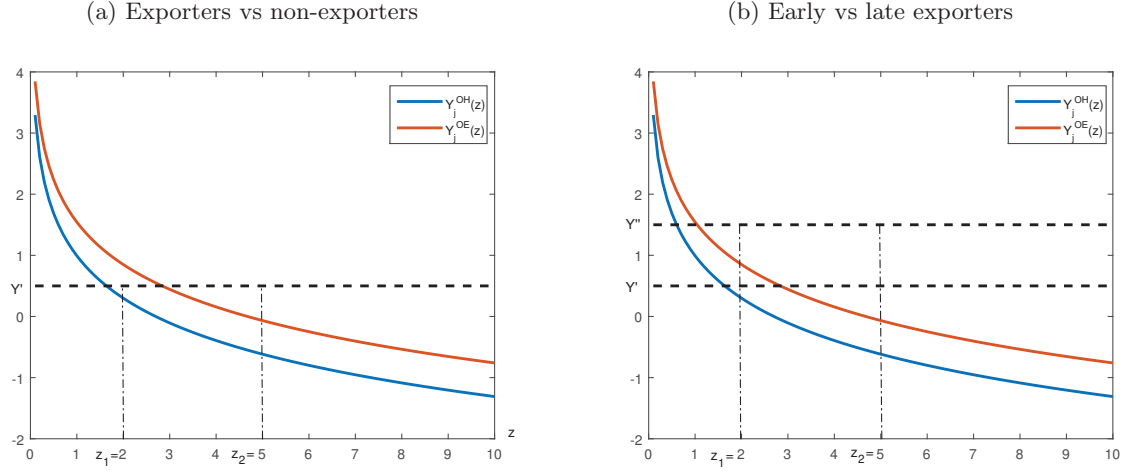
Proposition 2. *For a given firm with productivity z , the affiliate entry threshold is decreasing in the host-market preference shifter, $\partial Y_j^{OH}(z)/\partial \lambda_{jj} \leq 0$, and increasing in the entry cost, $\partial Y_j^{OH}(z)/\partial F_j^h \geq 0$.*

Proof. See Appendix C.

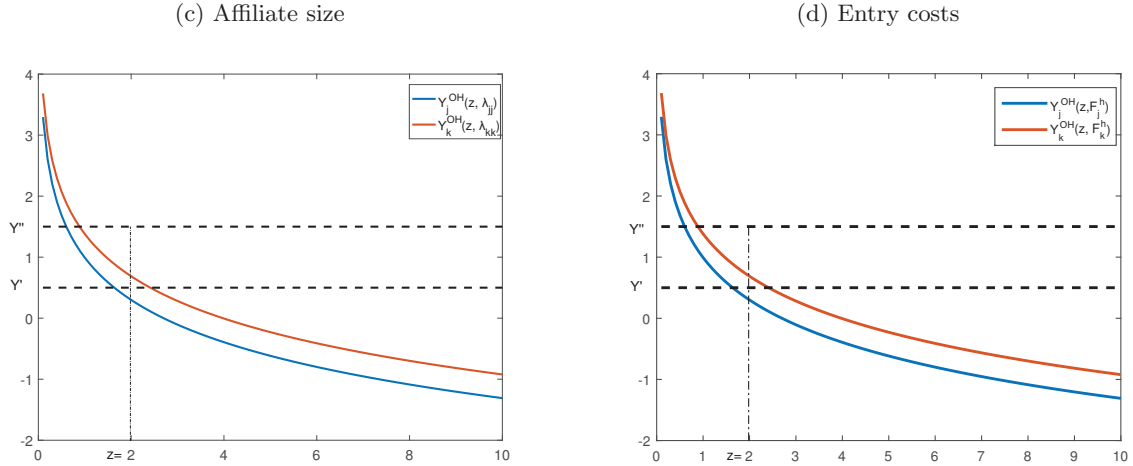
¹⁷ Notice that if $(f_j^h + \rho F_j^h)/\rho - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) < 0$, then $Y_j^{OH}(z) < 0$. In this case, a firm with productivity z opens an affiliate in j for any realization of Y_j because the value of its potential export network is larger than the cost of opening the affiliate.

Figure 3: Model's implications.

Affiliate size, export status, and the timing of entry.



Host market characteristics and the timing of entry.



Proposition 2 relates to the expansion strategies of a MNE across countries. Since entry thresholds are decreasing in the preference shifter, the model predicts that—keeping host market size constant and conditional on aggregate productivity increasing over time ($\tilde{\mu} \geq 0$)—a MNE first opens its largest affiliates and subsequently opens its smaller affiliates. Similarly, since entry thresholds are increasing in entry costs, the model predicts that a MNE first opens affiliates in markets that are less costly to enter.

The lower panels of Figure 3 illustrate the predictions of Proposition 2. Figure 3c plots entry thresholds in two host countries of the same size, ($Q_k = Q_j$) but with different taste shifters

($\lambda_{kk} < \lambda_{jj}$), so that $Y_k^{OH}(z, \lambda_{kk}) \geq Y_j^{OH}(z, \lambda_{jj})$. Firm z only opens an affiliate in country j when the realization of the aggregate shock is Y' . When the realization of the shock grows to $Y'' > Y'$, the firm can also afford to open an affiliate in country k , illustrating that, controlling for factor costs and host country size, a MNE opens its largest affiliates first. Figure 3d plots entry thresholds in two host countries with different entry costs, $F_k^h > F_j^h$, so that $Y_k^{OH}(z, F_k^h) \geq Y_j^{OH}(z, F_j^h)$. Firm z opens an affiliate in country j when the realization of the composite shock is Y' . When the realization of the composite shock increases to $Y'' > Y'$, the firm can also afford to open an affiliate in country k .

Numerical simulations of the model reveal that the implications described in this section also hold in the general case where fixed costs are “large”, and hence, exit thresholds are active.

4 Calibration

We calibrate the model to match the expansion of US MNEs during the period 1987-2011, in the top-ten host countries for US FDI (Brazil, Canada, China, France, United Kingdom, Germany, Ireland, Japan, Mexico, and Singapore). We set the values of preference and technology parameters using estimates from the literature and direct observations from the data. Then, we jointly calibrate the rich set of barriers to MNE expansion included in the model to match static and dynamic moments from the BEA data.

4.1 Procedure

We set the elasticity of substitution $\eta = 5$, in line with estimates in the literature (Broda and Weinstein, 2006). We need to set the time preference rate to $\rho = 0.1$ so that it does not violate the technical condition that ensures that the present discounted value of profits in market j does not diverge ($\rho > \tilde{\mu}_j, \forall j$).¹⁸ We assume that the distribution of firm productivities is Pareto, with location parameter normalized to $b = 1$ and shape parameter $\vartheta = 4.5$, consistent with estimates in the literature (Simonovska and Waugh, 2014).

We use data on expenditure-based real GDP growth across countries, from the Penn World Tables 9.0, to calibrate the composite-shock process, for each country in our sample. The composite shock

¹⁸ The value of $\rho = 0.1$ might appear high, but its interpretation includes economic magnitudes other than just the time preference rate. For example, if the model included an exogenous death rate, this variable would be added to the time preference rate and the technical condition would allow for a lower time preference rate. Since the solution of the model would be unchanged, we prefer not to add unnecessary parameters and rather to assume a high value for the time preference rate.

Y_j captures the effect on profits of both US aggregate productivity and aggregate demand in country j . We set the drift of the process, $\tilde{\mu}_j$, to match real GDP growth in country j . Matching $\tilde{\sigma}_j$ to the standard deviation of real GDP growth, however, would generate too little volatility to induce reasonable firm dynamics. For this reason, we first set the standard deviation of US aggregate productivity, σ , to match the standard deviation of labor productivity among US firms, and the standard deviation of the aggregate demand shock in country j , σ_j , and its correlation with the US aggregate shock, γ_j , to match, respectively, country j 's standard deviation of real GDP growth and its correlation with US GDP growth. We then use (8), together with $\eta = 5$, to recover the values of $\tilde{\sigma}_j$. To initialize the shock processes, we normalize the initial value of the US productivity shock to $Z(0) = 1$, and the US demand shock to $Q_{US}(0) = 1$. We then set $Q_j(0)$ to be equal to country j 's GDP relative to US GDP, $\forall j$. Finally, we set (exogenous) wages in the model to match real GDP per unit of equipped labor, from Klenow and Rodríguez-Clare (2005), an average over the period 1995-2000. Appendix Table D.1 shows the results for each of the top ten host countries for US FDI.

It remains to calibrate the preference shifters, λ_{ijk} , and the parameters related to the costs of MNE expansion. These costs are: the fixed and sunk costs of affiliate opening, f_j^h and F_j^h , for $j = 1, \dots, 10$; the fixed and sunk costs of affiliate exports, f_{jk}^e and F_{jk}^e , for $j = 1, \dots, 10$, $k = 1, \dots, 10, US$, and $k \neq j$; and the iceberg trade costs of affiliate exports, τ_{jk} , for $j = 1, \dots, 10$, $k = 1, \dots, 10, US$, and $k \neq j$.

Due to data limitations, we need to make some symmetry assumptions.¹⁹ First, to represent possible taste differences between domestic goods and goods produced by foreign firms, we assume that $\lambda_{kkk} = 1$ and $\lambda_{ijk} = \lambda_j \neq 1$, for $i = US$ and for all k . These taste shifters allow us to generate different market shares for domestic firms and US MNEs in a host country. Second, we assume that the fixed and sunk costs of affiliate exports are symmetric across all destination countries, except for the United States: $f_{jk}^e = f_j^e$ and $F_{jk}^e = F_j^e$, for $j = 1, \dots, 10$, $j = 1, \dots, 10$, $k \neq j$, and $k \neq US$. Third, we assume that iceberg trade costs for destinations for which we do not have any bilateral affiliate export data are proportional to bilateral distance and to an exporter-specific dummy which is chosen to exactly match the aggregate export share from country j to all destinations.²⁰

Additionally, since aggregate demand grows over time at the rate μ_j , we assume that the fixed and sunk costs of MNE activities in each host country j also grow at the deterministic rate μ_j . Hence, we need to calibrate the initial values of the fixed and sunk costs for each host country. Without this growth adjustment, frictions to MNE activities would become irrelevant to the firm's decision

¹⁹ As mentioned in Section 2.1, the BEA data do not record affiliate exports by destination country, except for the United States, and for a handful of countries (Canada, Japan, and the United Kingdom) in benchmark-survey years.

²⁰ The distance elasticity is calculated by running a standard gravity equation with two sets of fixed effects and assuming that the trade elasticity is 4.

over long time horizons; otherwise, the adjustment does not affect any other property of the model.

We are left with 117 parameters to calibrate for which we target 117 moments from the data. Even though the model does not have a one-to-one mapping from each parameter to each moment, and parameters are jointly calibrated, because of the model’s closed-form solutions, it is relatively easy to isolate the moment that drives the identification of a given parameter. More precisely, the intensive margin of exports, given by export sale shares, drives the identification of the iceberg trade cost τ_{jk} , while affiliate entry rates and the share of MNE affiliates in each country help identify the sunk and fixed MNE entry costs, F_j^h and f_j^h , respectively.²¹ Similarly, export entry rates and the share of exporting affiliates help identify the sunk and fixed export costs, F_j^e and f_j^e , respectively. Finally, the ratio of affiliate horizontal sales in country j to parent US sales helps identify the taste shifter λ_j .

We choose the values of the parameters that best fit the data moments, for each country. To this end, we simulate the model 100 times, each time for a different realization of the vector of aggregate shocks. Each simulation amounts to solving the model for 1,000 firms and 30 years. Computationally, this entails solving $N + N^2$ systems of four equations in four unknowns, for each firm and time period, as well as solving for the equilibrium price index every period.

4.2 Model’s fit

Table 4 reports simulated and data moments taking averages across the top-ten host countries for US FDI and across years. Appendix Tables D.4-D.9 report the full set of simulated and data moments, while Appendix Tables D.2 and D.3 show the calibrated parameters by country. We construct moments from the data using the sample of all affiliates operating in the top-ten host countries for US FDI. This sample includes 83,214 affiliate-year observations, which account for 68.8 percent of all sales by foreign affiliates of US MNEs.

Table 4 shows that the model matches quite well both the static and dynamic targeted moments. We also include in the table two sets of non-targeted moments: moments related to the affiliate size advantage and exit moments.

The moments capturing the affiliate size advantage are related to the analytical predictions of the model described in Section 3.4. Proposition 1 implies that, controlling for the affiliates’ host market, affiliates that export have larger horizontal sales than affiliates that do not export. In the data, the

²¹ Since the share of affiliates, affiliate entry rates, and affiliate exit rates are linearly dependent, it is enough to target two out of the three moments. In the calibration, we target the share of affiliates and affiliate entry rates, and leave affiliate exit rates as non-targeted moments.

Table 4: Moments: model versus data, averages.

	data	model
Targeted Moments		
1. Static moments: intensive margin		
1.1 Affiliate sales share to host country	0.026	0.026
1.2 Affiliate sales share to the US	0.139	0.139
1.3 Affiliate sales share to third countries	0.288	0.296
1.4 Affiliate sales share to Canada	0.015	0.014
1.5 Affiliate sales share to the UK	0.069	0.087
1.6 Affiliate sales share to Japan	0.033	0.026
2. Static moments: extensive margin		
2.1 Share of MNEs with affiliates in j	0.287	0.283
2.2 Share of affiliates in j exporting to US	0.566	0.566
2.3 Share of affiliates in j exporting to third countries	0.650	0.646
3. Dynamic moments: entry		
3.1 Share of MNEs opening affiliates in j	0.035	0.021
3.2 Share of affiliates in j that start exporting to the US	0.030	0.024
3.3 Share of affiliates in j that start exporting to third countries	0.031	0.027
Non-Targeted Moments		
4. Static moments: size advantage		
4.1 Exporter size advantage	6.27	6.97
4.2 Early-exporter size advantage	3.68	5.54
4.3 First-affiliate size advantage	2.57	2.13
5. Dynamic moments: exit		
5.1 Share of MNEs shutting down affiliates in j	0.113	0.083
5.2 Share of affiliates in j that stop exporting to the US	0.025	0.042
5.3 Share of affiliates in j that stop exporting to third countries	0.027	0.037

Note: Averages across host countries and years. Data moments for Japan, Canada, and the United Kingdom are averages over benchmark-year surveys only. Shares' denominators are: in 1.1, US parent's sales; and in 1.2-1.6, total horizontal sales of affiliates in j ; in 2.1, the total number of MNEs; in 2.2 and 2.3, the total number of affiliates in j ; in 3.1, total number of MNEs in period before entry; and in 3.2 (3.3), total number of affiliates in j in period before export entry into US (third countries); in 5.1, total number of affiliates in j in period before exit; and in 5.2 (5.3), the total number affiliate in j that export to the US (third countries) in the period before stopping the activity. In 4.1, exporter size advantage refers to the average size of exporting MNE affiliates relative to the average size of non-exporting MNE affiliates, an average across countries and years; in 4.2, affiliate early-exporter size advantage refers to the average size of exporting MNE affiliates that start exports early in life relative to the average size of exporting MNE affiliates that start exports later in life; and in 4.3, first-affiliate size advantage refers to the ratio of the size of first foreign affiliate of a MNE (relative to GDP in the affiliate host market) to the size of subsequent foreign affiliates of the same MNE (relative to GDP in the affiliate host market), an average across MNEs and years. For moments in 4., size refers to horizontal affiliate sales; early versus late exporters refers to affiliates that are born with exports versus the ones that start exporting later. Calculations in the calibrated model trim the upper and lower 10th decile of the simulated firm-level data.

average horizontal sales of an affiliate that exports from birth are 6.3 times larger than the average horizontal sales of an affiliate that never exports, averaging across affiliates' host markets. Our calibrated model generates an exporter premium among MNE affiliates of around seven. Similarly, the model predicts that affiliates that start exporting earlier in their life have larger horizontal sales than affiliates that start exporting later. In the data, the average horizontal sales of an affiliate that starts exporting early in life are 3.7 times larger than the average horizontal sales of an affiliate that starts exporting later in life averaging across affiliates' host markets.²² Our calibrated model generates an early-exporter premium of 5.5. Appendix Table D.9 shows results by country.

