Multinational Expansion in Time and Space*

Stefania Garetto
BU, CEPR, and NBER

Lindsay Oldenski
Georgetown University

Natalia Ramondo†
BU and NBER

October 4, 2021

Abstract

This paper studies the expansion patterns of multinational enterprises (MNEs) in time and space. Informed by a set of facts that we document using a panel of US MNEs, we develop a multi-country dynamic model of MNEs featuring a rich structure of costs to MNE expansion into host and export markets. We introduce a novel compound-option formulation of the MNE problem that captures in a tractable way the rich heterogeneity observed in the data. Using the calibrated model, our quantitative application reveals that the nature of frictions to MNE activities in time and space matters for predicting the effects of globalization shocks.

JEL Codes: F1.

Key Words: Multinational firms, foreign direct investment, firm dynamics, sunk costs.

*We thank our discussants Javier Cravino, Ana Maria Santacreu, Michael Sposi, and Stephen Yeaple for very helpful comments. We have also benefited from comments from George Alessandria, Costas Arkolakis, Ariel Burstein, Jonathan Eaton, Oleg Itskhoki, Sam Kortum, Andrei Levchenko, Eduardo Morales, Ezra Oberfield, and Andrès Rodríguez-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, with output disclosed under BEA Approval Numbers BEA-FY21-E0011-R0002 to R0004 and other unnumbered approved disclosure requests. The views expressed are those of the authors and do not reflect official positions of the US Department of Commerce.

†E-mail: garettos@bu.edu; lindsay.oldenski@georgetown.edu; nramondo@bu.edu.

1 Introduction

Many important questions in international economics involve the complex activities of multinational enterprises (MNEs) across time and over space. Consider the recent rise in protectionism worldwide, of which the debate on the United Kingdom abandoning the European Union, “Brexit”, is one example. Under Brexit, would MNEs pull out from the United Kingdom? Would MNEs located in the EU be affected? How would different implementations of Brexit affect MNEs’ expansion patterns? And how would long-run changes differ from short-run adjustments? Providing sound answers to these and other similar questions requires an understanding of the dynamics of MNE expansion across locations and of the nature of the costs these firms face.

Despite their importance, the behavior of MNEs and their affiliates across 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 multi-country dynamic settings where MNEs can separate the locations of production and sales. This paper contributes to 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 foreign affiliates of US MNEs. The model is aimed at answering quantitatively questions about the effects of policy shocks on MNE expansion, which require both a rich spatial and dynamic structure.

Our analysis uses a long panel of US MNEs and their foreign affiliates from the Bureau of Economic Analysis (BEA).[2] We start by documenting a set of facts about the behavior of MNEs and their affiliates in time and space. First, almost all MNE affiliates are born with sales to the host market, which remains the main destination market of the affiliate; exports start later in the affiliate’s life. Second, affiliate sales, relative to parent sales, are flat over the affiliate’s life, except for a one-time jump in the year when the affiliate starts exporting. Third, the probability of MNE entry into a market is indistinguishable from that probability conditional on having entered similar markets (e.g. markets in the same continent or with a similar income level)—a pattern we refer to as weak “extended gravity”. Finally, sales of existing affiliates to their host market do not decrease when sibling affiliates (i.e. affiliates of the same MNE located in a different country) start exporting to that market and when new sibling affiliates are created in a different host market.

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

[2] US MNEs and their foreign affiliates are an empirically relevant sample 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 (Yeaple, 2013).

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.

While the static components of our model are standard and follow Melitz (2003), the way we formulate the dynamic problem of the MNE 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). Concretely, our model is based on a *compound option* structure: opening an affiliate in a country is an option, which, if exercised, gives access to a set of additional options, such as exporting from the affiliate to any other location. Since we observe that almost all affiliates in the data have some horizontal sales at birth, we further assume that firms that decide to do Foreign Direct Investment (FDI) must set up an affiliate and sell to the local market prior to considering exporting from that affiliate. Finally, we couple this dynamic structure with a standard Armington assumption stating that varieties are firm and location specific. This assumption is consistent with the observation that horizontal affiliate sales and export sales from sibling affiliates to a market co-exist.

Together, these ingredients are key for achieving tractability of the model while, at the same time, preserving the rich heterogeneity necessary for quantitative analysis. First, the compound option structure allows us to introduce dynamic interdependence in the location choices of the MNE: the decision to open an affiliate in a country depends on the set of countries that the affiliate can export to. Hence, shocks to the cost of operating, or establishing, an export network from a given market will affect the MNE entry decisions into that market. Second, due to the continuous-time formulation of the real-option problem, the value functions can be solved in closed form as simple additive functions of the firm's realized profit flow, the option value of expansion, and the option value of exit. The tractability of the firm's problem makes it possible to aggregate firms' outcomes and solve for the evolution of the price index in each country. Finally, the Armington assumption implies that the MNE entry decision into a market is separable across locations, consistent with our evidence of weak extended gravity in affiliate entry.[3] Nonetheless, because the price indices endogenously change with trade barriers, our model displays the substitution forces of the proximity-concentration tradeoff at the aggregate level, as first documented by Brainard

[3] In terms of the implications for affiliate entry, the Armington assumption is *isomorphic* to assuming that firm productivity is a random variable drawn from a multivariate Pareto distribution independent over varieties and locations, as in Arkolakis et al. (2018).

(1997).[4]

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. Affiliate 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 type of frictions. The calibrated model is also able to reproduce non-targeted observations related to the geography of MNE sales, which we document in the BEA data. We document and match selection patterns of MNE affiliates within and across host markets, as well as sorting patterns of MNE entry across host markets, in the spirit of the facts documented in Eaton et al. (2011) for French exporters.

Armed with the calibrated model, we perform various counterfactual exercises with the goal of evaluating the effects of changes in different frictions on the dynamics of MNE expansion. Among our exercises, we show how including the possibility of affiliate exports—captured by the compound option structure—interacts with changes in frictions to deliver different long-run responses of MNE activity compared to what a standard dynamic model of only horizontal FDI predicts. Since our sample includes affiliates located in the United Kingdom, Ireland, Germany, and France, we use the withdrawal of the United Kingdom from the European Union, Brexit, as our quantitative exercise, and increase different types of export costs between the UK and EU countries. Our model predicts that such an increase in export costs would have a static effect, a dynamic effect (coming from the inclusion of sunk costs), and an equilibrium price index effect. The strength of each effect on aggregate firm dynamics depends 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 per-period fixed costs would decrease both the number and sales of UK-based US affiliates to the EU. These different responses could not be captured with a static model of the MNE since one-time sunk and fixed per-period costs would be indistinguishable. Put together, our quantitative exercises point to the importance of considering both the time *and* space dimensions when evaluating the changes in the MNE expansion decisions after a shock.