Additionally, the model has predictions about the expansion patterns of a MNE. Proposition 2 implies that MNEs open their largest affiliates first. In the data, on average, the horizontal sales of a first affiliate of a MNE are 2.6 times larger than the horizontal sales of the MNE subsequent affiliates. The model generates a first-affiliate size premium of 2.1.²³

Finally, as shown in the last panel of Table 4, the exit rates in the model are close to the ones we observe in the data. The model slightly underpredicts affiliate exit, and overpredicts exit from export markets.

4.3 The costs of MNE expansion

We now evaluate the magnitude of the costs of MNE expansion in time and space. Since model-based magnitudes are hard to interpret, in Table 5 we express the calibrated MNE costs as shares of firm revenues and in monetary values. Appendix Tables D.11 and D.12 report costs as shares of revenues by host country, while Appendix Table D.13 shows sales, in US dollars, as observed in the BEA data, by host country.

On average, opening an affiliate involves spending a very low share of a firm's US parent revenues. An affiliate's fixed operating costs range from about two percent of the affiliate's horizontal sales, for the largest affiliates, to about 19 percent, for the smallest affiliates. In monetary terms, affiliate export operations appear less costly than affiliate horizontal operations. As expected, affiliate exports to the United States have lower costs than exports to other destination markets, especially in terms of sunk costs. This result is intuitive as the United States is the origin country of the affiliates.

²² Appendix Figure D.1 shows size distributions for the different groups of affiliates.

²³ Proposition 2 also states that MNEs first open affiliates located in countries where entry is less costly. Even though measuring entry costs directly in the data is difficult, we proxy them with the commonly used World Bank Doing Business indicators and provide suggestive evidence that supports the model's prediction, in Appendix Table D.10.

Table 5: Calibrated MNE costs: shares of sales and monetary values, average across host countries.

	As % of sales			In thousands \$
	5th	50th	95th	
Sunk affiliate entry cost F_j^h , as % of US parent sales	0.05	0.02	0.00	161.0
Fixed affiliate entry cost f_j^h , as % of horizontal sales	18.7	12.2	2.30	1,525
Sunk export cost F_{jk}^e , as % of horizontal sales				
To United States	0.33	0.16	0.03	15.60
To other destinations	3.00	0.99	0.18	97.30
Fixed export cost f_{jk}^e , as % of exports sales				
To United States	12.7	10.1	1.70	98.50
To other destinations	12.8	11.6	3.30	814.6

Note: For F_j^h and f_j^h , we consider variables in the year of affiliate entry. For F_{jk}^e and f_{jk}^e , we consider variables in the year the affiliate first exports to the destination. Percentiles are with respect to affiliate sales in the calibrated model. Cost shares in the model are converted into thousands of US dollars using sales values for the median affiliate in each of the top-ten host countries included in the calibration, from the BEA, and averaged across host countries. For confidentiality purposes, median sales are an average of the 9 observations around the median.

The relative magnitudes of different fixed and sunk costs of affiliate operations and affiliate exports are informative about the heterogeneity of the frictions to MNE activities, both by country and by type of affiliate sales. Table 6 shows that opening affiliates in Brazil, Canada, China, Japan and Mexico is much more costly than starting export operations from them. Favorable trade regimes, such as NAFTA, or tax policies favoring exports but restricting horizontal FDI, as in China, can rationalize this pattern. Ratios lower than one for countries like Ireland and Singapore may be explained by favorable “tax-haven”-like policies that attract FDI.

There is a large amount of heterogeneity in the fixed costs of setting up an affiliate relative to the fixed cost of exporting from it. Nonetheless, all the calibrated ratios are above one, indicating that the operating cost of horizontal activities is higher than the operating cost of export operations in all host countries.

The third and fourth columns in Table 6 shows the peculiarity of the United States as an export destination for affiliates of US MNEs. If the United States were an “average” export market for affiliates, the reported ratios would be close to one. This is true only for the sunk export cost paid by affiliates located in the United Kingdom where the United States appears similar to other export destinations. Surprisingly, the United States seems to be a more expensive market for affiliates located in Canada and Mexico. This may well be because, for example, affiliates located in Canada also export large magnitudes to Mexico, which is also part of NAFTA, and they may be

Table 6: Calibrated MNE costs: by type, destination, and host country.

	F_j^h/F_j^e	f_j^h/f_j^e	$F_{jk}^e/F_{j,US}^e$	$f_{jk}^e/f_{j,US}^e$
	(1)	(2)	(3)	(4)
Brazil	117	4.80	4.20	0.82
Canada	NaN	4.95	0.00	0.63
China	2.10	1.80	0.45	0.75
France	0.31	3.10	470	2.26
United Kingdom	0.01	2.80	1.01	1.91
Germany	0.15	2.80	30.2	2.45
Ireland	0.43	1.90	NaN	0.95
Japan	42.2	6.40	36.4	1.21
Mexico	12.6	4.70	0.06	0.70
Singapore	0.03	1.07	6.21	0.84

Note: In column 1 (column 2), the denominator is an average of the sunk (fixed) costs of exporting to the United States and to any other destination k , from country j . In column 3 (column 4), the ratio is between the sunk (fixed) export cost from j to any destination k , relative to the costs of exporting to the United States. "NaN" are ratios in which the denominator is almost zero.

associated with lower costs than with the United States.

5 Quantitative Analysis

Armed with the calibrated model, we explore the implications of counterfactual scenarios on aggregate firm dynamics. Since the top-ten host countries for US FDI in our sample include the United Kingdom, Ireland, Germany, and France, we use, as our first counterfactual exercise, the possibility of the United Kingdom abandoning the European Union (EU), "Brexit". Different proposed implementations of Brexit have as a common element the increase in export costs between the United Kingdom and other countries in the EU. In our model, this increase can be captured either as an increase in the iceberg trade cost, $\tau_{UK,j}$, an increase in the fixed export cost, $f_{UK,j}^e$, or an increase in the sunk export cost, $F_{UK,j}^e$, for $j = \text{France, Germany, Ireland}$. We can think about "shallow Brexit" scenarios, where only one type of friction is affected, or "deep Brexit" scenarios, where all the export barriers increase simultaneously. Additionally, in the Brexit context, we evaluate the importance of including the endogenous price response to the various shocks by comparing the aggregate dynamics of the model with and without endogenous prices. In our model, the magnitude of the price change is also informative about changes in the real wage—i.e., our welfare measure.

In a second counterfactual, we raise the cost of MNE activities in the United Kingdom, captured in our model by the per-period cost of MNE operations, f_{UK}^h . We use this exercise to evaluate

quantitatively the role played by the compound-option structure of our model. To such end, we evaluate the response to a cost shock in a calibrated model with and without export platforms.

For each counterfactual exercise, we simulate the model for 30 periods and impose a permanent change in one or more of parameters at $t = 15$.

5.1 Brexit

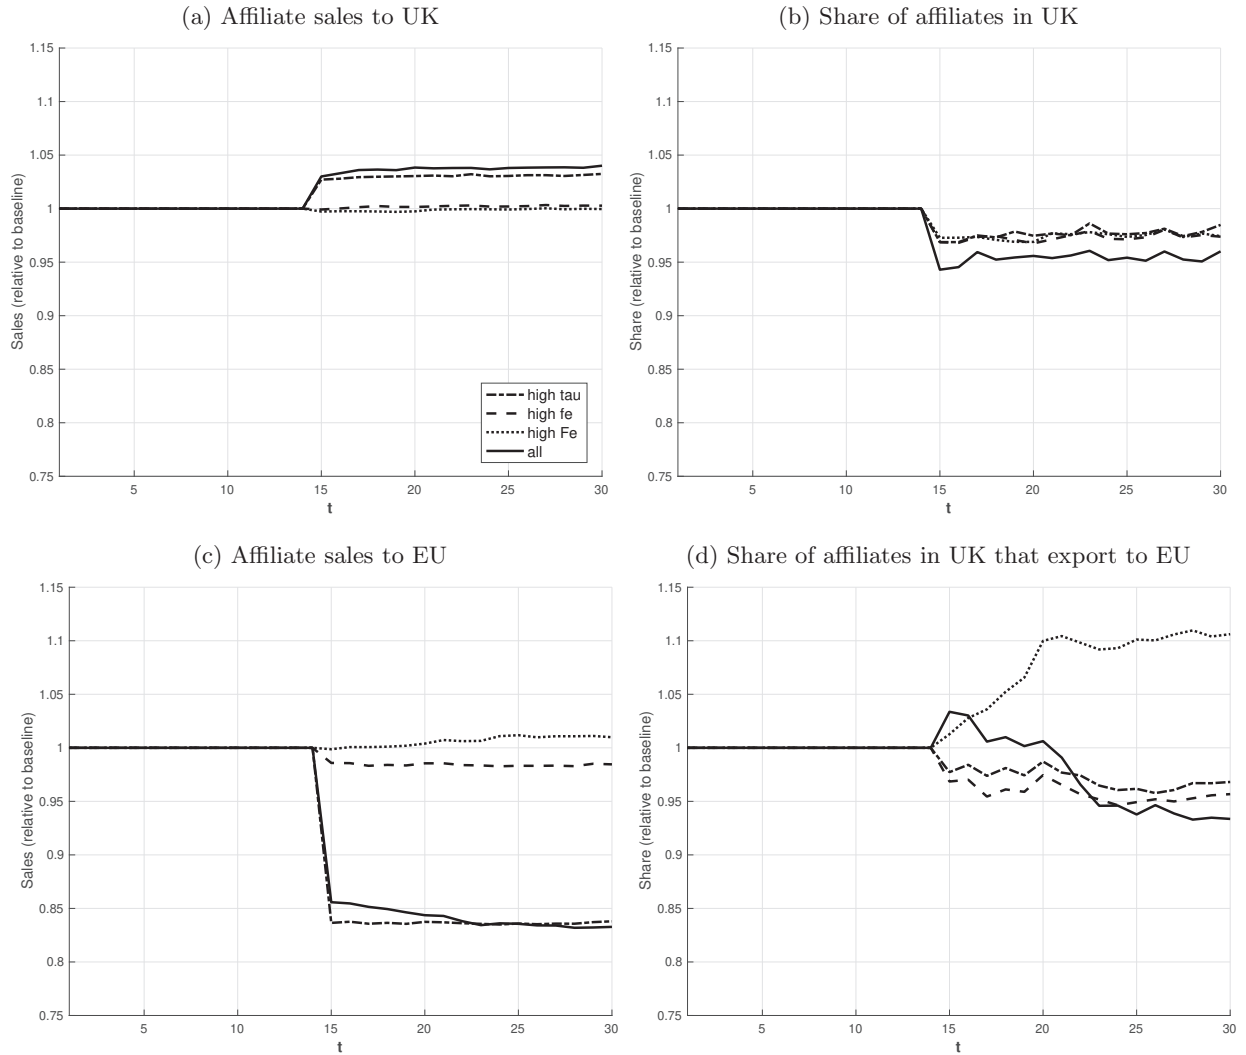
First, we increase, one at the time, the barriers to export from (to) the United Kingdom to (from) the EU countries in our sample (“shallow Brexit”): $\tau_{UK,j}$, $f_{UK,j}^e$, and $F_{UK,j}^e$, for $j = \text{Ireland, Germany, and France}$. In order to make the results of the three exercises comparable, we increase each trade friction by an amount that results in a 20 percent increase in the total per-period cost of FDI, $\left(f_{UK,j}^e + \rho F_{UK,j}^e\right) \tau_{UK,j}^{\eta-1}$. For “deep Brexit”, we increase all export barriers at once. Each barrier is increased by the same amount as in the “shallow Brexit” exercises.

Intuitively, increasing trade barriers between the United Kingdom and other EU countries has three main effects. First, exporting from the United Kingdom to the EU becomes more costly, so that export sales from UK-based affiliates to countries in the EU decline, decreasing the incentive to open affiliates in the United Kingdom due to the smaller, and more costly, available network of exporting destinations. Analogously, exporting from the EU to the United Kingdom also becomes more costly, so that export sales from affiliates in EU countries to the United Kingdom also decline, decreasing the incentive to open affiliates in those countries. These are static, partial equilibrium effects. In addition, increases in trade costs affect the affiliate export band of inaction, which in turn affects affiliate export entry and exit. Finally, increases in trade frictions have the effect of raising the price index in the destination countries, encouraging more export entry from the United Kingdom. Our quantitative results combine the effects of these three forces.

Figure 4 shows the effects of “shallow Brexit” on affiliates in the United Kingdom. Results are presented as deviations from the baseline scenario—i.e., the calibrated model. In all cases, higher trade costs between the United Kingdom and other EU countries reduce the share of affiliates based in the United Kingdom. Before the shock, about half of US MNEs had affiliates located in the United Kingdom, as shown in Appendix Table D.6. All three “shallow Brexit” shocks cause a permanent decrease in the share of US MNE affiliates in the United Kingdom of around three percent (Figure 4b). The cumulative increase of all frictions (“deep Brexit”) drives a decrease in the share of affiliates in the United Kingdom of five percent. Naturally, the affiliates that exit are the smallest and account for only a few percentages of US affiliate sales in the United Kingdom. This observation, together with the contemporaneous increase in the price index in the United

Kingdom, shown in Figure 7, explains the somewhat counterintuitive result in Figure 4a whereby affiliates sales to the United Kingdom increase when variable trade costs increase. The effects of changing trade costs on horizontal activities in the United Kingdom are small, since trade frictions affect those activities only indirectly through the compound option.

Figure 4: Brexit: US MNEs' affiliates in the United Kingdom.



Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

The lower panels of Figure 4 show the effect of “shallow Brexit” on the export sales and export participation rates of UK-based affiliates of US MNEs to Ireland, Germany, and France. These plots illustrate how different frictions to MNE activities have very different quantitative effects on

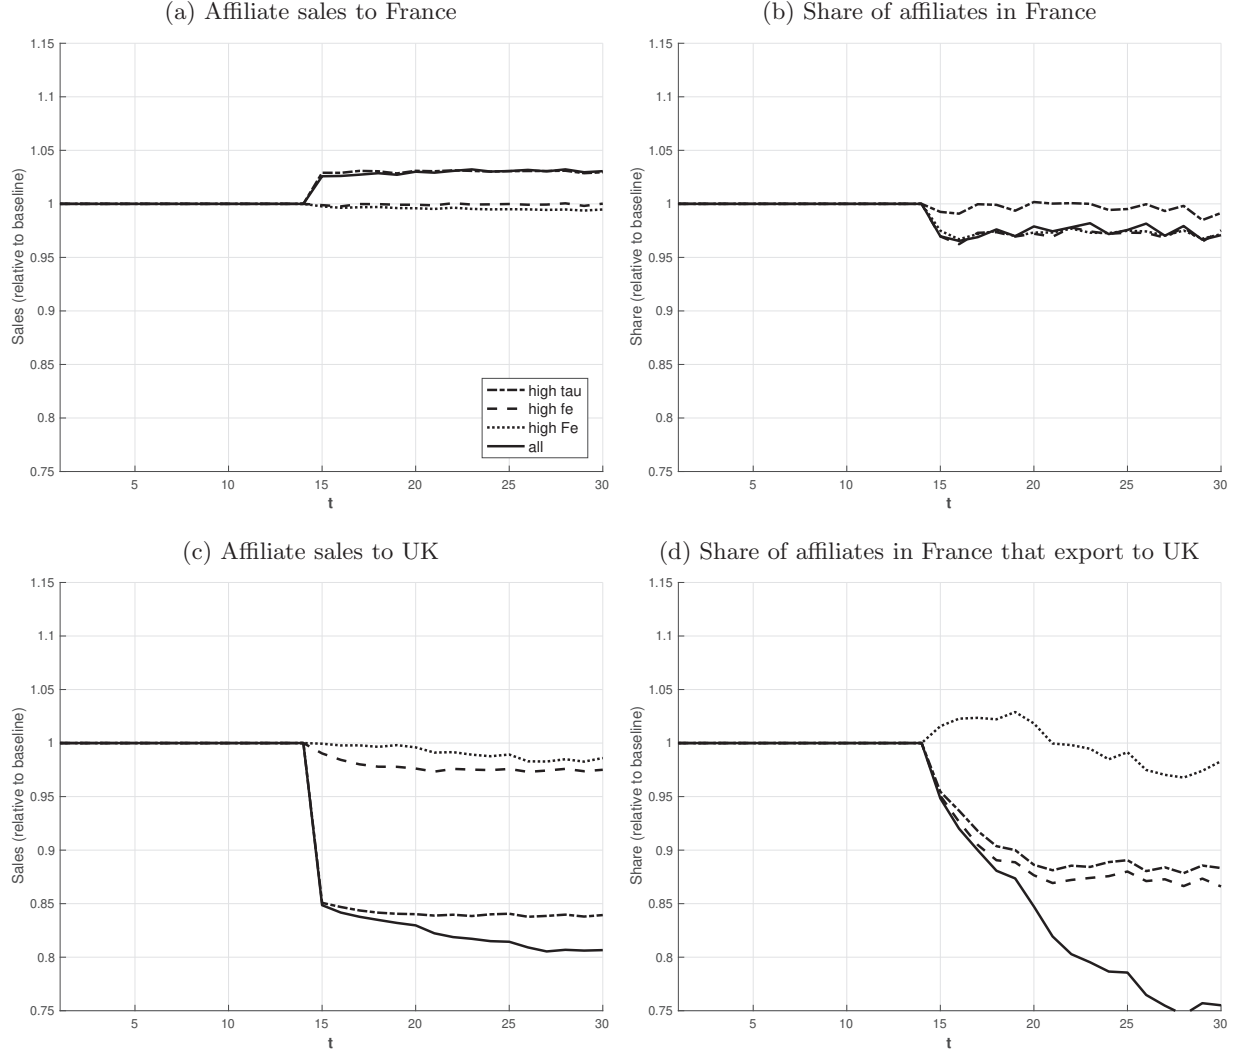
affiliate exports and participation rates, even if the changes in those frictions are associated with the same increase in the per-period cost of FDI. More precisely, the observed decline in export sales, when either $f_{UK,j}^e$ or $\tau_{UK,j}^e$ increase, comes from affiliates that stop exporting. The decline in export sales also includes an intensive-margin decline for the case of a shock to $\tau_{UK,j}$. Consequently, the change that has the highest impact on affiliate export sales is the increase in the per-period iceberg trade cost: an increase in $\tau_{UK,j}$ corresponding to a 20 percent increase in the cost of FDI produces a 15 percent decline in UK-based affiliates' export sales, a much larger decline than the one produced by increases in fixed export costs. Conversely, an increase in the sunk export cost, $F_{UK,j}^e$, produces a small increase in affiliate export sales. As Appendix Figure E.5 shows, an increase in the sunk export cost increases the export band of inaction, driving a decline in both affiliate export entry and exit rates. The decline in the exit rate is the most pronounced, giving rise to the increase in affiliate export sales observed in Figure 4c, and in the share of exporting affiliates observed in Figure 4d. Except for the case of an increase in the sunk export cost, the Brexit shock has the effect of reducing the share of affiliates that export. The increase in the fixed export cost produces the largest decline in the export participation rate because this cost is intimately related to the affiliates' decision to exit a market.

The case of “deep Brexit” is interesting as it shows important differences between the aggregate dynamics after a shock in the short and long run. In particular, in the short run the effect of the decrease in sunk export costs dominates and the share of affiliates in the United Kingdom that export to EU increases. Fifteen periods after the shock, this share decreases by around seven percentage points relative to the pre-shock levels.

The results for affiliates located in Ireland, Germany, and France are qualitatively similar to the ones for affiliates located in the United Kingdom. There are, however, some important quantitative differences. We show the results for affiliates of US MNEs located in France in Figure 5, and relegate the results for Ireland and Germany to Appendix Figures E.3 and E.4.

As expected, higher trade costs between France and the United Kingdom reduce the incentives to locate in France, and the share of affiliates of US MNEs located in France declines. As already seen for the United Kingdom, this decline in the extensive margin is accompanied by an increase in the intensive margin of horizontal sales in France for the scenario in which $\tau_{UK,FR}$ increases. This increase is driven by the corresponding increase in the price index in France. Naturally, the share of affiliates in France that export to the United Kingdom as well as their export sales drop after the Brexit shock in almost all specifications. An increase in the sunk export cost generates a non-monotonic response of the share of exporting affiliates, which first increases (due to a decline

Figure 5: Brexit: US affiliates in France.



Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

in exit rates) and then decreases.

Summing up, the results in Figures 4 and 5 show that increasing a variable, fixed, or sunk cost of export has a *quantitatively* different effect on aggregate firm dynamics. Even though the increase in the per-period costs of FDI is the same in all cases, the type of trade barrier that changes matters for both the extensive and intensive margins of firms’ decisions. When interpreting our results, it is important to keep in mind that: on the one hand, the lack of reallocation across markets due to the independence assumption may have the effect of overstating the losses from Brexit; on the other hand, exogenous wages imply that the aggregate price effects in our model are stronger than

what they would be in a full general equilibrium model, understating losses.

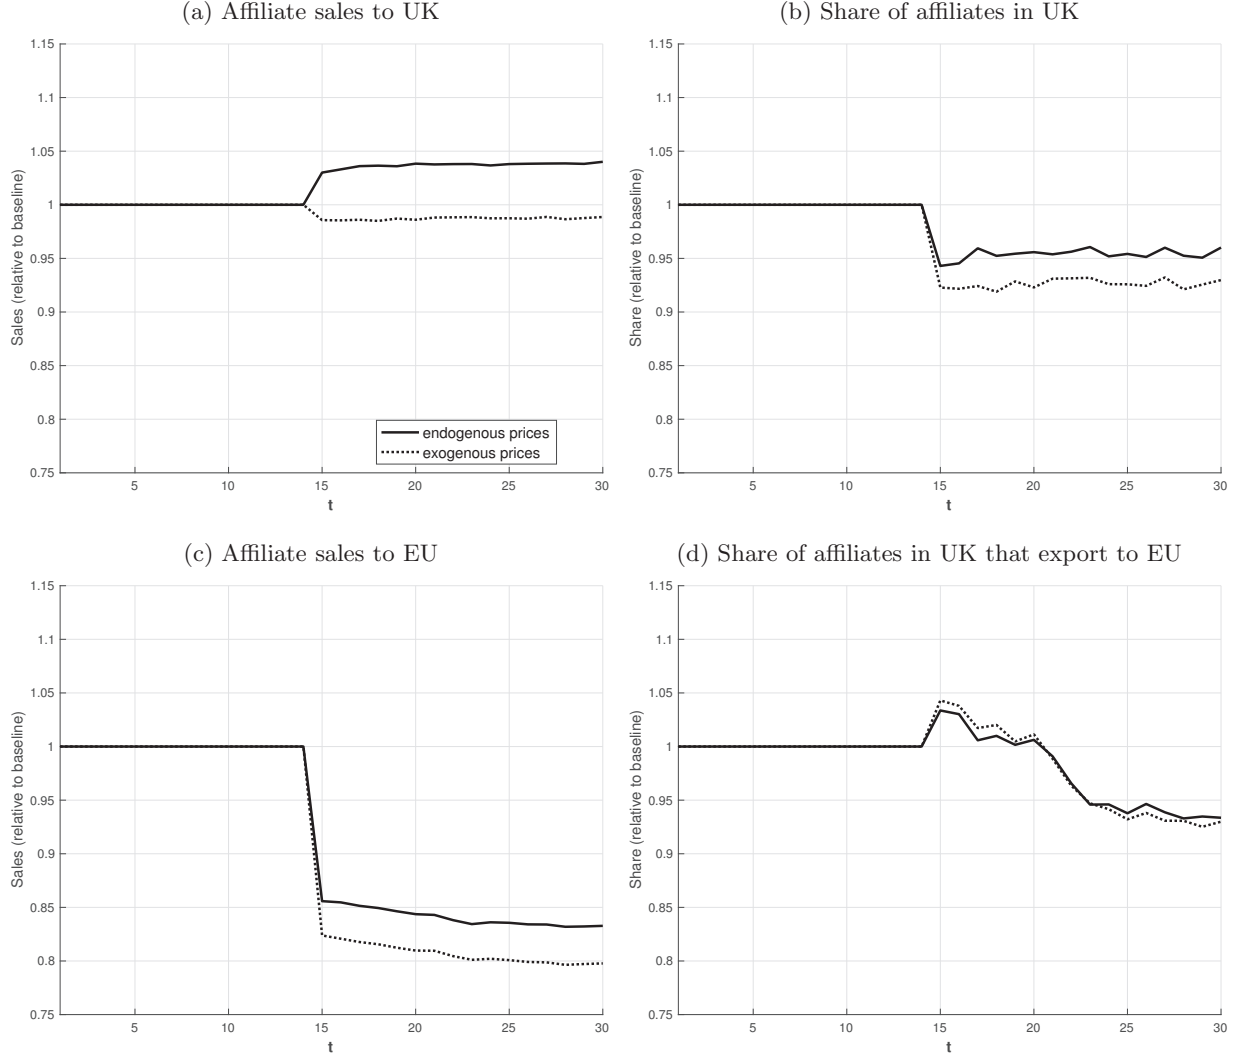
The tractability of our model allows us to solve numerically for the aggregate price index in a multi-country dynamic setup, so that the results of our counterfactual exercises incorporate the effects of changes in the price indexes on firms’ decisions. How important are these price effects *quantitatively*? Are they strong enough to affect the aggregate dynamics of MNEs? In Figure 6, we show the dynamics of aggregate outcomes after the “deep Brexit” shock, under endogenous and exogenous prices. It is clear that the endogenous response of prices acts as a buffer to the decrease in affiliate sales, both horizontal and exports, as well as to the decrease in the share of US MNEs operating in the United Kingdom. In the case of horizontal sales, shown in Figure 6a, the effect is strong enough to reverse the pattern from a two-percent decrease to a three-percent increase.

Appendix Figure E.1 shows results with exogenous prices for US MNE affiliates in the United Kingdom, for each of the Brexit counterfactual exercises.

It is straightforward to evaluate the welfare losses from the different Brexit scenarios from our counterfactual price changes. Figure 7 shows that, as expected, the largest welfare loss is experienced under “deep Brexit”, with a decrease in real income that goes from 0.7 percent in the short run to more than 0.8 percent after fifteen years. In the context of “shallow Brexit”, an increase in variable trade costs produces the highest loss, while an increase in sunk export costs produces the smallest loss. These magnitudes are not small given that our model is restricted to the behavior of MNEs from the United States. The increase in trade barriers would presumably affect local and other non-US MNE exporters. Appendix Figure E.2 shows the change in prices after each Brexit shock, for each of the countries involved. As expected, changes in prices—and corresponding welfare losses—are highest for Ireland, whose economy is deeply connected to the UK economy.

Our Brexit counterfactual illustrates the importance of considering the global structure of the MNE for accurately assessing the consequences of shocks to the costs of MNE expansion. In a model in which MNEs could only serve the host market of operation, one would erroneously conclude that Brexit, in the form of higher trade barriers, would not have any impact on the behavior in time and space of US MNEs. We explore in more detail the role played by export platforms in the next counterfactual.

Figure 6: Brexit: US affiliates in the United Kingdom. Endogenous vs exogenous prices.




Note: The shock refers to increasing all three export barriers from/to the United Kingdom to/from country j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

5.2 The role of the compound option

We now evaluate quantitatively the role played by the compound-option structure of our model. To such end, we analyze the effects of an increase in the barriers to MNE activities in a model with and without the compound option—or analogously, a model with and without MNE affiliate exports. We increase the fixed cost of MNE activities in the United Kingdom, f_{UK}^h , by 20 percent.

Our baseline model without the compound option is a calibrated version of the model in which affiliates of US MNEs cannot sell to countries other than their host country (i.e., no exports). In

Figure 7: Brexit: Price changes in the United Kingdom.



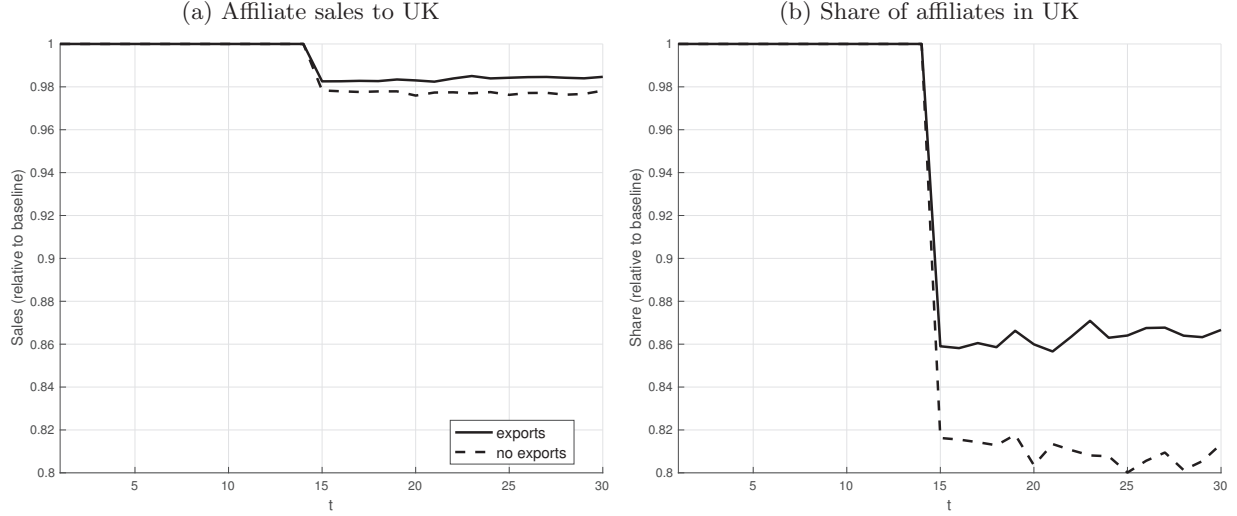
Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

this calibration, we target the same moments as for the model with exports (moments 1.1, 2.1, and 3.1 in Table 4).²⁴

Figure 8 shows the dynamics of horizontal affiliate sales and the share of US affiliates in the United Kingdom after the shock to f_{UK}^h at $t = 15$. Results are shown as deviations from the respective baselines—i.e., the calibrated models with and without the compound option. In the model with only horizontal sales (dashed lines), MNEs do not have the option of exporting part of their output. Hence, the incentives to open an affiliate in a country coming from the possibility of using that host country as an export platform are precluded. Without the possibility of exporting to other markets from the United Kingdom, an increase in barriers to MNE activities in the United Kingdom would decrease the presence of US affiliates in the country by 25 percent more (in the short run) than in the case in which exports are possible. Moreover, in the model without exports, this share would continue to decline over time, while in the model with export platforms the share of US affiliates would start increasing in the long run.

²⁴ The calibrated model without export platforms does not fit the relevant data moments as well as the baseline model. While the share of MNEs with affiliates in the United Kingdom is 53.7 (55.4) percent in the model without exports (data), the share of MNEs opening affiliates the United Kingdom is 2.2 (6.9) percent in the model (data), and the share of UK affiliates shutting down is 8.8 (4.5) percent in the model (data).

Figure 8: The role of the compound option: US affiliates in the United Kingdom.



Note: At $t = 15$, the fixed cost of MNE activities in the UK, f_{UK}^h increases by 20 percent. “Exports” refers to the calibrated model where MNE affiliates can export, while “no exports” refers to the calibrated model where MNE affiliates cannot export. Results are shown as deviations from the calibrated models with and without export platforms.

The result in Figure 8, once again, points to the importance of the compound-option structure when making dynamic quantitative predictions on the effects of shocks to MNE activities.

6 Conclusions

This paper studies the expansion patterns of the multinational enterprise (MNE) in time and space. Our facts, documented using a long panel of US MNEs, guide us on the development of a dynamic model of the MNE that is tractable and, at the same time, rich enough to capture the complexity of MNE activities observed in the data. The model features heterogeneity in firm productivity, persistent aggregate shocks, and a rich structure of MNE costs. Importantly, MNE affiliates can decouple their locations of production and sales, and endogenously choose to enter or exit both the host and the export markets. The novel feature of the model is the introduction of a compound option formulation, which is key to quantifying the costs of MNE expansion. Our counterfactual exercises reveal that the compound-option structure is important for understanding the reallocation of MNE activity in time and space after a shock. Importantly, these exercises also reveal that the nature of the frictions to MNE activities (variable, fixed, or sunk) is important for understanding the patterns of aggregate firm dynamics after a shock.