Our paper is related to the existing literature in several ways. First, most contributions in the literature have analyzed MNE behavior in space, but not in time. Papers such as Ramondo and Rodríguez-Clare (2013), Tintelnot (2017), Fan (2020), Arkolakis et al. (2018), Alviarez (2019), and Head and Mayer (2019) have made substantial progress in building static general equilibrium

[4] It is worth noting that while the aggregate predictions of the proximity-concentration tradeoff are widely supported by the data, the evidence in support of the tradeoff at the firm level is mixed and not conclusive.

models with rich geography, in particular, by allowing MNEs to set up affiliates in countries that differ from the destinations of their sales.[5]

Making progress in dynamic setups, while keeping the spatial complexity of the static models, requires restricting the problem of the MNE to retain tractability. The sharp patterns that we document from observing MNE affiliates over time guide us on how to simplify this problem: thanks to our novel compound option structure, coupled with the assumption on the sequentiality of MNE decisions, and a standard Armington assumption, we are able to reduce the choice set of firms in a way that is consistent with the data, while keeping the model amenable to quantitative analysis. In this way, we are able to make substantial progress towards modeling the dynamics of MNE expansion, without sacrificing the spatial richness of static models.

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. Gumpert et al. (2020) focus on the life-cycle dynamics of exporters and MNEs as alternative ways of serving a foreign market. Given the nature of their question, the analysis does not consider export platforms, and focuses on life-cycle, rather than aggregate, firm dynamics.[6] Fillat and Garetto (2015) build a dynamic two-country model of exporters and MNEs, where they introduce the idea that MNE activities can be treated as a real option that gets exercised once an affiliate is opened abroad.[7] 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. Our model treats MNE activities as a compound, rather than a simple, option. In this way, we preserve the tractability of the multi-country dynamic problem, and expand on the spatial dimension by separating the locations of MNEs' production and sales.[8]

Third, our paper is naturally related to the large literature on export dynamics, which has been

[5]In static setups, the model of trade and FDI in Helpman et al. (2004) and its quantitative version in Irarrazabal et al. (2013) assume that the locations of production and sales of the MNE coincide (i.e. MNE activities are restricted to horizontal sales). While these two models feature the proximity-concentration trade-off at the plant-level, the models of trade and FDI in Ramondo and Rodríguez-Clare (2013) and Arkolakis et al. (2018) feature, as our model, that tradeoff at the aggregate, not plant, level.

[6]Our facts on US MNEs complement the facts in Gumpert et al. (2020), which compare the life-cycle dynamics of (non-MNE) exporters and MNEs, for France and Norway. While the BEA data are very detailed on the activity of affiliates abroad of US MNES, they do not provide any information about (non-MNE) US exporters. In contrast, the French and Norwegian data used in Gumpert et al. (2020) contain detailed information on all firms operating in those economies, while they lack detailed information on the activity of MNEs outside their home country.

[7]Impullitti et al. (2013) use a real option framework to study the entry and exit patterns of exporters.

[8]Other papers in the MNE literature limit both the spatial and dynamic dimension of the analysis by considering only horizontal FDI sales and only two periods (see, for instance, Ramondo et al., 2013; Egger et al., 2014; Conconi et al., 2016).

primarily concerned with quantifying the various costs of export activities and their welfare implications.[9] An important difference of our approach is that the nature of the MNE problem is more complex 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 allows us to solve the complex spatial problem of the MNE in a dynamic setup. In this way, we complement the literature on export dynamics by quantifying the frictions to MNE expansion and their implications in terms of aggregate firm dynamics and welfare.

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. These differences may be indicative of new US firms facing a very different set of frictions in the domestic and foreign markets.

2 Evidence on US MNE Expansion

We organize our empirical evidence into two sets of facts: facts about the expansion of affiliate sales over time; and facts about the location choices of the MNE over space.

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.[10] We also remove affiliates and parents with zero total sales. Finally, we consolidate affiliates belonging to the same parent

[9] 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) calculate the welfare gains from trade in a dynamic setting that matches well the life-cycle export facts. Arkolakis (2016) presents rich evidence on firm selection and export growth that supports dynamic theories of endogenous entry costs. Finally, Fitzgerald et al. (2019) show that the life-cycle growth patterns of export prices and quantities are quite different.

[10] 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.

Table 1: Summary Statistics.

	Horizontal sales	Export sales
No. of observations	132,493	132,493
with positive sales	127,220 (96%)	88,060 (67%)
of pure type	44,433 (34%)	5,273 (4.0%)
Sales accounted by pure type	16%	7.7%
Average share of total affiliate sales	72%	28%
Average affiliate sales over parent sales	5.6%	4.3%

Note: Observations are at the affiliate-year level, for new majority-owned affiliates 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.

and operating in the same country and 3-digit industry.[11]

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 affiliate exports to the United States and to third markets.[12] Every five years the BEA conducts a more detailed benchmark survey, which further distinguishes affiliate exports to Canada, the United Kingdom, and Japan.[13]

Table 1 shows the number of observations with positive horizontal and export sales in our sample. 96 percent of our affiliate-year observations have some horizontal sales, while about 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 four percent. Since affiliates that only export are very few and account for a small share of total affiliate sales, the model that we present in Section 3 does not include pure exporting affiliates.

Table 1 also reports that, on average, about 72 percent of affiliate sales are directed to the country where the affiliate is located, while the remaining 28 percent are exports. Furthermore, on average, both horizontal and export affiliate sales are small relative to parent sales (around six and four

[11]The reporting rules at the BEA 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. See Appendix A for details.

[12]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, however, the available break-down of affiliate exports by destination in our quantitative analysis.

[13]Affiliate exports to Canada are collected in 1989, 1994, 1999, 2004, and 2009. Affiliate exports to the UK and Japan are collected in 1989, 1994, and 1999.

percent). In the next section we show the evolution of these numbers over the life of the affiliate. Appendix A provides more details on the data coverage and sample construction.