References

- Alessandria, G. and H. Choi (2007). Do sunk costs of exporting matter for net exports dynamics? *Quarterly Journal of Economics* 122(1), 289–336.
- Alessandria, G., H. Choi, and K. Ruhl (2018). Trade adjustment dynamics and the welfare gains from trade. NBER Working Paper #20663.
- Antràs, P. and S. R. Yeaple (2014). Multinational firms and the structure of international trade. *Handbook of International Economics* 4, 55–130.
- Arkolakis, C. (2016). A unified theory of firm selection and growth. *Quarterly Journal of Economics* 131(1), 89–155.
- Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (2018). Innovation and production in the global economy. *American Economic Review* 108(8), 2128–2173.
- Baldwin, R. and P. Krugman (1989). Persistent trade effects of large exchange rate shocks. *Quarterly Journal of Economics* 104(4), 635–654.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *Quarterly Journal of Economics* 121(2), 541–585.
- Conconi, P., A. Sapir, and M. Zanardi (2016). The internationalization process of firms: from exports to FDI. *Journal of International Economics* 99, 16–30.
- Cravino, J. and A. Levchenko (2017). Multinational firms and international business cycle transmission. *Quarterly Journal of Economics* 132(2), 921–962.
- Das, S., M. J. Roberts, and J. R. Tybout (2007). Market entry costs, producer heterogeneity, and export dynamics. *Econometrica* 75(3), 837–873.
- Davis, S. J., J. Haltiwanger, and S. Schuh (1996). Small business and job creation: Dissecting the myth and reassessing the facts. *Small Business Economics* 8(4), 297–315.
- Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014). The role of entrepreneurship in US job creation and economic dynamism. *Journal of Economic Perspectives* 28(3), 3–24.
- Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2016). Where has all the skewness gone? the decline in high-growth (young) firms in the US. *European Economic Review* 86, 4–23.
- Dixit, A. K. and R. S. Pindyck (1994). *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

- Eaton, J., M. Eslava, M. Kugler, and J. Tybout (2008). The margins of entry into export markets: Evidence from Colombia. In E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*. Cambridge, MA: Harvard University Press.
- Egger, P., M. Fahn, V. Merlo, and G. Wamser (2014). On the genesis of multinational foreign affiliate networks. *European Economic Review* 65, 136–163.
- Fan, J. (2017). Talent, geography, and offshore R&D. Mimeo, Penn State University.
- Fillat, J. L. and S. Garetto (2015). Risk, returns, and multinational production. *Quarterly Journal of Economics* 130(4), 2027–2073.
- Fillat, J. L., S. Garetto, and L. Oldenski (2015). Diversification, cost structure, and the risk premium of multinational corporations. *Journal of International Economics* 96(1), 37–54.
- Fitzgerald, D., S. Haller, and Y. Yedid-Levi (2017). How exporters grow. Mimeo, Federal Reserve Bank of Minneapolis.
- Ghironi, F. and M. J. Melitz (2005). International trade and macroeconomic dynamics with heterogeneous firms. *The Quarterly Journal of Economics* 120(3), 865–915.
- Gravelle, J. (2015). Tax havens: International tax avoidance and evasion. *Congressional Research Service* 7(5700), 1–55.
- Gumpert, A., A. Moxnes, N. Ramondo, and F. Tintelnot (2018). The life-cycle dynamics of exporters and multinational firms. *NBER Working Paper* (24013).
- Head, K. and T. Mayer (2019). Brands in motion: How frictions shape multinational production. *American Economic Review*, Forthcoming.
- Impullitti, G., A. A. Irarrazabal, and L. D. Oromolla (2013). A theory of entry and exit into exports markets. *Journal of International Economics* 90(1), 75–90.
- Klenow, P. and A. Rodríguez-Clare (2005). Externalities and growth. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 1A, Chapter 11, pp. 817–861. Elsevier.
- Kueng, L., M.-J. Yang, and B. Hong (2017). Sources of firm life-cycle dynamics: Differentiating size vs. age effects. Mimeo, University of Washington.
- Leahy, J. V. (1993). Investment in competitive equilibrium: the optimality of myopic behavior. *Quarterly Journal of Economics* 108(4), 1105–1133.

- Luttmer, E. G. J. (2007). Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics* 122(3), 1103–1144.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Morales, E., G. Sheu, and A. Zahler (2018). Extended gravity. *Review of Economic Studies*, Forthcoming.
- Ramondo, N., V. Rappoport, and K. J. Ruhl (2013). The proximity-concentration tradeoff under uncertainty. *Review of Economic Studies* 80(4), 1582–1621.
- Roberts, M. J. and J. R. Tybout (1997). The decision to export in Colombia: An empirical model of entry with sunk costs. *The American Economic Review* 87(4), 545–564.
- Ruhl, K. J. and J. Willis (2017). New exporter dynamics. *International Economics Review* 58(3).
- Simonovska, I. and M. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics* 92(1), 34–50.
- Tintelnot, F. (2017). Global production with export platforms. *Quarterly Journal of Economics* 132(1), 157–209.

Appendix

A Data Description

Reporting thresholds. The BEA collects firm-level data on the operations of US multinational enterprises (MNEs) in its annual surveys of US direct investment abroad. All US-located firms that have at least one foreign affiliate and that meet a minimum size threshold are required by law to respond to these surveys. These minimum size thresholds are in terms of affiliate sales and differ over time. In general, reporting thresholds increased in recent years, reaching US\$60 million by 2011. Additionally, benchmark survey years (i.e., years in which the survey is more comprehensive), which occur every 5 years, have lower reporting thresholds. Table A.1 shows the reporting thresholds for the years in our sample.

Table A.1: BEA minimum survey exemptions levels.

survey year	Minimum exemption levels (in US\$ millions)	survey year	Minimum exemption levels (in US\$ millions)
1987-88	10	2000-03	30
1989	3	2004	25
1990 -93	15	2005-07	40
1994	3	2008	60
1995-98	20	2009	25
1999	7	2010-11	60

Note: Exemption levels are for majority-owned foreign affiliates. Benchmark survey years are **in bold**.

Tax havens. Our sample contains affiliates that do not operate in tax haven countries. Affiliates in tax haven countries are likely to open for different reasons than production purposes, and to be subject to different cost structures than affiliates in non-tax haven countries. We exclude countries defined as tax havens by Gravelle (2015), except for Ireland, Switzerland, Hong Kong, and Singapore that meet some of the criteria for tax haven status but also have a substantial amount of US MNE production. Table A.2 reports the list of countries that we exclude from our sample.

Industry classification. Each foreign affiliate is assigned an industry classification based on its primary activity according to the BEA International Surveys Industry (ISI) system, which closely follows the 3-digit Standard Industrial Classification (SIC) system. The BEA uses 3-digit SIC-

Table A.2: Tax haven countries excluded from our sample.

Anguilla	Turks and Caicos	Monaco
Antigua and Barbuda	US Virgin Islands	San Marino
Aruba	Belize	Maldives
Bahamas	Costa Rica	Mauritius
Barbados	Panama	Seychelles
British Virgin Islands	Bermuda	Bahrain
Cayman Islands	Macau	Cook Islands
Dominica	Andorra	Marshall Islands
Grenada	Channel Islands	Samoa
Montserrat	Cyprus	Nauru
Netherlands Antilles	Gibraltar	Niue
St. Kitts and Nevis	Isle of Man	Tonga
St Lucia	Liechtenstein	Vanuatu
St Vincent and Grenadines	Malta	Liberia

Note: From Gravelle (2015).

based ISI codes for years prior to 1999. From 1999 onward, they use 4-digit NAICS-based ISI codes. For consistency, we convert the NAICS-based codes to 3-digit SIC-based ISI codes for the relevant years.

Unit of observation. According to the BEA definition, an affiliate is a business enterprise operating in a given host country; it thus can operate several plants in different locations within the host country. The BEA rules permit consolidated reporting for distinct plants located in the same country that operate in the same narrowly defined industry or otherwise are integral parts of the same business operation. We consolidate observations of enterprises belonging to the same parent company and operating in the same country and 3-digit industry. We group these enterprises' activities together and refer to them as a single affiliate.

B Additional Facts

Table B.1: Affiliate-to-parent sales ratio by sales type, robustness. OLS.

	Dependent variable: affiliate-to-parent sales ratio					
	First affiliate	Subsequent affiliate	Non-GVC affiliate	GVC affiliate	Greenfield affiliate	M&A affiliate
	(1)	(2)	(3)	(4)	(5)	(6)
D(age=2)	-0.0014 (0.0055)	0.0016 (0.0012)	-0.004 (0.0056)	0.0019 (0.0019)		
D(age=3)	0.0001 (0.0066)	0.0016* (0.0009)	-0.0142 (0.0141)	0.0039*** (0.0014)	0.0044 (0.0029)	0.0050* (0.0028)
D(age=4)	0.0003 (0.0067)	0.0022** (0.0009)	-0.0157 (0.0147)	0.0045*** (0.0014)	0.0079** (0.0031)	0.0038 (0.0031)
D(age=5)	0.0033 (0.0027)	0.0026*** (0.0009)	0.0111 (0.0122)	0.0032** (0.0013)	0.0069* (0.0039)	0.0034 (0.0031)
D(age=6)	0.0065 (0.0040)	0.0030** (0.0013)	-0.0020 (0.0053)	0.0057*** (0.0019)	0.0066* (0.0037)	0.0053* (0.0029)
D(age=7)	0.0018 (0.0032)	0.0029** (0.0013)	-0.0009 (0.0040)	0.0035** (0.0014)	0.0046 (0.0028)	0.0047* (0.0024)
D(age=8)	0.0013 (0.0036)	0.0029** (0.0013)	-0.0064 (0.0077)	0.0039*** (0.0012)	0.0056* (0.0029)	0.0049** (0.0022)
D(age=9)	0.0034 (0.0027)	0.0042*** (0.0015)	-0.0050 (0.0053)	0.0051*** (0.0014)	0.0089 (0.0054)	0.0060*** (0.0021)
D(age=10)	0.0119 (0.0094)	0.0039** (0.0016)	0.0279 (0.0291)	0.0043*** (0.0014)	0.0082 (0.0058)	0.0056*** (0.0019)
log global employment	-0.0254 (0.0208)	-0.0089 (0.0073)	-0.0418 (0.0347)	-0.0078 (0.0074)	0.0133 (0.0158)	0.0089 (0.0150)
Observations	17,360	20,720	10,320	27,760	2,214	3,564
R ² (overall)	0.01	0.027	0.01	0.02	0.01	0.02

Note: Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. The dependent variable *affiliate to parent sales ratio* refers to affiliate sales relative to the domestic sales of the US parent. First affiliate refers to the first foreign affiliate opened by the parent, while subsequent affiliate refers to second or higher. GVC affiliate refers to affiliates with positive intra-firm trade, while non-GVC affiliate refers to affiliates with zero intra-firm trade. M&A affiliate refers to affiliates created through a merger or acquisition of an existing firm, while greenfield affiliate refers to a new firm. All specifications include affiliate and country-year fixed effects. Standard errors, clustered at the parent level, are in parenthesis. Levels of significance are denoted *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

Table B.2: Intensive and extensive margins of sales. OLS.

Dependent variable	Intensive margin of sale shares				Extensive margin of sale shares			
	horizontal sales		export sales		horizontal sales		export sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
affiliate age	-0.002 (0.002)	-0.012*** (0.001)	-0.005** (0.002)	0.005*** (0.001)	0.00003 (0.001)	-0.001 (0.0006)	0.014*** (0.003)	0.029*** (0.002)
country-year fe	yes	yes	yes	yes	yes	yes	yes	yes
industry fe	yes	no	yes	no	yes	no	yes	no
affiliate fe	no	yes	no	yes	no	yes	no	yes
Observations	36,135	36,135	25,958	25,958	38,080	38,080	38,080	38,080
R-square	0.079	0.013	0.092	0.000	0.042	0.0001	0.081	0.036

Note: Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. In columns (1)-(4), the dependent variable is horizontal (export) sales, as a share of total affiliate's sales, for affiliates with positive horizontal (export) sales; in columns (5)-(8), the dependent variable is the share of affiliates with positive horizontal (export) sales. Standard errors, clustered at the parent level, are in parenthesis. Levels of significance are denoted *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

Table B.3: Intensive and extensive margins of sales, pure-type affiliates at birth. OLS.

Dependent variable	Intensive margin of sale shares				Extensive margin of sale shares			
	horizontal sales		export sales		horizontal sales		export sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
affiliate age	-0.014*** (0.001)	-0.096*** (0.001)	-0.036*** (0.004)	-0.021*** (0.003)	-0.046*** (0.004)	-0.044*** (0.002)	-0.059*** (0.005)	-0.038*** (0.004)
country-year fe	yes	yes	yes	yes	yes	yes	yes	yes
industry fe	yes	no	yes	no	yes	no	yes	no
affiliate fe	no	yes	no	yes	no	yes	no	yes
Observations	19,910	19,910	3,590	3,590	19,910	19,910	3,590	3,590
R-square	0.147	0.020	0.288	0.133	0.245	0.099	0.268	0.125

Note: Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. In columns (1)-(4), the dependent variable is horizontal (export) sales, as a share of total affiliate's sales, for affiliates born with only horizontal (export) sales; in columns (5)-(8), the dependent variable is the share of affiliates born with only horizontal (export) sales. Standard errors, clustered at the parent level, are in parenthesis. Levels of significance are denoted *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

Table B.4: Unconditional and conditional probability of affiliate entry, GVC vs non-GVC affiliates.

	Unconditional	Continent	Border	Language	Income	All
GVC Affiliates						
Canada	0.021	0.022 (0.733)	-	0.024 (0.000)	0.021 (0.360)	-
United Kingdom	0.026	0.028 (0.001)	0.030 (0.204)	0.027 (0.426)	0.027 (0.008)	0.030 (0.204)
Germany	0.024	0.026 (0.000)	0.030 (0.000)	0.028 (0.024)	0.025 (0.000)	0.028 (0.024)
Ireland	0.010	0.011 (0.003)	0.012 (0.012)	0.011 (0.001)	0.010 (0.005)	0.012 (0.012)
China	0.028	0.038 (0.000)	0.051 (0.000)	0.049 (0.000)	0.052 (0.000)	0.057 (0.000)
France	0.022	0.025 (0.000)	0.029 (0.000)	0.024 (0.059)	0.023 (0.000)	0.029 (0.000)
Brazil	0.017	0.022 (0.000)	0.027 (0.000)	0.025 (0.070)	0.024 (0.000)	0.019 (0.713)
Singapore	0.017	0.024 (0.000)	0.045 (0.000)	0.018 (0.000)	0.017 (0.292)	0.046 (0.000)
Mexico	0.025	0.030 (0.000)	0.028 (0.699)	0.034 (0.000)	0.031 (0.000)	0.024 (0.936)
Japan	0.016	0.022 (0.000)	-	-	0.017 (0.092)	-
Non-GVC Affiliates						
Canada	0.009	0.014 (0.483)	-	0.012 (0.002)	0.0073 (0.233)	-
United Kingdom	0.009	0.011 (0.516)	-	0.008 (0.668)	0.0091 (0.830)	-
Germany	0.010	0.011 (0.687)	0.014 (0.308)	0.017 (0.552)	0.0096 (0.552)	0.0165 (0.552)
Ireland	-	-	-	-	-	-
China	0.006	0.007 (0.656)	0.007 (0.795)	0.010 (0.444)	0.0217 (0.249)	-
France	0.006	0.007 (0.015)	0.014 (0.014)	0.009 (0.248)	0.006 (0.023)	0.0238 (0.069)
Brazil	0.005	0.007 (0.521)	0.033 (0.192)	-	0.007 (0.547)	-
Singapore	-	-	-	-	-	-
Mexico	0.006	0.010 (0.176)	-	0.024 (0.131)	0.0210 (0.143)	-
Japan	0.005	-	-	- (0.386)	0.0038	-

Note: Probability of opening affiliates in the top-ten most popular destinations for US MNEs. Conditional probabilities refer to the probability of observing a MNE opening an affiliate in country i given that the parent already has an affiliate in a “similar” country. The sample is restricted to parents with at least two affiliates worldwide. GVC affiliate refers to affiliates with positive intra-firm trade, while non-GVC affiliate refers to affiliates with zero intra-firm trade. Results for non-GVC affiliates in Ireland and Singapore are not shown for confidentiality reasons. Conditional probabilities in **bold** are not significantly different from the relevant unconditional probability.

Table B.5: Unconditional and conditional probability of affiliate entry, different samples.

	Unconditional	Continent	Border	Language	Income	All
Parents with at least five affiliates worldwide						
Canada	0.0252	0.0258 (0.724)	– –	0.0258 (0.000)	0.0250 (0.286)	– –
United Kingdom	0.0300	0.0302 (0.661)	0.0324 (0.506)	0.0301 (0.865)	0.0230 (0.520)	0.0324 (0.506)
Germany	0.0345	0.0353 (0.000)	0.0353 (0.284)	0.0319 (0.210)	0.0346 (0.000)	0.0319 (0.210)
Ireland	0.0166	0.0169 (0.025)	0.0174 (0.311)	0.0169 (0.163)	0.0167 (0.000)	0.0174 (0.311)
China	0.0451	0.0489 (0.000)	0.0568 (0.000)	0.0579 (0.000)	0.0581 (0.000)	0.0599 (0.004)
France	0.0335	0.0341 (0.000)	0.0364 (0.000)	0.0322 (0.263)	0.0335 (0.000)	0.0339 (0.758)
Brazil	0.0273	0.0290 (0.212)	0.0294 (0.413)	0.0253 (0.674)	0.0308 (0.001)	0.0188 (0.091)
Singapore	0.0289	0.0322 (0.000)	0.0507 (0.000)	0.0293 (0.061)	0.0290 (0.000)	0.0508 (0.000)
Mexico	0.0340	0.0354 (0.351)	0.0313 (0.739)	0.0367 (0.121)	0.0361 (0.107)	0.0243 (0.247)
Japan	0.0269	0.0288 (0.007)	– –	– –	0.0269 (0.598)	– –
Parents with at least ten affiliates worldwide						
Canada	0.0291	0.0291 (0.964)	– –	0.0292 (0.017)	0.0291 (0.286)	– –
United Kingdom	0.0273	0.0275 (0.000)	0.0264 (0.810)	0.0276 (0.218)	0.0273 (0.610)	0.0264 (0.809)
Germany	0.0354	0.0357 (0.000)	0.0349 (0.508)	0.0306 (0.027)	0.0354 (0.000)	0.0306 (0.027)
Ireland	0.0224	0.0224 (0.862)	0.0220 (0.680)	0.0225 (0.002)	0.0224 (0.105)	0.0220 (0.680)
China	0.0587	0.0600 (0.069)	0.0616 (0.382)	0.0661 (0.007)	0.0598 (0.687)	0.0620 (0.515)
France	0.0395	0.0396 (0.697)	0.0401 (0.214)	0.0376 (0.123)	0.0395 (0.116)	0.0382 (0.283)
Brazil	0.0321	0.0299 (0.112)	0.0277 (0.084)	0.0286 (0.505)	0.0323 (0.808)	0.0183 (0.007)
Singapore	0.0432	0.0439 (0.238)	0.0564 (0.008)	0.0431 (0.556)	0.0432 (0.610)	0.0564 (0.008)
Mexico	0.0423	0.0412 (0.535)	0.0282 (0.085)	0.0393 (0.088)	0.0388 (0.009)	0.0247 (0.039)
Japan	0.0361	0.0365 (0.478)	– –	– –	0.0361 (0.865)	– –

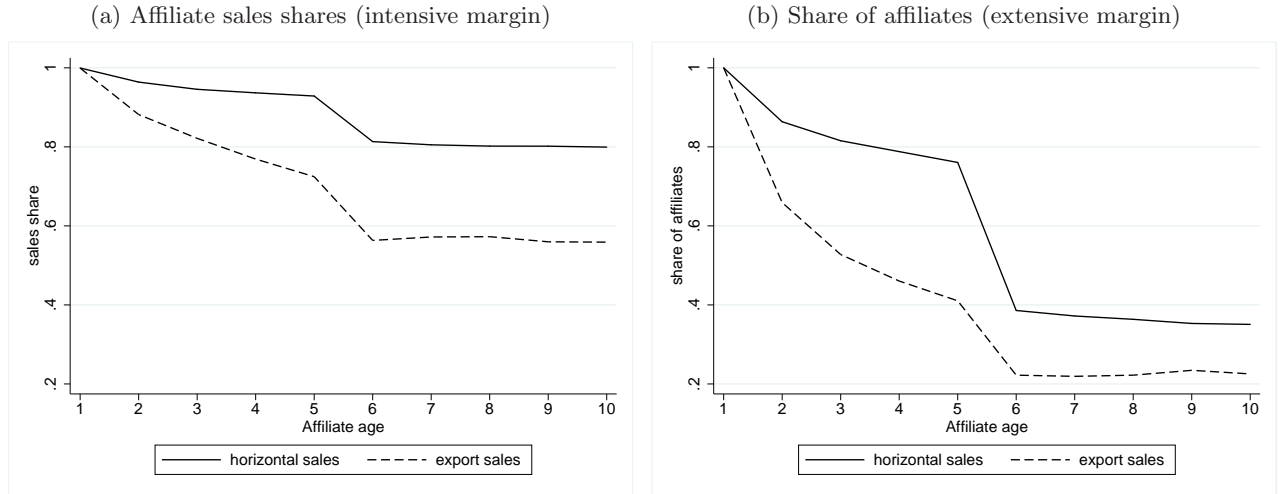
Note: Probability of opening affiliates in the top-ten most popular destinations for US MNEs. Conditional probabilities refer to the probability of observing a MNE opening an affiliate in country i given that the parent already has an affiliate in a “similar” country. The sample is restricted to parents with at least five (ten) affiliates worldwide in the upper (lower) panel. Conditional probabilities in **bold** are not significantly different from the relevant unconditional probability.

Table B.6: MNE shock structure, OLS.

Dependent variable	log of horizontal affiliate sales			
country-industry fixed effect	yes	yes	yes	yes
US GDP	yes	yes	yes	yes
Host country GDP	yes	yes	yes	yes
parent fixed effect	no	yes	yes	no
parent sales	no	no	yes	yes
affiliate fixed effect	no	no	no	yes
Adjusted R-squared	0.24	0.27	0.29	0.79


Notes: Sample of *all* affiliates born during the sample period. Number of observations: 153,773.

Figure B.1: Intensive and extensive margins of sales, by activity type. Pure-type affiliates at birth.



Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Horizontal and export sales refer, respectively, to sales to the market where the affiliate is located, and to sales to markets outside the local market. (B.1a): average sales, as a share of total affiliate sales, include affiliates with positive horizontal and export sales, respectively, for the subset of affiliates with only horizontal and only export sales at birth, respectively). (B.1b): number of affiliates, as a share of the total number of affiliates, include affiliates with positive horizontal and export sales, respectively.

Figure B.2: Sales and affiliate entry age, by activity type.



Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Average horizontal (export) sales shares, by affiliate first age with positive horizontal (export) sales.

C Model's Solution

C.1 Dynamic firm's problem

The compound option structure of the model implies that it can be solved backwards. We start from the problem of a firm that already has an affiliate in country j and has to decide whether to export to any country $k \neq j$. Once this problem is solved, the value of an affiliate in j is determined. Then we can solve the problem of a firm that has to decide whether to open an affiliate in each country j . For each step of the solution, we follow the method outlined in Dixit and Pindyck (1994).

To solve the affiliate export problem, we start by solving for the value functions $V_{jk}^o(z, X, \mathbf{Q})$ and $V_{jk}^o(z, X, \mathbf{Q})$ in their continuation region. Writing the Bellman equation in (14) that describes the value of the option to export to country k for a firm with an affiliate in country j in the continuation region, taking the limit as $\Delta t \rightarrow 0$, and applying Ito's Lemma, yields the no-arbitrage condition

$$\rho V_{jk}^o(z, Y_k) = \tilde{\mu}_k Y_k V_{jk}^{'o}(z, Y_k) + \frac{\tilde{\sigma}_k^2}{2} Y_k^2 V_{jk}^{''o}(z, Y_k), \quad (\text{C.1})$$

where we acknowledge that the option value of exporting to k only depends on the realization of the composite shock Y_k . Guessing a solution for the value function and applying the method of undetermined coefficients, the value of the option of exporting to country k for an affiliate in

country j has general solution given by

$$V_{jk}^o(z, Y_k) = A_{jk}^o(z)Y_k^{\alpha_k} + B_{jk}^o(z)Y_k^{\beta_k}, \quad (\text{C.2})$$

where $\alpha_k < 0$ and $\beta_k > 1$ are the roots of $\frac{1}{2}\tilde{\sigma}_k^2\xi^2 + \left(\tilde{\mu}_k - \frac{\tilde{\sigma}_k^2}{2}\right)\xi - \rho = 0$. As $Y_k \rightarrow 0$, the option of exporting becomes worthless, so it must be that $A_{jk}^o(z) = 0$. Conversely, the option of exporting becomes more attractive as Y_k increases, so it must be that $B_{jk}^o(z) > 0$.

Similarly, writing the Bellman equation in (15) that describes the value of exporting to country k from an affiliate in country j in the continuation region, taking the limit as $\Delta t \rightarrow 0$, and applying Ito's Lemma, yields the no-arbitrage condition

$$\rho V_{jk}^e(z, Y_k) = \pi_{jk}(z, Y_k) - f_{jk}^e + \tilde{\mu}_k V_{jk}^{'e}(z, Y_k) + \frac{\tilde{\sigma}_k^2}{2} Y_k^2 V_{jk}^{''e}(z, Y_k). \quad (\text{C.3})$$

Guessing a solution for the value function and applying the method of undetermined coefficients, the value of the option of exporting to country k for an affiliate in country j has general solution given by

$$V_{jk}^e(z, Y_k) = A_{jk}^e(z)Y_k^{\alpha_k} + B_{jk}^e(z)Y_k^{\beta_k} + \frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f_{jk}^e}{\rho}. \quad (\text{C.4})$$

Notice that, as $Y_k \rightarrow 0$, there is value from the possibility of endogenously stopping to export, so it must be that $A_{jk}^e(z) > 0$. Also, as Y_k increases, the value of exports converges to the discounted profit flow (i.e., there is no further expansion option), so it must be that $B_{jk}^e(z) = 0$.

To completely solve the affiliate export problem, we need to solve for the policy functions, which are thresholds for the realizations of the composite shock that induce the affiliate to start and stop exporting. For each country pair (j, k) and for each firm with productivity z , the parameters $B_{jk}^o(z) > 0$, $A_{jk}^e(z) > 0$, and the export entry and exit thresholds, denoted by Y_{jk}^{OE} and Y_{jk}^{EO} , respectively, can be recovered from the following system of value-matching conditions,

$$V_{jk}^o(z, Y_{jk}^{OE}) = V_{jk}^e(z, Y_{jk}^{OE}) - F_{jk}^e \quad \text{and} \quad V_{jk}^o(z, Y_{jk}^{EO}) = V_{jk}^e(z, Y_{jk}^{EO}), \quad (\text{C.5})$$

and smooth-pasting conditions,

$$V_{jk}^{'o}(z, Y_{jk}^{OE}) = V_{jk}^{'e}(z, Y_{jk}^{OE}) \quad \text{and} \quad V_{jk}^{'o}(z, Y_{jk}^{EO}) = V_{jk}^{'e}(z, Y_{jk}^{EO}), \quad (\text{C.6})$$

where $V'(\cdot)$ denotes the derivative of a value function with respect to the composite shock.

To determine the value of an affiliate in j , we still need to solve for the value of horizontal sales. Writing the Bellman equation in (13) that describes the value of horizontal sales for an affiliate in

country j in the continuation region, taking the limit as $\Delta t \rightarrow 0$, and applying Ito's Lemma, yields the no-arbitrage condition

$$\rho V_j^h(z, Y_j) = \pi_{jj}(z, Y_j) - f_j^h + \tilde{\mu}_j Y_j V_j'^h(z, Y_j) + \frac{\tilde{\sigma}_j^2}{2} Y_j^2 V_j''^h(z, Y_j). \quad (\text{C.7})$$

Guessing a solution for the value function and applying the method of undetermined coefficients, the value of horizontal sales for an affiliate in country j has general solution given by

$$V_j^h(z, Y_j) = A_j^h(z) Y_j^{\alpha_j} + B_j^h(z) Y_j^{\beta_j} + \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f_j^h}{\rho}. \quad (\text{C.8})$$

Notice that, as $Y_j \rightarrow 0$, there is value from the possibility of shutting down the affiliate, so it must be that $A_j^h(z) > 0$. As Y_j increases, the value of horizontal sales converges to the discounted profit flow, so it must be that $B_j^h(z) = 0$.

The value of an affiliate in country j , $V_j^a(z, \mathbf{Y})$ is completely characterized up to the option value parameter $A_j^h(z)$:

$$V_j^a(z, \mathbf{Y}) = A_j^h(z) Y_j^{\alpha_j} + \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f_j^h}{\rho} + \sum_{k \in \mathcal{A}_j(z)} \left[\frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f_{jk}^e}{\rho} + A_{jk}^e(z) Y_k^{\alpha_k} \right] + \sum_{k \notin \mathcal{A}_j(z)} \left[B_{jk}^o(z) Y_k^{\beta_k} \right], \quad (\text{C.9})$$

where $\mathcal{A}_j(z)$ denotes the set of export markets in which an affiliate of a firm with productivity z located in country j exports.

To solve the affiliate opening problem, we still need to solve for the option value of opening an affiliate, and for the policy functions. Writing the Bellman equation in (12) that describes the value of the option to open an affiliate in country j in the continuation region, taking the limit for $\Delta t \rightarrow 0$, and applying Ito's Lemma, yields the no-arbitrage condition

$$\rho V_j^o(z, Y_j) = \tilde{\mu}_j Y_j V_j'^o(z, Y_j) + \frac{\tilde{\sigma}_j^2}{2} Y_j^2 V_j''^o(z, Y_j). \quad (\text{C.10})$$

Guessing a solution for the value function and applying the method of undetermined coefficients, the value of the option of opening an affiliate in country j has general solution given by

$$V_j^o(z, Y_j) = A_j^o(z) Y_j^{\alpha_j} + B_j^o(z) Y_j^{\beta_j}. \quad (\text{C.11})$$

As $Y_j \rightarrow 0$, the option of opening an affiliate becomes worthless, so it must be that $A_j^o(z) = 0$. Conversely, the option of opening an affiliate becomes more attractive as Y_j increases, so it must be that $B_j^o(z) > 0$.

Let Y_j^{OH} and Y_j^{HO} denote the thresholds for the realization of the composite shock that induce a firm to open or shut down an affiliate in country j , respectively. It is important to point out that, when a firm decides to open an affiliate in a country, it considers not only the value of its horizontal sales, but also the option value of potential exports to any destination country. For this reason, the value-matching and smooth-pasting conditions that deliver the parameters $A_j^h(z)$, $B_j^o(z)$, and the policy functions Y_j^{OH} and Y_j^{HO} entail tangency conditions linking the option value function $V_j^o(z, Y_j)$ and the total value of the affiliate $V_j^a(z, \mathbf{Y})$,

$$V_j^o(z, Y_j^{OH}) = V_j^a(z, Y_j^{OH}, \mathbf{Y}_{-j}) - F_j^h \quad , \quad V_j^o(z, Y_j^{HO}) = V_j^a(z, Y_j^{HO}, \mathbf{Y}_{-j}), \quad (\text{C.12})$$

$$V_j'^o(z, Y_j^{OH}) = V_j'^a(z, Y_j^{OH}, \mathbf{Y}_{-j}) \quad , \quad V_j'^o(z, Y_j^{HO}, \mathbf{Y}_{-j}) = V_j'^a(z, Y_j^{HO}, \mathbf{Y}_{-j}), \quad (\text{C.13})$$

where \mathbf{Y}_{-j} denotes the vector of composite shocks in countries other than j .