2.2 Affiliate Sales Over Time

In this section we present evidence on the composition and size of affiliate sales over time. To establish the evidence from time series data, we need a set of affiliates that we can observe for a number of years from birth. Thus, in this section only, we focus on affiliates that open during our sample period and that survive for at least ten consecutive years. 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.[14]

I. MNE affiliates start with sales in their host market and expand into export markets.

We document that affiliates are mostly specialized in horizontal sales at birth, and that many of them incorporate export sales later in life.

Figure 1 shows the evolution of the intensive and extensive margins of horizontal and export sales of MNE affiliates. The intensive margin refers to the average share of horizontal (export) sales in total affiliate sales for firms with positive horizontal (export) sales. The extensive margin refers to the share of affiliates with positive horizontal (export) sales.[15] Figure 1a shows that, 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.[16] Figure 1b shows that while the share of affiliates with horizontal sales is stable at more than 95 percent, the share of exporting affiliates increases from 50 to 70 percent during the first ten years of life of the affiliate. Together, these findings suggest that, for horizontal activities, changes in sales shares are due to the intensive margin, while export shares increase only because of affiliates that start exporting. Over time, many affiliates incorporate export sales into their activities, but they never stop selling to their host market.

The patterns in Figure 1 are confirmed by OLS regressions that include a battery of fixed effects, as shown in Appendix Table B.1. Estimates that include affiliate fixed effects show that, on average,

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

[15]Since the BEA does not report affiliate exports by destination in their annual surveys, we define the extensive margin as whether or not an affiliate exports, instead of considering the extensive margins of each potential affiliate export destination.

[16]Note that these shares add up to more than 100 percent. This is because we exclude zeros separately for each type of sales, and the set of affiliates with positive horizontal sales is larger than the set of affiliates with positive export sales.

Figure 1: Affiliate sales and number of affiliates: horizontal vs export sales.

(a) Affiliate sales shares (b) Share of affiliates

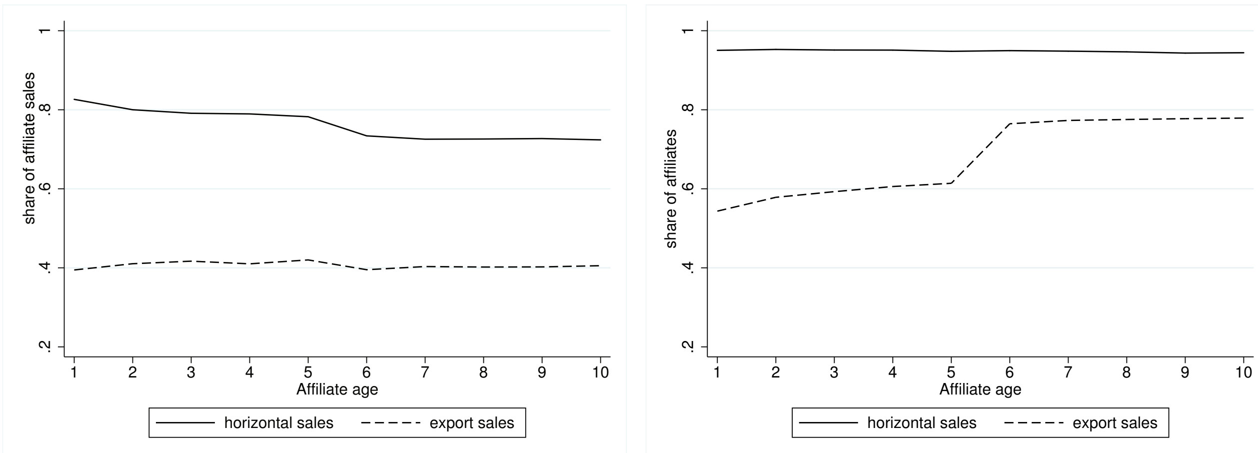

Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Horizontal and export sales refer, respectively, to sales to the market where the affiliate is located, and to sales to markets outside the local market. (1a): average sales, as a share of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

horizontal sales shares decrease, and export sales share increase, during the life of an affiliate, while the share of affiliates with exports is higher among older affiliates.

The facts documented in Figure 1 motivate an important assumption of our model: all affiliates start operations with horizontal sales; some affiliates also export from birth; and others expand into export markets later in life.

II. Affiliate sales as a share of parent sales grow at the time of export entry.

First, we show in Figure 2a that the size of affiliate horizontal sales, relative to parent sales, is close to the sample mean at the time of entry: affiliates do not start "small". Second, the growth of the affiliate in the host market, relative to the parent, is minimal. Third, Figure 2b shows that there is a jump in affiliate sales, relative to parent sales, in the year when the affiliate starts exporting.

While Figure 2 shows the raw data, these patterns are robust to including country-year and affiliate fixed effects in an OLS regression (see Appendix Table B.2).[17]

[17] Additionally, results are robust to using different subsamples of affiliates (see Appendix Table B.3): first, initial affiliates do not grow faster than subsequent affiliates of the same parent, suggesting that the age of the firm at the time an affiliate opens is not driving the results; second, affiliates with intra-firm exports do not grow differently than affiliates without intra-firm exports, suggesting that being part of a global value chain or not does not drive our baseline result; and finally, affiliates established through greenfield FDI versus mergers and acquisitions (M&A) both have flat sales ratios, suggesting that the lack of growth is not driven by pre-existing affiliates that became part of a new firm through M&A.

Figure 2: Affiliate sales relative to parent sales.