C.2 Price indexes

We now show how to solve for the vector of price indexes P_k , for $k = 1, \dots, N$, and for the laws of motion governing the evolution of affiliate operations over time across countries. As we only consider MNEs from the United States, we compute the price index as an aggregate of the prices associated with transactions of US MNEs and of domestic firms:

$$P_k^{1-\eta} = \lambda_{kkk} P_{kkk}^{1-\eta} + \lambda_{US,kk} P_{US,kk}^{1-\eta} + \sum_{j \neq k} \lambda_{US,jk} P_{US,jk}^{1-\eta}, \quad (\text{C.14})$$

where P_{kkk} denotes the aggregate price of domestic varieties, $P_{US,kk}$ denotes the aggregate price of varieties produced via the horizontal operations of US affiliates in k , and $P_{US,jk}$ denotes the aggregate price of varieties produced by US affiliates in j and exported to k :

$$P_{kkk}^{1-\eta} = \int_{\Omega_{kkk}} \left(\frac{\eta}{\eta-1} \frac{w_k}{z} \right)^{1-\eta} dG_k(z), \quad (\text{C.15})$$

$$P_{US,kk}^{1-\eta} = \int_{\Omega_{US,kk}} \left(\frac{\eta}{\eta-1} \frac{w_k}{zZ} \right)^{1-\eta} dG_{US}(z), \quad (\text{C.16})$$

$$P_{US,jk}^{1-\eta} = \int_{\Omega_{US,jk}} \left(\frac{\eta}{\eta-1} \frac{\tau_{jk} w_j}{zZ} \right)^{1-\eta} dG_{US}(z) \quad (\text{C.17})$$

and Ω_{kkk} , $\Omega_{US,kk}$, $\Omega_{US,jk}$, denote the corresponding sets of varieties produced. $G_k(z)$ denotes the exogenous distribution of productivity of firms from country k .

The price indexes depend directly on the US productivity shock $Z = e^X$ and indirectly —via

the integration sets— on the demand shocks Q_k . Moreover, as shown below, the integration sets themselves depend on the entry and exit thresholds, which in turn depend on the price indexes. To solve the aggregation problem, we appeal to the equivalence result shown in Leahy (1993): when solving the entry and exit problem, each firm takes aggregate prices and the sets of firms operating in each country as given, and does not take into account the effect of its own entry and exit decisions on these variables.

We assume that the mass of firms in each country k , M_k , is constant. The endogenous mass of affiliates of US firms located in j , $M_{US,j}$, is given by continuing plus new affiliates,

$$M'_{US,j} = M_{US,j} \cdot (1 - G_{US}(z_{US,j}^{HO})) + (M_{US} - M_{US,j}) \cdot (1 - G_{US}(z_{US,j}^{OH})), \quad (\text{C.18})$$

where $z_{US,j}^{OH}$ ($z_{US,j}^{HO}$) is the productivity threshold that induces a US firm to open (shut down) an affiliate in j . Notice that $z_{US,j}^{OH}(Y_j)$ ($z_{US,j}^{HO}(Y_j)$) is the inverse of the threshold in the realization of the shock $Y_j^{OH}(z)$ ($Y_j^{HO}(z)$), whose existence is guaranteed by the monotonicity property of the thresholds shown in Proposition 1.

Similarly, the mass of affiliates of US firms located in j that export to k is given by continuing plus new exporters to k ,

$$M'_{US,jk} = M_{US,jk} \cdot (1 - G_{US}(z_{US,jk}^{EO})) + (M_{US} - M_{US,jk}) \cdot (1 - G_{US}(z_{US,jk}^{OE})), \quad (\text{C.19})$$

where $z_{US,jk}^{OE}$ ($z_{US,jk}^{EO}$) is the productivity threshold that induces an affiliate of a US firm in j to start (stop) exporting to k .

C.3 Proofs

In this section, we provide proofs for Propositions 1 and 2. To this end, notice that, when $A_j^e(z) = A_{jk}^e(z) = 0$, the systems of value-matching and smooth pasting conditions (C.5)-(C.6) and (C.12)-(C.13) only identify the entry thresholds Y_j^{OH} , Y_{jk}^{OE} and can be solved in closed form, delivering the expressions in (22) and (23). Additionally, the option value parameters $B_j^o(z)$, $B_{jk}^o(z)$ can be characterized in closed form,

$$B_j^o(z) = \beta_j^{-\beta_j} (\beta_j - 1)^{\beta_j - 1} \cdot \left(\frac{f_j^h + \rho F_j^h}{\rho} - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) \right)^{1 - \beta_j} \cdot \left(\frac{\kappa_{jj}(z)}{\rho - \tilde{\mu}_j} \right)^{\beta_j}, \quad (\text{C.20})$$

$$B_{jk}^o(z) = \beta_k^{-\beta_k} (\beta_k - 1)^{\beta_k - 1} \cdot \left(\frac{f_{jk}^e + \rho F_{jk}^e}{\rho} \right)^{1 - \beta_k} \cdot \left(\frac{\kappa_{jk}(z)}{\rho - \tilde{\mu}_k} \right)^{\beta_k}, \quad (\text{C.21})$$

where $\kappa_{jk}(z) \equiv H(\tau_{jk}w_j/z)^{1-\eta}P_k^\eta\lambda_{jk}$ and $\mathbf{V}_j^E(z, \mathbf{Y}_{-j}) = \sum_{k \in \mathcal{A}_j(z)} \left[\frac{\kappa_{jk}(z)Y_k}{\rho - \tilde{\mu}_k} - \frac{f_{jk}^e}{\rho} \right] + \sum_{k \notin \mathcal{A}_j(z)} \left[B_{jk}^o(z)Y_k^{\beta_k} \right]$.

Equations (C.20) and (C.21) reveal that the option value of opening an affiliate (exporting from an affiliate) is decreasing in both the fixed and sunk costs of affiliate opening (exporting). In addition, the option value of opening an affiliate is increasing in the value of the potential export network of the affiliate, highlighting the effects of the compound option mechanism. Finally, both the option values of opening and exporting from an affiliate are increasing in firm productivity z .

Proof of Proposition 1. Affiliate's variable profits and the value of an affiliate's export network are increasing in firm productivity,

$$\frac{\partial \kappa_{jk}(z)}{\partial z} > 0, \quad (\text{C.22})$$

$$\frac{\partial \mathbf{V}_j^E(z, \mathbf{Y}_{-j})}{\partial z} > 0. \quad (\text{C.23})$$

Given (C.22) and (C.23), the proof for the affiliate export entry threshold is immediate from taking derivatives in (23),

$$\frac{\partial Y_{jk}^{OE}(z)}{\partial z} = \underbrace{\left(\frac{\beta_k}{\beta_k - 1} \right) \cdot \left(\frac{f_{jk}^e + \rho F_{jk}^e}{\rho} \right) \cdot (\rho - \tilde{\mu}_k)}_{\geq 0} \cdot \underbrace{\left(\frac{1}{-\kappa_{jk}(z)^2} \right)}_{< 0} \cdot \underbrace{\frac{\partial \kappa_{jk}(z)}{\partial z}}_{> 0} \leq 0.$$

We compute the derivative in the scenario in which $(f_j^h + \rho F_j^h)/\rho - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) > 0$, so that $Y_j^{OH}(z) > 0$ and the entry problem is well-defined. Taking derivative in (23) yields

$$\frac{\partial Y_j^{OH}(z)}{\partial z} = \underbrace{\left(\frac{\beta_j}{\beta_j - 1} \right)}_{> 0} \cdot \left\{ \underbrace{\frac{-\partial \mathbf{V}_j^E(z, \mathbf{Y}_{-j})}{\partial z} \cdot \frac{\rho - \tilde{\mu}_j}{\kappa_{jj}(z)}}_{< 0} + \underbrace{\left(\frac{f_j^h + \rho F_j^h}{\rho} - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) \right)}_{\geq 0} \cdot \underbrace{\left(\frac{\rho - \tilde{\mu}_j}{-\kappa_{jj}(z)^2} \right) \cdot \frac{\partial \kappa_{jj}(z)}{\partial z}}_{< 0} \right\} \leq 0.$$

Proof of Proposition 2. Affiliate's variable profits are increasing in the taste shifter,

$$\frac{\partial \kappa_{jj}(z)}{\partial \lambda_{jj}} > 0. \quad (\text{C.24})$$

We compute the derivative in the scenario in which $(f_j^h + \rho F_j^h)/\rho - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) > 0$, so that $Y_j^{OH}(z) > 0$ and the entry problem is well-defined. Given (C.24), the proof is immediate from

taking derivatives in (22),

$$\frac{\partial Y_j^{OH}(z)}{\partial \lambda_{jj}} = \underbrace{\frac{\beta_j}{\beta_j - 1} \cdot \left(\frac{f_j^h + \rho F_j^h}{\rho} - \mathbf{V}_j^E(z, \mathbf{Y}_{-j}) \right)}_{\geq 0} \cdot \underbrace{\left(\frac{\rho - \tilde{\mu}_j}{-\kappa_{jj}(z)^2} \right)}_{< 0} \cdot \underbrace{\frac{\partial \kappa_{jj}(z)}{\partial \lambda_{jj}}}_{> 0} \leq 0.$$

Additionally, taking derivatives with respect to F_j^h in (22) yields

$$\frac{\partial Y_j^{OH}(z)}{\partial F_j^h} = \frac{\beta_j}{\beta_j - 1} \cdot \left(\frac{\rho - \tilde{\mu}_j}{\kappa_{jj}(z)} \right) > 0.$$

D Calibration: Additional Results

Table D.1: Calibrated parameters: shock processes and wages.

	$\tilde{\mu}_j$	$\tilde{\sigma}_j$	γ_j	$Q_j(0)$	w_j
Brazil	0.051	0.130	0.032	0.096	0.711
Canada	0.030	0.136	0.661	0.073	0.831
China	0.064	0.122	0.035	1.162	0.307
France	0.025	0.127	0.370	0.125	0.981
United Kingdom	0.025	0.136	0.314	0.128	0.930
Germany	0.026	0.128	0.648	0.222	0.825
Ireland	0.048	0.144	0.560	0.003	1.188
Japan	0.027	0.133	0.383	0.354	0.719
Mexico	0.036	0.127	0.083	0.078	0.699
Singapore	0.080	0.137	0.213	0.003	0.950
United States	0.028	0.116	1.000	1.000	1.000

Note: $\tilde{\mu}_j$ and $\tilde{\sigma}_j$ refer to the drift and standard deviation, respectively, of the composite shock Y_j . γ_j refers to the correlation between the US aggregate productivity shock and country j 's aggregate demand shock. $Q_j(0)$ denotes the initial value of the demand shock process in country j . w_j refers to country j 's wage relative to the US.

Table D.2: Calibrated parameters: fixed costs, sunk costs, and taste shifters.

	f_j^h	F_j^h	$f_{j,US}^e$	$F_{j,US}^e$	f_{jk}^e	F_{jk}^e	λ_j
Brazil	0.0113	0.0023	0.0028	0.0000	0.0023	0.00002	0.051
Canada	0.0121	0.0002	0.0040	0.0008	0.0024	0.0000	0.161
China	0.0016	0.0000	0.0012	0.0000	0.0009	0.0000	0.013
France	0.0165	0.0007	0.0023	0.0000	0.0053	0.0022	0.097
United Kingdom	0.0122	0.0000	0.0023	0.0003	0.0044	0.0003	0.104
Germany	0.0167	0.0004	0.0025	0.0001	0.0060	0.0024	0.088
Ireland	0.0027	0.0003	0.0015	0.0000	0.0014	0.0007	0.010
Japan	0.0327	0.0024	0.0042	0.0000	0.0051	0.0001	0.097
Mexico	0.0047	0.0001	0.0014	0.0001	0.0010	0.00001	0.020
Singapore	0.0029	0.00002	0.0033	0.0002	0.0028	0.0009	0.013

Note: f_j^h (F_j^h) is the fixed (sunk) cost of opening an affiliate in country j . $f_{j,US}^e$ ($F_{j,US}^e$) is the fixed (sunk) cost of exporting from j to the United States. f_{jk}^e (F_{jk}^e) is the fixed (sunk) cost of exporting from j to a destination k other than the United States. λ_j is the taste shifter associated with goods produced by US MNEs in country j .

Table D.3: Calibrated parameters: bilateral iceberg-trade costs.

	BRA	CAN	CHN	FRA	GBR	GER	IRL	JPN	MEX	SGP	USA
BRA	1.000	2.579	1.000	2.470	3.467	2.419	2.360	2.551	1.825	2.684	5.415
CAN	1.675	1.000	1.000	1.828	2.516	1.780	1.684	1.821	1.221	2.204	3.230
CHN	7.933	9.552	1.000	7.663	11.446	7.301	7.338	6.313	6.747	6.125	16.55
FRA	2.265	1.773	1.000	1.000	1.409	1.138	1.311	1.544	2.102	2.625	3.638
GBR	2.569	1.620	1.000	1.244	1.000	1.328	1.286	1.375	2.359	2.976	3.369
GER	2.829	2.199	1.059	1.450	1.573	1.000	1.691	1.815	2.608	3.205	4.090
IRL	3.404	2.009	1.312	2.060	1.737	2.085	1.000	1.745	3.090	3.989	4.120
JPN	2.174	2.565	1.000	2.166	2.984	2.070	2.062	1.000	1.776	1.733	5.236
MEX	2.411	3.171	1.105	3.030	3.659	2.946	2.832	2.753	1.000	3.345	5.685
SGP	1.863	2.170	1.000	1.985	2.093	1.903	1.919	1.545	1.758	1.000	4.145
USA	2.359	2.481	1.264	2.002	2.057	2.079	2.231	2.162	2.056	2.578	1.000

Note: Trade costs from country j (rows) to destination k (columns).

Table D.4: Static moments: intensive margin. Model versus data.

Share of:	Affiliate sales to host market		Affiliate sales to the US		Affiliate sales to third countries	
	data	model	data	model	data	model
Brazil	0.018	0.018	0.072	0.072	0.142	0.140
Canada	0.048	0.048	0.261	0.261	0.113	0.111
China	0.003	0.003	0.099	0.099	0.219	0.219
France	0.038	0.039	0.075	0.075	0.364	0.364
United Kingdom	0.052	0.053	0.111	0.111	0.371	0.366
Germany	0.047	0.048	0.092	0.092	0.413	0.413
Ireland	0.002	0.002	0.242	0.242	0.515	0.605
Japan	0.034	0.034	0.047	0.047	0.130	0.130
Mexico	0.011	0.011	0.173	0.173	0.128	0.128
Singapore	0.003	0.003	0.222	0.222	0.488	0.488
Average	0.026	0.026	0.139	0.139	0.288	0.296

Note: Affiliate sales to the host market are expressed as a share of the parent's US sales. Affiliate sales to the US and to third countries are expressed as a share of affiliate sales in the host market. Calculations are conditional on affiliate entry, but unconditional on affiliate exports. Averages across years.

Table D.5: Static moments: intensive margin. Selected destinations. Model versus data.

Share of:	Affiliate sales to Canada		Affiliate sales to the United Kingdom		Affiliate sales to Japan	
	data	model	data	model	data	model
Brazil	0.008	0.015	0.008	0.005	0.004	0.005
Canada	n.a.	n.a.	0.006	0.005	0.003	0.005
China	0.008	0.004	0.002	0.003	0.037	0.046
France	0.010	0.006	0.091	0.122	0.006	0.005
United Kingdom	0.012	0.011	n.a.	n.a.	0.009	0.010
Germany	0.009	0.003	0.079	0.136	0.010	0.003
Ireland	0.053	0.042	0.386	0.395	0.082	0.037
Japan	0.006	0.006	0.001	0.005	n.a.	n.a.
Mexico	0.007	0.011	0.001	0.010	0.002	0.009
Singapore	0.024	0.024	0.047	0.106	0.147	0.116
Average	0.015	0.014	0.069	0.087	0.033	0.026

Note: Affiliate sales to destination j are expressed as share of sales in the affiliate host market. Calculations are conditional on affiliate entry, but unconditional on affiliate exports. Averages across (benchmark) years.

Table D.6: Static moments: extensive margin. Model versus data.

Share of:	MNEs with affiliates in j		Affiliates in j exporting to the US		Affiliates in j exporting to third countries	
	data	model	data	model	data	model
Brazil	0.198	0.189	0.515	0.515	0.674	0.668
Canada	0.544	0.539	0.725	0.722	0.478	0.478
China	0.184	0.181	0.382	0.382	0.548	0.545
France	0.312	0.310	0.539	0.542	0.747	0.746
United Kingdom	0.554	0.560	0.605	0.605	0.739	0.748
Germany	0.367	0.364	0.608	0.608	0.760	0.760
Ireland	0.122	0.122	0.575	0.574	0.760	0.751
Japan	0.155	0.150	0.468	0.468	0.578	0.541
Mexico	0.302	0.305	0.647	0.649	0.494	0.498
Singapore	0.129	0.113	0.597	0.596	0.724	0.721
Average	0.287	0.283	0.566	0.566	0.650	0.646

Note: MNEs with affiliates in j are expressed as shares of the total number of US MNEs. Exporting affiliates are expressed as shares of the total number of affiliates in j . Calculations are conditional on affiliate entry. Averages across years.