(a) Horizontal sales (b) Total sales, new exporters

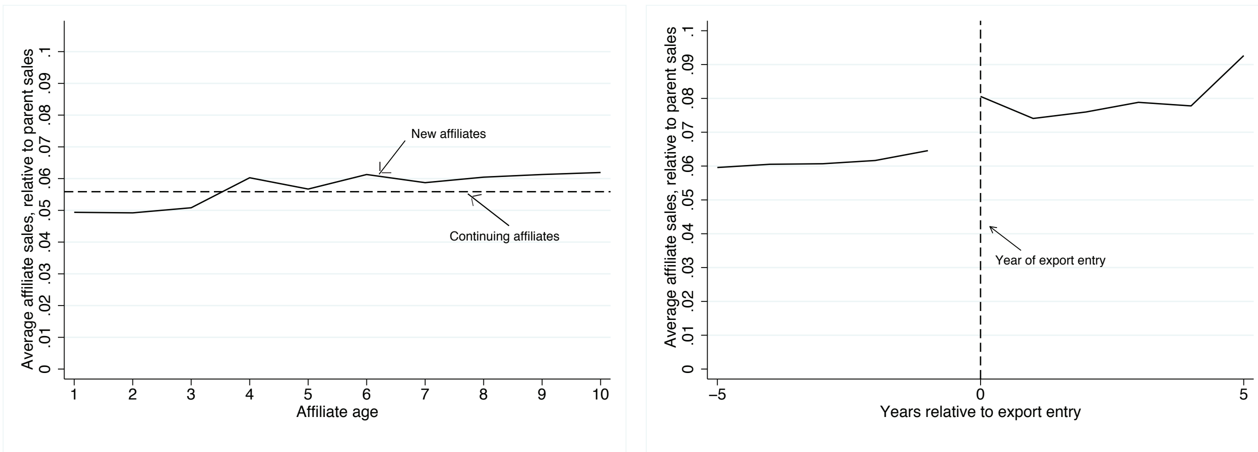

Notes: Figure 2a reports the average value of affiliate horizontal sales relative to the domestic sales of the US parent, for new affiliates in each of the first ten years of life, and compares it with the average value for all affiliates of all ages. Figure 2b reports the average value of affiliate total sales, relative to the domestic sales of the US parent, for a subsample of affiliates that are born with only horizontal sales and start exporting at a later age.

Taken together, the results in Figure 2 convey the message that affiliates sales as a share of parent sales grow significantly only when affiliates expand to other destination markets by exporting. Our model will capture this feature of the data, which will be important for the model tractability.

2.3 Geography, Affiliate Entry, and Affiliate Sales

In this section, we document how US MNEs expand across space. We use the *full* sample of affiliates of US MNEs. First, we show that the location of a new affiliate barely depends on the location of preexisting affiliates of the same MNE, both in terms of geographic proximity and of other measures of similarity between markets. Second, we document that the sales of affiliates (to unaffiliated parties) in a host country co-exist with export sales from affiliates of the same MNE located in other countries.

III. Affiliate entry follows a weak "extended gravity" pattern.

Figure 3 considers the sample of affiliates located in the ten most popular host countries for US MNEs and belonging to US parents with at least two foreign affiliates. For a given US parent, there is an extremely small difference between the unconditional probability of opening an affiliate in a country and the probability of opening an affiliate conditional on already having an affiliate in a country located in the same continent, or in a country with similar income per capita (Appendix

Table B.4 also includes as "similarity" measures sharing a border and sharing a language).[18] This is particularly true for rich host economies. Differences are slightly larger, but still small in magnitude and in many cases statistically insignificant, for less developed host economies, which are typically more engaged in global value chains (GVCs).[19]

Figure 3: Unconditional and conditional probability of affiliate entry.

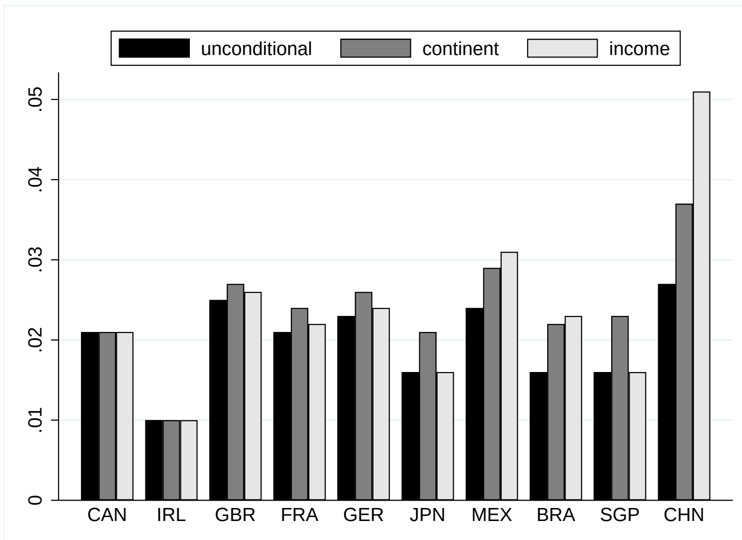

Note: Probabilities of affiliates' entry into the top-ten most popular destinations of US MNEs. Conditional probabilities refer to the probability of observing an MNE opening an affiliate in a country given that the parent already has an affiliate in another country in the same continent or in a country with similar income per capita. Similarity in terms of income per capita follows the group classification from the World Bank. The sample is restricted to parents with at least two affiliates worldwide.

This finding is in stark contrast with the analogous finding for exporter entry in Morales et al. (2019). For instance, they find that the unconditional probability of exporting to a given country is 0.7 percent and increases to 2.8 percent if the firm is already exporting to a country in the same continent. We find that the unconditional probability of opening an affiliate in the United Kingdom, for example, is 2.5 percent and increases to only 2.7 percent if the MNE already has an affiliate in the same continent. In general, while differences between conditional and unconditional probabilities for exporter entry range between 2 and 4 *times*, differences for MNE entry range between 2 and 20 *percent*.

Finally, it is worth noticing that our finding does not contradict the fact that affiliate entry exhibits

[18]Differences between unconditional and conditional probabilities become smaller when we restrict the sample to MNEs with more than five, and more than ten, affiliates (results available upon request).

[19]Appendix Table B.5 shows that the weak pattern of extended gravity is more pronounced among non-GVC affiliates (i.e., affiliates with zero intra-firm exports) than among GVC affiliates (i.e., affiliates with positive intra-firm imports). However, differences between conditional and unconditional entry probabilities are still small and often insignificant for both GVC and non-GVC affiliates.

gravity: MNEs open affiliates in closer and larger markets first (see Appendix Table B.6).

IV. Horizontal sales of existing affiliates do not change with sibling activities.

Our last fact presents evidence on the lack of changes in the horizontal sales of an affiliate when the activity of other affiliates of the same MNE changes.

First, an MNE may serve a destination from affiliates operating in the same 3-digit industry and located in different host countries (sibling affiliates). Since the BEA benchmark survey years record affiliate export data by destination for affiliates located in Canada, the United Kingdom, and Japan, for this exercise we restrict our sample to affiliates that are located in, or export to, those destinations.[20] We also restrict affiliate sales to *unaffiliated* parties. In this way, we make sure that our finding is not contaminated by GVCs within the MNE. Figure 4a reports the average horizontal sales (to unaffiliated parties) across affiliates located in country j, five years before and after an affiliate of the same parent and in the same industry, but located in a different country, starts exporting to unaffiliated parties in country j. As the figure shows, the horizontal sales of an affiliate do not significantly decrease when a sibling affiliate in the same industry starts exporting to that market from a different location. [21]

Second, an affiliate may shrink (or expand) in its country of operation when a new sibling affiliate in the same 3-digit industry opens in a different host country. Figure 4b reports the average horizontal sales across affiliates located in host market j, five years before and after a new affiliate of the same parent starts operating in the same industry but a different host country. Our estimates show that horizontal sales of existing affiliates do not change when their parent firm opens a sibling affiliate in the same industry but in a different country.[22]

The patterns of MNE entry and sales in a market documented in this section motivate the Armington assumption in the model, which implies that the conditional and unconditional probabilities of MNE entry into a market are similar, that horizontal sales and export sales from sibling affiliates may coexist in a destination market, and that entry of new sibling affiliates into a market does not affect the sales of existing affiliates operating in a different market.

[20] A large share of MNEs have affiliate exports to a country where they also operate affiliates: of the 20,359 affiliates that exported to Canada in 1999, 64 percent belong to a US parent that also has affiliates located in Canada; 70 percent of the 5,017 affiliates that exported to the United Kingdom have siblings located there; for Japan, this share was 47 percent, out of the 5,224 affiliates that exported to this market in 1999. Note that these numbers only include exports and horizontal sales to unrelated parties, to avoid concerns about within-firm GVCs.

[21] Exports to Canada are observed in 1989, 1994, 1999, 2004, and 2009. Exports to the UK and Japan are observed in 1989, 1994, and 1999.

[22] This result also holds when we consider affiliate total sales (available upon request).

Figure 4: Horizontal sales of existing affiliates and siblings activities. OLS.

(a) Years from export entry of sibling into same country (b) Years from opening of sibling in different country

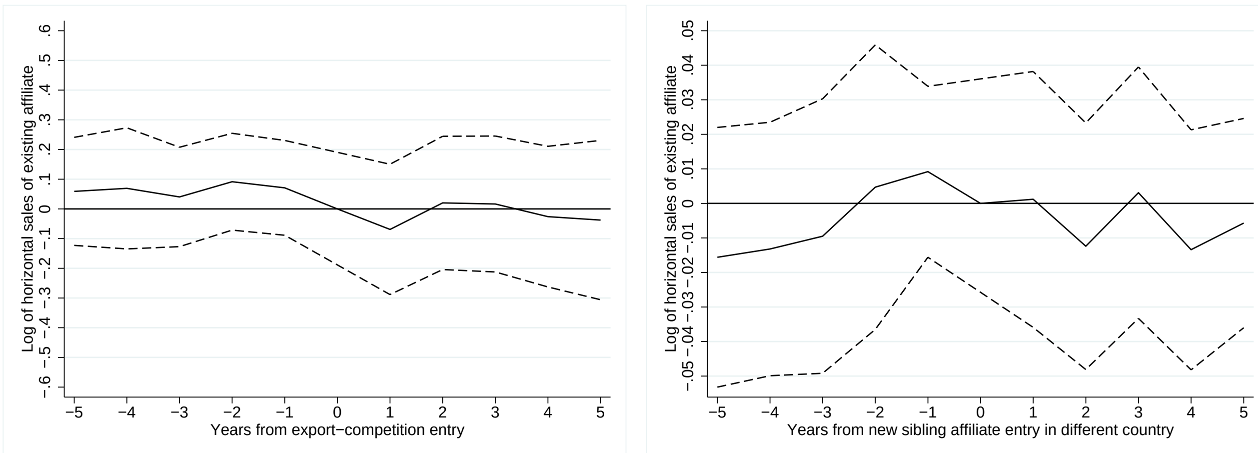

Notes: Figure 4a includes affiliates located in $j = Canada, UK, Japan$ (N = 5,092), while Figure 4b includes all affiliates of US MNEs in manufacturing (N = 132,493). OLS coefficients (solid line) from regressing the log of horizontal sales for affiliate a in country j and 3-digit industry h belonging to parent p at time $t \in \{-5, ...5\}$ on a set of dummies indicating: in Figure 4a, time from export entry (into j) of an affiliate located in country $k \neq j$ belonging to the same parent p and industry h, with $j = Canada, UK, Japan$; in Figure 4b, time from opening an affiliate in country $k \neq j$ belonging to the same parent p and industry h. We include the log of the MNE global sales, the log of the US parent sales, affiliate fixed effects and year fixed effects. Standard errors are clustered at the affiliate level. Dashed lines denote 95% confidence intervals.

3 A Dynamic Model of MNE Expansion

We build a dynamic model where MNEs open affiliates across countries over time. MNE 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.

The main innovation of the model is the introduction of a compound option formulation that allows us to characterize the richness of the decisions of MNEs 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 linear preferences over a composite good,

$$U_k = \int_0^{\infty} e^{-\rho t} Q_k(t) dt, \tag{1}$$

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) $\eta > 1$,

$$Q_k(t) = \left[\sum_i \sum_j \int_{\Omega_{ijk}(t)} \lambda_{ij}^{\frac{1}{\eta}} q_{ijk}(v,t)^{\frac{\eta-1}{\eta}} dv \right]^{\frac{\eta}{\eta-1}}. \tag{2}$$

The term $q_{ijk}(v,t)$ 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 variable λ_{ij} denotes a taste shifter.

Assumption 1 (**Armington**). *Varieties consumed and produced are firm-location specific.*

As in Armington (1969), Assumption 1 states that consumers perceive differently varieties produced in different locations by the same firm, a standard assumption in the literature. For example, consumers perceive Möet Chandon champagne produced in France as different from Chandon sparkling wine produced by the same firm in the United States. The implications of this assumption for affiliate entry are isomorphic to assuming that firm productivity is a random variable drawn from an independent multivariate Pareto distribution over varieties and production locations, as in Arkolakis et al. (2018).[23]Each country is populated by a continuum of firms. The US is the only source of MNEs: US firms decide whether to operate only in their home market or to also establish affiliates abroad. US firms can also export abroad from their US (parent) location. For this reason, we remove the index i that denotes a variety's origin country, and we use the subscript US to denote the parent's domestic operations.