Table D.7: Dynamic moments: entry. Model versus data.

Share of:	MNEs opening affiliates in j		Affiliates in j that start exporting to:			
	data	model	the United States		third countries	
			data	model	data	model
Brazil	0.021	0.012	0.032	0.023	0.037	0.028
Canada	0.060	0.031	0.024	0.018	0.036	0.028
China	0.029	0.017	0.027	0.019	0.040	0.031
France	0.038	0.022	0.033	0.032	0.025	0.020
United Kingdom	0.069	0.048	0.029	0.021	0.027	0.027
Germany	0.046	0.026	0.030	0.030	0.025	0.019
Ireland	0.015	0.006	0.037	0.030	0.030	0.043
Japan	0.020	0.014	0.032	0.023	0.030	0.028
Mexico	0.036	0.024	0.029	0.025	0.033	0.022
Singapore	0.019	0.005	0.026	0.019	0.028	0.022
Average	0.035	0.021	0.030	0.024	0.031	0.027

Note: MNEs opening affiliates in j are expressed as shares of the total number of US MNEs in the period before entry. Affiliates that start exporting are expressed as shares of the total number of affiliates in j in the period before entry. Calculations are conditional on affiliate entry. Averages across years.

Table D.8: Dynamic moments: exit. Model versus data.

Share of:	MNEs shutting down affiliates in j		Affiliates in j that stop exporting to			
	data	model	the United States		third countries	
			data	model	data	model
Brazil	0.102	0.074	0.027	0.043	0.032	0.037
Canada	0.125	0.083	0.021	0.022	0.028	0.058
China	0.082	0.105	0.020	0.054	0.031	0.051
France	0.117	0.088	0.029	0.058	0.021	0.018
United Kingdom	0.128	0.126	0.026	0.031	0.023	0.037
Germany	0.122	0.088	0.029	0.048	0.024	0.017
Ireland	0.119	0.046	0.030	0.055	0.030	0.053
Japan	0.111	0.084	0.029	0.048	0.029	0.046
Mexico	0.114	0.088	0.021	0.032	0.026	0.038
Singapore	0.115	0.044	0.023	0.029	0.026	0.014
Average	0.113	0.083	0.025	0.042	0.027	0.037

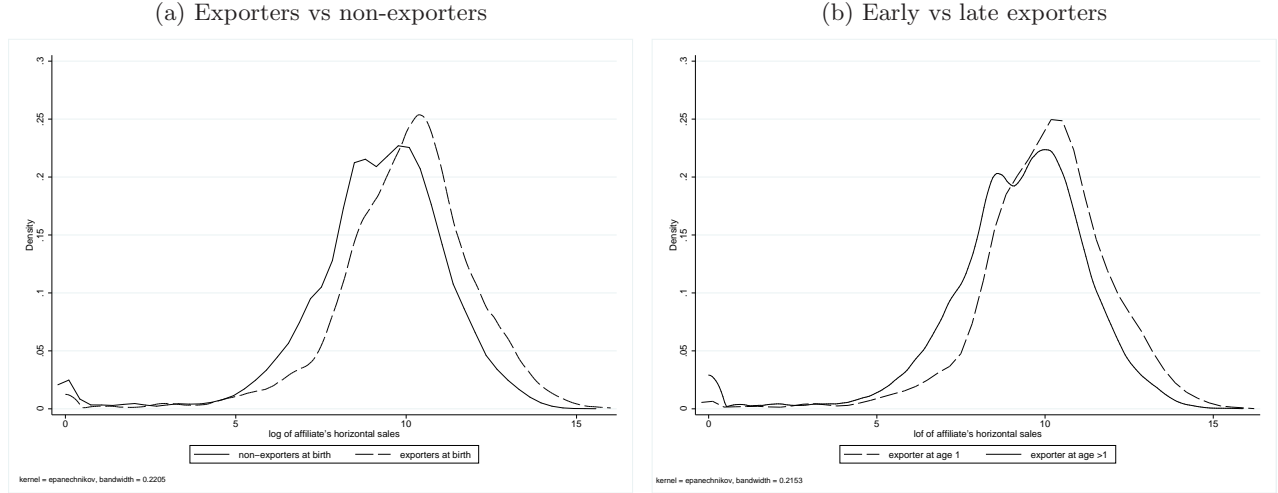
Note: MNEs shutting down affiliates in j are expressed as shares of the total number of affiliates in j in the period before exit. Affiliates that stop exporting are expressed as shares of the total number of affiliates in j that export in the period before entry. Calculations are conditional on affiliate entry. Averages across years.

Table D.9: Exporter and early-exporter size advantage. Model versus data.

	Exporter size advantage		Early exporter size advantage	
	data	model	data	model
Brazil	6.31	7.43	3.49	5.55
Canada	3.39	6.81	2.52	5.56
China	7.84	7.57	3.08	5.71
France	4.46	7.08	1.95	5.81
United Kingdom	1.93	6.91	1.52	5.56
Germany	5.47	6.97	4.24	5.56
Ireland	8.02	6.09	8.57	5.22
Japan	12.49	7.06	2.31	5.39
Mexico	4.21	7.00	2.08	5.67
Singapore	8.59	6.78	7.04	5.38
Average	6.27	6.97	3.68	5.54

Note: Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Exporter size advantage refers to the average size of exporting MNE affiliates relative to the average size of non-exporting MNE affiliates, an average across countries and years. Early-exporter size advantage refers to the average size of exporting MNE affiliates that start exports early in life relative to the average size of exporting MNE affiliates that start exports later in life. Size refers to horizontal affiliate sales; early versus late exporters refers to affiliates that are born with exports versus the ones that start exporting later. Calculations in the calibrated model trim the upper and lower 10th decile of the simulated firm-level data.

Figure D.1: Exporter and early-exporter size advantage, OLS.



Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Kernel density of log horizontal sales for affiliates that: are born with exclusively horizontal sales (non-exporters) and those with exports (exporters), in (D.1a); start exporting in their first year of life and those that start after their first year of life, in (D.1b).

Table D.10: First-affiliate size and cost advantage, OLS.

Dependent variable	probability of being the first affiliate of a MNE			
	(1)	(2)	(3)	(4)
Log of horizontal sales	0.013*** (0.003)	0.012*** (0.003)	0.012*** (0.003)	0.011*** (0.003)
Number of business procedures		-0.006*** (0.001)		-0.008*** (0.002)
Cost of starting business (% of GDP p.c.)			-0.006*** (0.0002)	-0.0004* (0.0002)
Obs	36,127	36,127	36,127	36,127
R-squared	0.291	0.295	0.293	0.297

Note: Observations at the affiliate-year level, for new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. The variables “Number of business procedures” and “Cost of starting a business” are from the World Bank’s Doing Business dataset. All specifications include host country GDP (from Penn World Tables, 9.0), year fixed effects, and parent fixed effects. Standard errors, clustered at the parent level, are in parenthesis. Levels of significance are denoted *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

Table D.11: Calibrated MNE costs, as share of sales, by country.

Sales percentiles	Sunk costs F_j^h (% of parent US sales)			Fixed costs f_j^h (% of horizontal sales)		
	5th	50th	95th	5th	50th	95th
Brazil	0.189	0.068	0.007	16.25	9.960	1.530
Canada	0.023	0.014	0.002	22.74	14.36	2.402
China	2.e-04	1.e-04	0.000	11.82	7.013	2.778
France	0.075	0.035	0.005	22.46	15.81	2.547
United Kingdom	2.e-04	1.e-04	0.000	23.93	14.90	2.386
Germany	0.037	0.021	0.003	22.60	15.52	2.483
Ireland	0.009	0.006	0.001	18.92	10.85	1.822
Japan	0.163	0.066	0.009	22.63	16.30	2.388
Mexico	0.010	0.003	4.e-04	17.86	12.68	2.269
Singapore	0.001	3.e-04	0.000	7.740	4.432	2.849
Average	0.051	0.021	0.003	18.69	12.18	2.345

Note: Sales evaluated at the year of affiliate entry.

E Quantitative Analysis: Additional Results

Table D.12: Calibrated MNE export costs, as share of sales, by country.

Sunk export costs F_{jk}^e (% of horizontal affiliate sales)							Fixed export costs f_{jk}^e (% of average affiliate exports)					
to United States			to other countries				to United States			to other countries		
Sales percentiles	5th	50th	95th	5th	50th	95th	5th	50th	95th	5th	50th	95th
Brazil	0.009	0.003	0.001	0.049	0.0172	0.002	11.89	9.569	1.695	12.07	11.85	3.365
Canada	1.499	0.765	0.121	0.000	0.000	0.000	11.84	9.822	1.604	14.45	12.52	3.657
China	0.027	0.013	0.003	0.020	0.008	0.001	12.96	9.511	1.751	12.67	11.39	3.371
France	0.008	0.003	4e-04	4.7667	1.667	0.206	12.13	9.901	1.466	12.17	11.69	3.024
United Kingdom	0.423	0.203	0.029	0.660	0.245	0.036	12.06	10.11	1.614	12.74	10.31	3.708
Germany	0.139	0.046	0.006	6.587	1.782	0.252	12.88	10.09	1.433	12.35	11.75	2.214
Ireland	0.000	0.000	0.000	11.28	3.706	0.634	13.92	10.88	1.894	11.68	10.01	3.527
Japan	0.002	4.e-04	1.e-04	0.102	0.019	0.003	13.71	11.09	1.99	13.74	13.14	4.12
Mexico	0.596	0.222	0.029	0.031	0.011	0.002	12.95	10.60	1.493	14.03	11.35	2.892
Singapore	0.580	0.328	0.086	6.059	2.404	0.705	12.16	9.653	1.708	12.39	12.02	3.207
Average	0.328	0.158	0.028	2.956	0.986	0.184	12.65	10.13	1.666	12.83	11.60	3.309

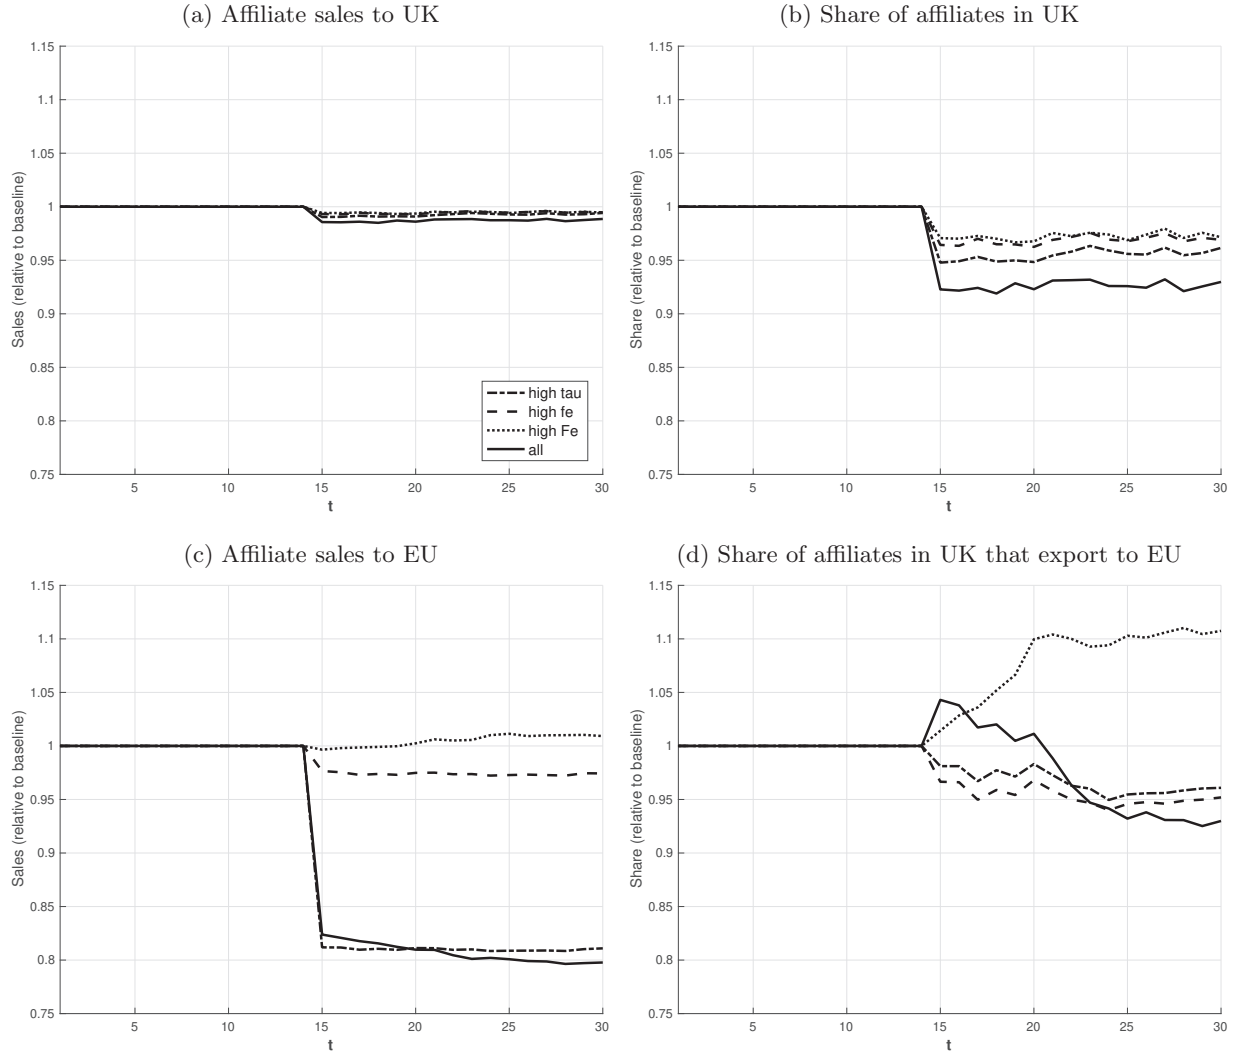
Note: Sales evaluated at the year of first export.

Table D.13: MNE sales distribution, in thousands of dollars, by country. BEA data.