Labor is the only factor of production. Each firm produces with a linear technology and operates under monopolistic competition. Each firm is characterized by a productivity parameter φ that determines the unit labor cost of the good produced. Each firm sets prices in each location to maximize profits from sales to each destination, $p_{jk}(\varphi) = \widetilde{\eta} c_{jk}(\varphi)$, with $\widetilde{\eta} \equiv \eta/(\eta - 1)$ and $c_{jk}(\varphi) \equiv w_j \tau_{jk}/\varphi$, for $j, k = US, 1, ...N$. 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 $\tau_{jj} = 1$, for $j = US, 1, ...N$.

[23]Firm productivity for good v across locations would be given by the random vector $\mathbf{z}(v) = [z_1(v), \ldots, z_N(v)]$, drawn from the multivariate Pareto distribution $1 - \sum_j T_{ij} z_j^{-\theta}$. The scale T_{ij} would take the place of the taste shifter λ_{ij} in (2). In the static model in Arkolakis et al. (2018), this case implies that the probability of opening an affiliate in market j is independent from the probability of opening an affiliate in market j'.

Variable profits from sales to k from location j are given by $\pi_{jk}(\varphi) = H(\tau_{jk}w_j/\varphi)^{1-\eta}\lambda_j P_k^{\eta} Q_k$, where $H \equiv \eta^{-\eta}(\eta-1)^{\eta-1}$ and P_k is the corresponding CES price index. For $j = k$, the variable π_{jk} denotes profits from either domestic sales ($j = US$) or horizontal affiliate sales ($j \neq US$). For $j \neq k$, the variable denotes profits from export sales, either from the parent ($j = US$) or from the foreign affiliates ($j \neq US$).

When a firm establishes an affiliate in a foreign country j, it has to pay a sunk entry cost $F_j^h > 0$ and a per-period fixed cost $f_j^h > 0$. We assume that there are no fixed or sunk costs associated with domestic sales, so that all firms produce and sell domestically. There are sunk costs $F_{jk}^e > 0$ to start exporting to country k, and per-period fixed export costs $f_{jk}^e > 0$, to be paid both from exporting affiliates located in country j and from exporting parents ($j = US$).[24]

Our setup has two important implications. First, Assumption 1, coupled with CES preferences and monopolistic competition, implies that there is no cannibalization of sales within the MNE, as goods produced in different locations are different goods. This is consistent with our fourth fact in Section 2, as well as with the evidence from the automobile industry in Head and Mayer (2019) in which each vehicle model (a "variety") is produced in—and sourced from—a single location. Second, consistent with our evidence on weak extended gravity in affiliate entry, Assumption 1 also implies that the MNE location decision is separable across locations.

However, as it will be evident below, the MNE entry decision into a location depends on its entry decision in other locations in a variety of ways. First, the formulation of the MNE problem as a compound option introduces dynamic interdependence in the MNE location choices: the choice of opening an affiliate in a country depends on the network of potential export destinations that the MNE can reach from that affiliate. Second, because of the (endogenous) responses of the price indexes to changes in the economy, the model displays substitution patterns in the location choices of MNEs that are consistent with the predictions of the proximity-concentration tradeoff in Brainard (1997): the higher trade costs and the lower plant-level scale economies, the larger the ratio of affiliate sales to export sales into a destination. Those substitution forces will be on display in our model calibration and at work in our counterfactual exercises.

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

[24] Appendix Table B.7 presents supporting evidence on the persistence in the affiliate operations and on the size of the US parent at affiliate entry and exit.

existing affiliates or from the US parent. Affiliate exports may include exporting back to the US 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, the value of exports 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 assumption we present next is guided by the empirical observations in Section 2.2.

Assumption 2 (Sequential MNE affiliate activities). *A new affiliate has to sell to its host market in order to eventually starting to export from there.*

Assuming sequential decisions 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.

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 firm-specific, drawn from a time-invariant distribution, $G(z)$, as in Melitz (2003). We assume that $Z = e^X$, where X is a Brownian motion with drift,

$$dX = \mu dt + \sigma dW, \tag{3}$$

for $\mu \in \Re$, $\sigma > 0$, and dW denoting a standard Wiener process. Additionally, we introduce host-country aggregate demand shocks by assuming that aggregate demand in destination country k evolves according to a geometric Brownian motion,

$$dQ_k = \mu_k Q_k dt + \sigma_k Q_k dW_k, \tag{4}$$

where $\mu_k \in \Re$, $\sigma_k > 0$, and dW_k denotes a standard Wiener process, possibly correlated with the Home aggregate productivity shock.[25] We assume that when a firm operates an affiliate in a foreign

[25] The shock process in (4) is analogous to assuming that foreign productivity X_k evolves according to a process

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).

Our shock structure is based on analytical and computational convenience, as well as on empirical observations. Analytically, the specification in (3)-(4) is equivalent to assuming that aggregate Home productivity and foreign demand growth behave like a random walk and that productivity and demand growth are independently and identically distributed. This functional form assumption guarantees the tractability of the model's solution. Concretely, 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

$$\tilde{\mu}_k = \mu_k + \mu(\eta-1) + \frac{\sigma^2}{2}(\eta-1)^2 + \gamma_k\sigma_k\sigma, \quad (5)$$

$$\tilde{\sigma}_k^2 = \sigma_k^2 + \sigma^2(\eta-1)^2 + 2(\eta-1)\gamma_k\sigma_k\sigma, \quad (6)$$

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. Computationally, relying only on aggregate shocks makes feasible the aggregation of individual firms' decisions and the computation of equilibrium price indexes for many countries. Additionally, we do not need to keep track of changes in the firms' productivity distribution over time, which significantly reduces the dimensionality of the state space. And finally, the introduction of country-specific demand shocks gives the model the flexibility to match the evolution of affiliate sales shares in different host countries. Empirically, this specification is consistent with the main sources of variation observed in the data for US MNE affiliates' sales; only a small share of that variation is driven by parent- and affiliate-level time-varying shocks (see Appendix Table B.8). Moreover, the persistence of the aggregate shock, together with aggregate productivity growing over time ($\mu \geq 0$), gives rise to the dynamic patterns documented in Figure 1: affiliates start serving their host market, and later on, they start expanding internationally.

Bellman equations. The state of the economy is described by the $(N+1)$-tuple $(X, \mathbf{Q})$, where $\mathbf{Q} = [Q_1, \ldots, Q_N]$. Let $\mathcal{V}(z, X, \mathbf{Q})$ denote the expected net present value of a US firm with productivity

that, in equilibrium, implies that foreign demand evolves according to (4).

z that follows an optimal policy when the state of the economy is $(X, \mathbf{Q})$,

$$\mathcal{V}(z, X, \mathbf{Q}) = V_{US}(z, X, \mathbf{Q}) + \sum_{j=1}^{N} \max\left\{V_j^o(z, X, \mathbf{Q}), V_j^a(z, X, \mathbf{Q})\right\}. \tag{7}$$

The function $V_{US}(z, X)$ denotes the value of US 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.

The value of an affiliate in country j is given by

$$V_j^a(z, X, \mathbf{Q}) = V_j^h(z, X, \mathbf{Q}) + \sum_{k \neq j} \max\left\{V_{jk}^o(z, X, \mathbf{Q}), V_{jk}^e(z, X, \mathbf{Q})\right\}. \tag{8}$$

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 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. Equations (7)-(8) reflect Assumption 2. 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.

The value of US operations, $V_{US}(z, X, Q)$, is given by

$$V_{US}(z, X, \mathbf{Q}) = V_{US,US}(z, X) + \sum_{k=1}^{N} \max\left\{V_{US,k}^o(z, X, \mathbf{Q}), V_{US,k}^e(z, X, \mathbf{Q})\right\}, \tag{9}$$

where $V_{US,US}(z, X)$ denotes the value of US domestic sales, $V_{US,k}^o(z, X, \mathbf{Q})$ denotes the option value of starting exporting to country k, and $V_{US,k}^e(z, X, \mathbf{Q})$ denotes the value of exports to country k.

Since all firms operate in the domestic market and sell to it, the value of domestic sales 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,

$$V_{US,US}(z, X) = \frac{1}{1 + \rho \Delta t}\left[\pi_{US,US}(z, X)\Delta t + E[V_{US,US}(z, X')|X]\right], \tag{10}$$

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

If a 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,

$$V_j^o(z, X, \mathbf{Q}) = \max\left\{\frac{1}{1+\rho\Delta t}E[V_j^o(z, X', \mathbf{Q}')|(X, \mathbf{Q})]; V_j^a(z, X, \mathbf{Q}) - F_j^h\right\}, \tag{11}$$

where $\mathbf{Q}'$ denotes the vector of realizations of demand shocks next period. 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, $V_j^a(z, X, \mathbf{Q})$. Because goods are firm- and location-specific, each firm evaluates entering each location separately.

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

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

An 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 by (8), 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

$$V_{jk}^o(z, X, \mathbf{Q}) = \max\left\{\frac{1}{1+\rho\Delta t}E[V_{jk}^o(z, X', \mathbf{Q}')|(X, \mathbf{Q})]; V_{jk}^e(z, X, \mathbf{Q}) - F_{jk}^e\right\}. \tag{13}$$

An 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

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

An 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, $V_{jk}^o(z, X, \mathbf{Q})$. The expressions for the option value of US exports, $V_{US,k}^o(z, X, \mathbf{Q})$, and the value of US exports, $V_{US,k}^e(z, X, \mathbf{Q})$, are analogous to the expressions in (13) and (14), for $j = US$.

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. Given the affiliate's location, 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 $Y_k \equiv e^{(\eta-1)X} Q_k$.

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

$$V^o_{jk}(z, Y_k) = B^o_{jk}(z) Y_k^{\beta_k}, \tag{15}$$

$$V^e_{jk}(z, Y_k) = \frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f^e_{jk}}{\rho} + A^e_{jk}(z) Y_k^{\alpha_k}. \tag{16}$$

The variables $B^o_{jk}(z) > 0$ and $A^e_{jk}(z) > 0$ are firm-specific, while $\alpha_k < 0$ and $\beta_k > 1$ are the roots of $\tilde{\sigma}_k^2 \xi^2/2 + \left(\tilde{\mu}_k - \tilde{\sigma}_k^2/2\right)\xi - \rho = 0$. The term $B^o_{jk}(z) Y_k^{\beta_k}$ in (15) represents the option value of exporting to country k and is increasing in the realization of the composite shock. Similarly, $A^e_{jk}(z) Y_k^{\alpha_k}$ in (16) 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 $B^o_{jk}(z) > 0$, $A^e_{jk}(z) > 0$, 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. The expressions in (15) and (16) are also valid for $j = US$ (US exports).

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,

$$V^h_j(z, Y_j) = \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f^h_j}{\rho} + A^h_j(z) Y_j^{\alpha_j}, \tag{17}$$

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

$$\begin{aligned} V^a_j(z, \mathbf{Y}) &= A^h_j(z) Y_j^{\alpha_j} + \frac{\pi_{jj}(z, Y_j)}{\rho - \tilde{\mu}_j} - \frac{f^h_j}{\rho} + \ldots \\ & \sum_{k \in \mathcal{A}_j(z)} \left[\frac{\pi_{jk}(z, Y_k)}{\rho - \tilde{\mu}_k} - \frac{f^e_{jk}}{\rho} + A^e_{jk}(z) Y_k^{\alpha_k}\right] + \sum_{k \notin \mathcal{A}_j(z)} \left[B^o_{jk}(z) Y_k^{\beta_k}\right], \end{aligned} \tag{18}$$

where $\mathcal{A}_j(z)$ is the subset of countries where an affiliate of firm z located in j exports to, and $\mathbf{Y} = [Y_1, \ldots, Y_N]$. Inspecting (18) makes clear that the compound option structure introduces interdependence in the MNE location choices: the value of an affiliate depends 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

$$V_j^o(z, Y_j) = B_j^o(z) Y_j^{\beta_j}. \tag{19}$$

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 for 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,

$$V_{US,US}(z, X) = \frac{\pi_{US,US}(z, X)}{\rho - \hat{\mu}}. \tag{20}$$

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

3.3 Price indexes

Thanks to the tractability of our multi-country model, we are able to solve for the dynamics of 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 reports the expressions for the price indexes for each country j, the law of motion of the mass of MNEs in each country j, and the law of motion of the mass of affiliates in j that export to a destination k.

The ability to solve for equilibrium price indexes derives from the choices we made about 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. (2019), 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 since we do not need to keep track of the evolution of the firm's productivity distribution. Finally, we assume that the wage in each destination is exogenous.