	Parent sales at age 1 of the affiliate	Affiliate horizontal sales		Affiliate export sales, avg across years	
		avg across years	at age of first export	total	to US
Brazil	1,194,252	17,013	14,404	2,814	576
Canada	475,426	14,085	12,871	4,916	2,453
China	887,290	10,329	5,024	5,025	738
France	709,151	17,443	14,468	7,586	622
Germany	657,868	19,761	15,358	12,370	937
United Kingdom	442,758	12,505	10,202	6,203	666
Ireland	974,741	5,892	4,731	16,317	1,371
Japan	575,766	15,017	11,187	2,058	374
Mexico	757,420	8,042	6,486	4,221	1,285
Singapore	859,366	5,097	3,927	8,692	705
Average	753,404	12,518	9,866	7,020	973

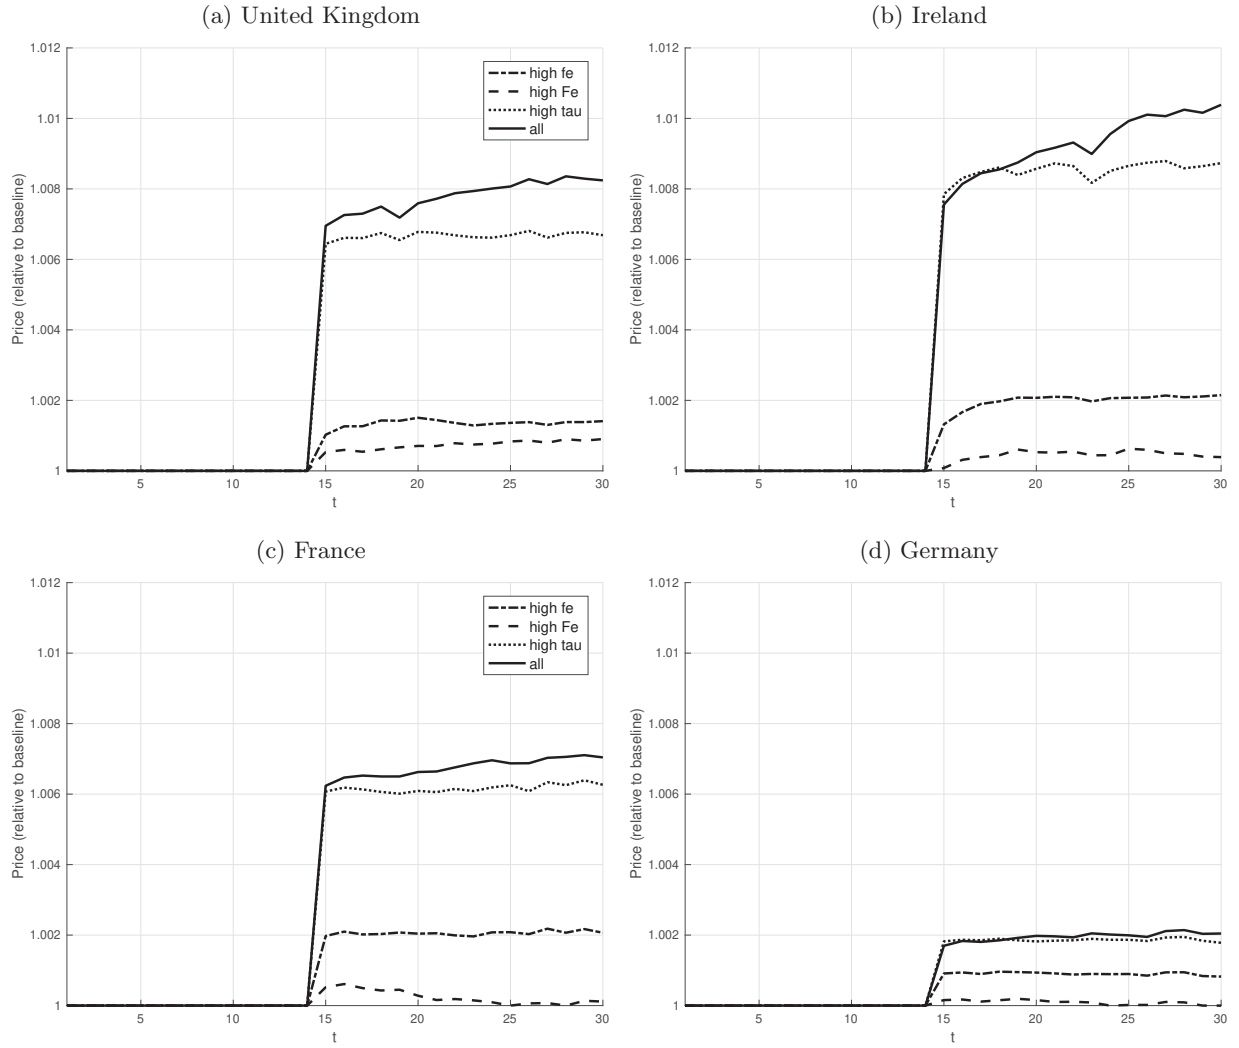
Note: Sales refer to the median firm and are in thousands of US dollars. For confidentiality purposes, magnitudes are averages across the nine observations around the median.

Figure E.1: Brexit: US affiliates in the United Kingdom. Exogenous price indexes.



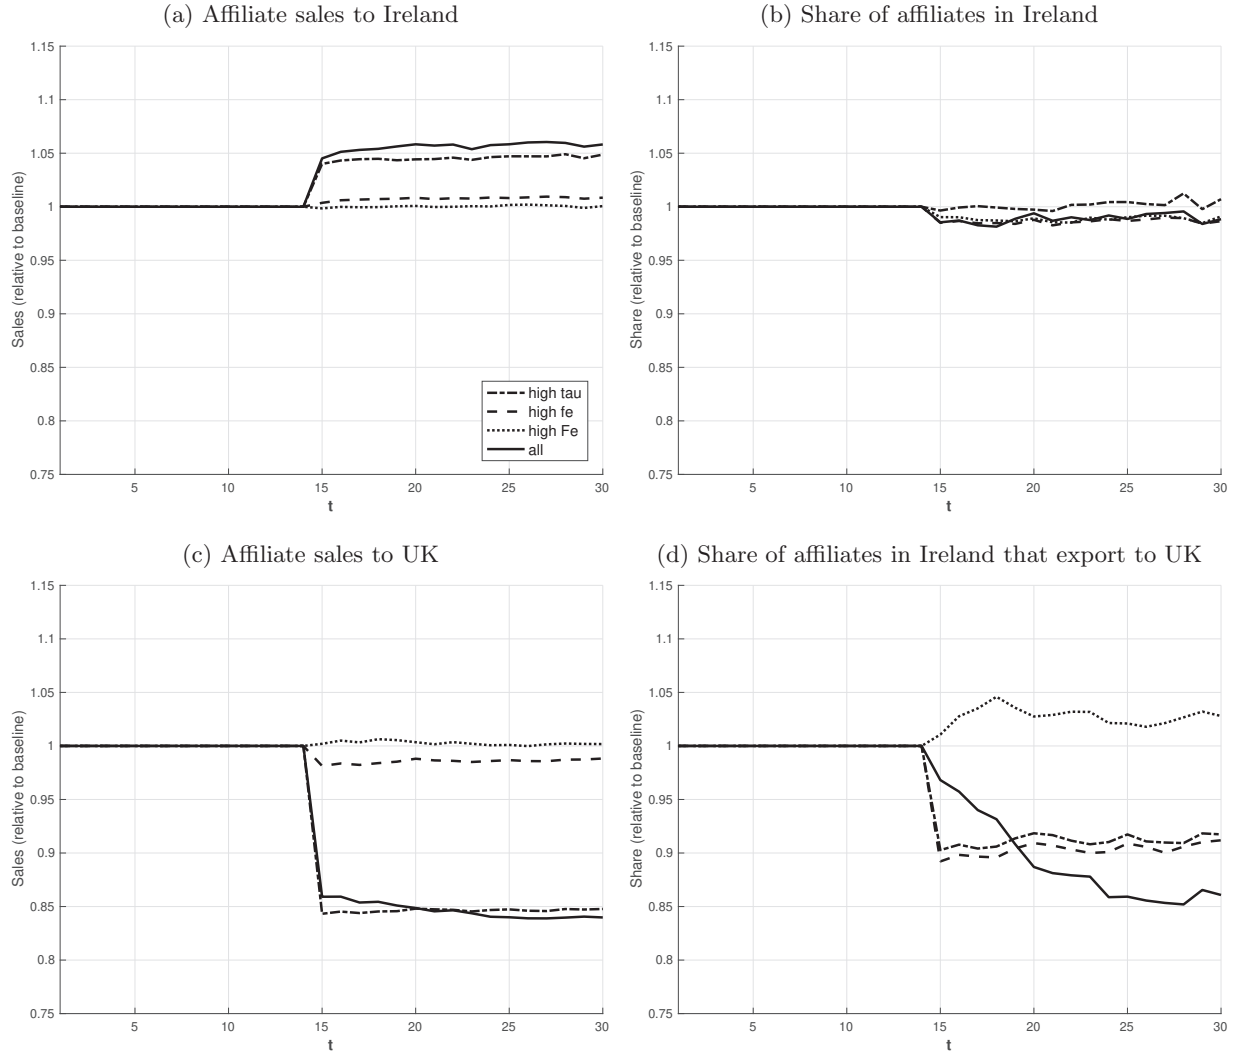
Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

Figure E.2: Brexit: Price changes.



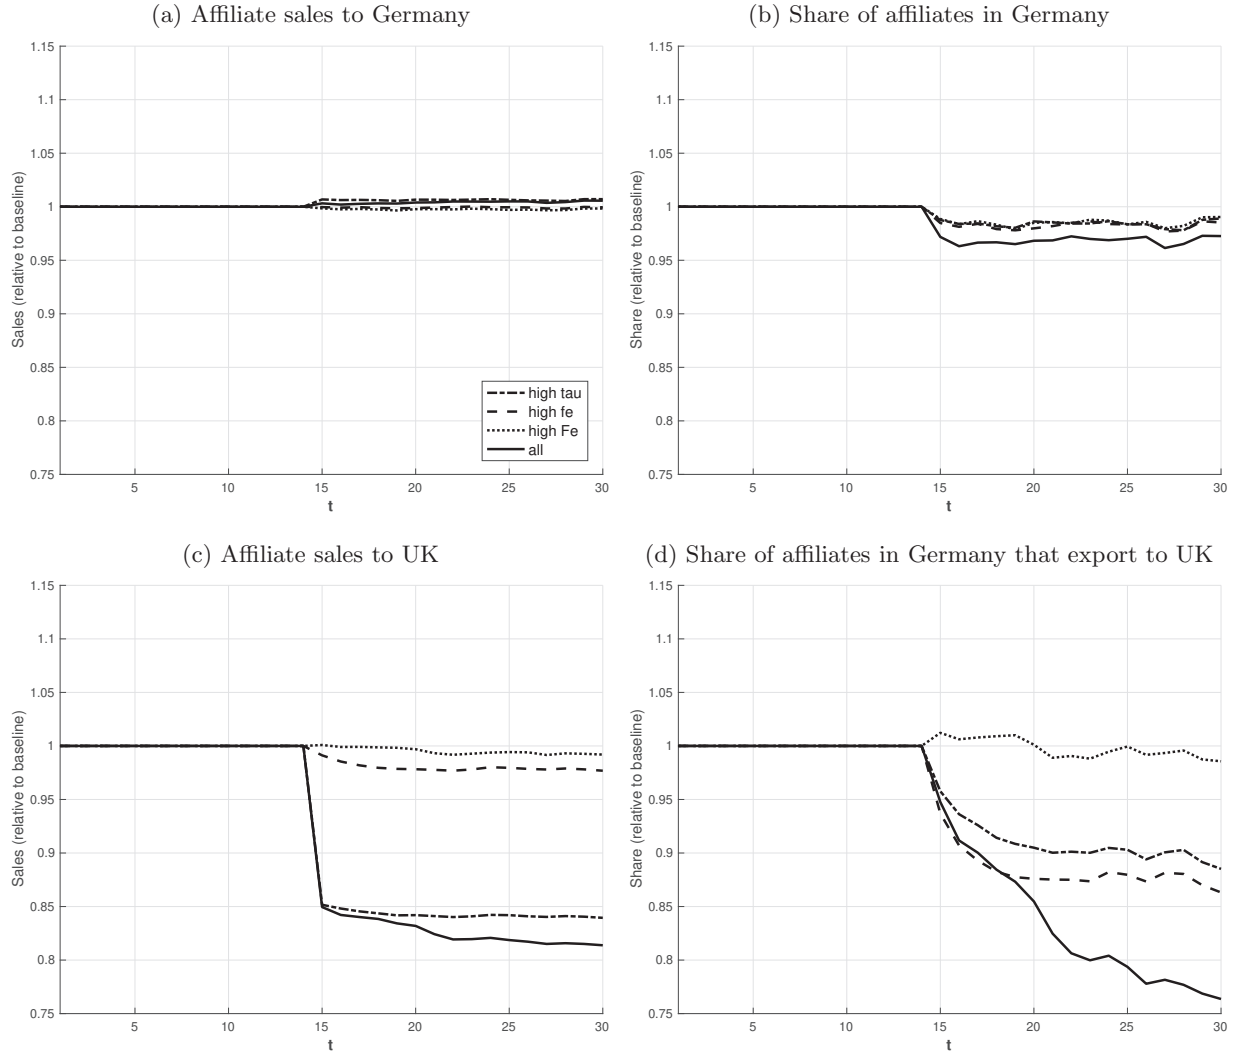
Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

Figure E.3: Brexit: US affiliates in Ireland.



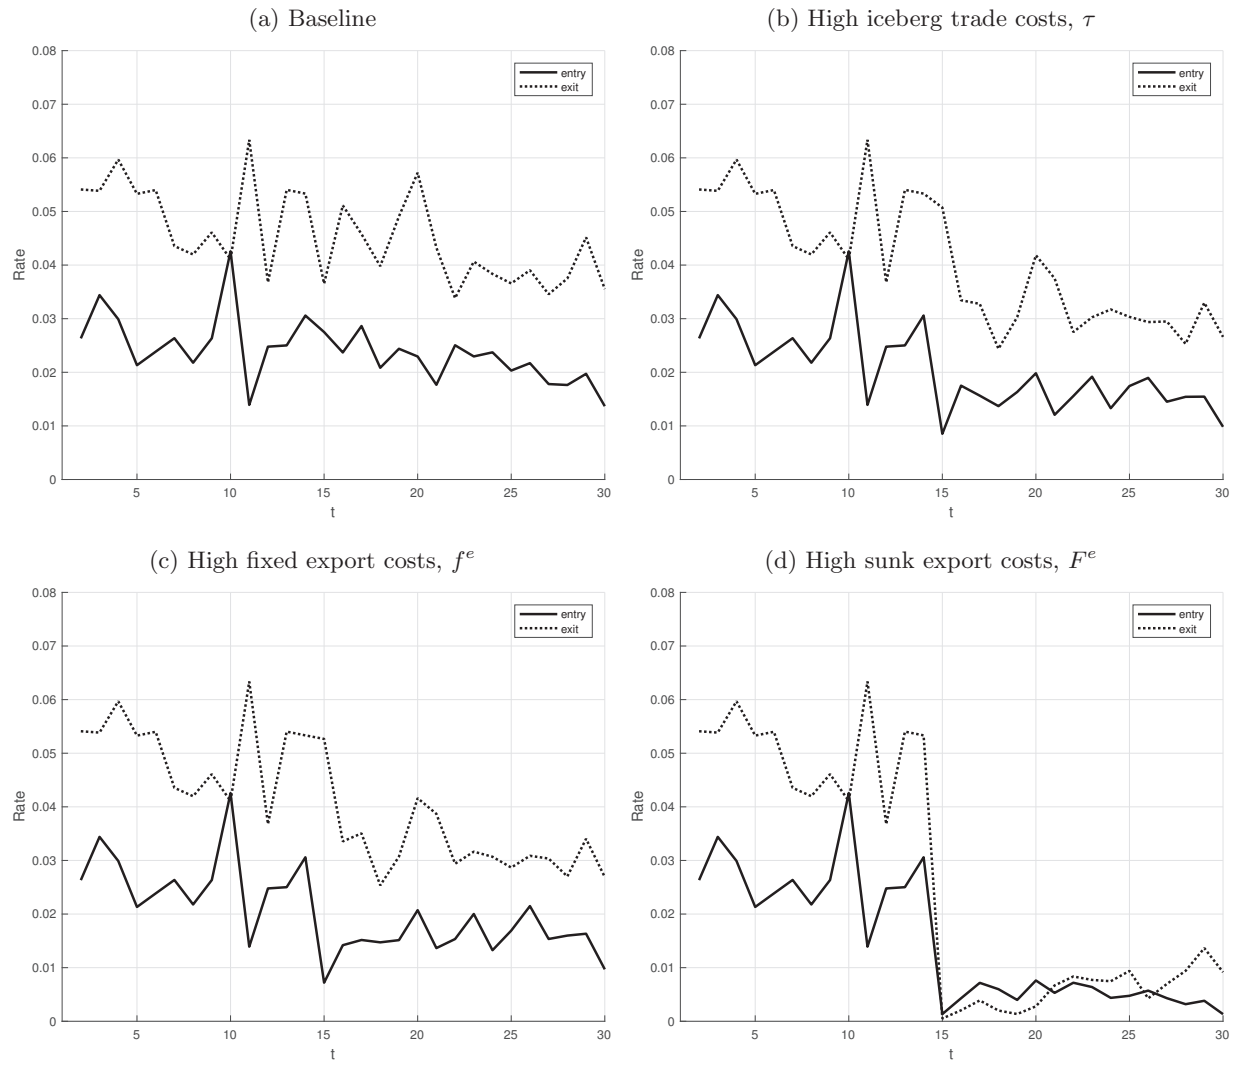
Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

Figure E.4: Brexit: US affiliates in Germany.



Note: “high X” refers to an increase in the barrier X from/to United Kingdom to/from country j (“shallow Brexit”). “All” refers to increasing all three export barriers from/to the United Kingdom to/from j at once (“deep Brexit”). Country j refers to Ireland, Germany, and France.

Figure E.5: Brexit: Entry and exit from the United Kingdom into EU.



Note: Price indexes are endogenous.