3.4 Model predictions

In this section, we derive analytical properties of the model linking firm-level productivity, host market characteristics, and affiliate entry and export 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. Under this assumption, the option-value terms $A^e_h(z)$ and $A^e_{jk}(z)$, in (16), (17), and (18), are zero. Hence, we can obtain closed-form solutions for the affiliate entry and export entry thresholds,

$$Y^{OH}_j(z) = \left(\frac{\beta_j}{\beta_j - 1}\right) \cdot \left(\frac{f^h_j + \rho F^h_j}{\rho} - \mathbf{V}^E_j(z, \mathbf{Y}_{-j})\right) \cdot \left(\frac{\rho - \tilde{\mu}_j}{\kappa_{jj}(z)}\right), \quad (21)$$

$$Y^{OE}_{jk}(z) = \left(\frac{\beta_k}{\beta_k - 1}\right) \cdot \left(\frac{f^e_{jk} + \rho F^e_{jk}}{\rho}\right) \cdot \left(\frac{\rho - \tilde{\mu}_k}{\kappa_{jk}(z)}\right). \quad (22)$$

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 $\mathbf{V}^E_j(z, \mathbf{Y}_{-j})$ denotes the total value of exporting from an affiliate in j for a firm with productivity z.[26] Details on the derivation of (21) and (22) are in Appendix C.

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

Proof. See Appendix C.

Under the assumption that, $\forall k \neq j$, $Y^{OH}_j(z) < Y^{OE}_{jk}(z)$ (the threshold for the realization of the composite shock 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 5 illustrate these predictions. The red and blue lines denote, respectively, the threshold for opening an affiliate in j, $Y^{OH}_j(z)$, and the threshold for starting exports from j to k, $Y^{OE}_{jk}(z)$. They are decreasing functions of the firm's productivity z, and hence, they are invertible functions. In Figure 5a, we assume that the realization of the aggregate shock is Y' and that we observe two firms with affiliates in the same host country j. Firm 1 with productivity z_1

[26] Notice that if $(f^h_j + \rho F^h_j)/\rho - \mathbf{V}^E_j(z, \mathbf{Y}_{-j}) < 0$, then $Y^{OH}_j(z) < 0$. 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.

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 $z_2 \geq z^e_j(Y') \geq z_1$, where $z^e_j(Y')$ denotes the productivity necessary to export when the realization of the composite shock is Y' (*i.e.*, the inverse of $Y^{OE}_{jk}(z)$). 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 5b, $z_1 \geq z^e_j(Y'')$ 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.

Figure 5: Model predictions.

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

(a) Exporters vs non-exporters

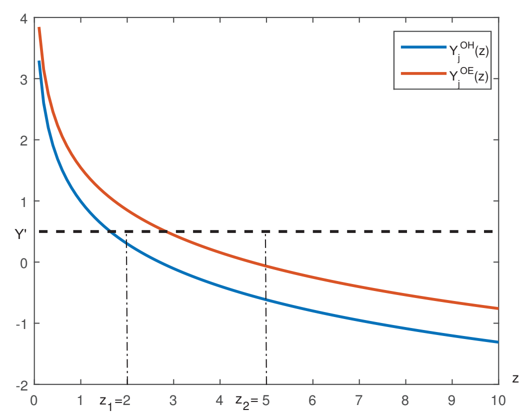

(b) Early vs late exporters

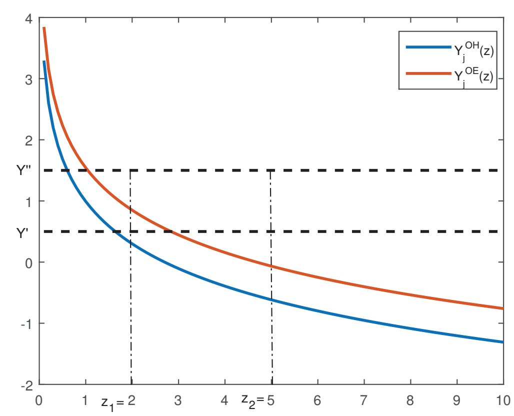

Host market characteristics and the timing of entry.

(c) Affiliate size

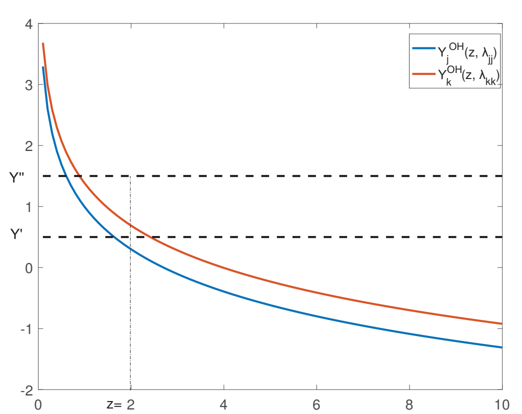

(d) Entry costs

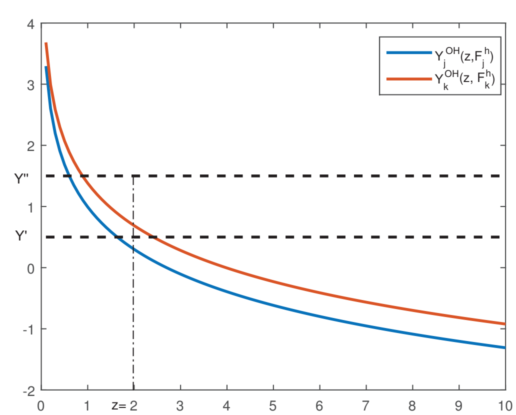

Proposition 1 also implies that the model exhibits sorting of MNEs across different host markets. First, MNEs with larger parent sales enter more foreign markets by opening foreign affiliates. Second, the mass of firms with affiliates in n host markets is decreasing in n, so that there is a negative relationship between the number of firms with affiliates in n markets and their parent sales (see Corollaries 1 and 2 in Appendix C).

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

Proof. See Appendix C.

Proposition 2 relates to the expansion strategies of an 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 or demand increasing over time ($\tilde{\mu} \geq 0$)—an 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 an MNE first opens affiliates in markets that are less costly to enter.[27]

The lower panels of Figure 5 illustrate the predictions of Proposition 2. Figure 5c 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 also opens an affiliate in country k, illustrating that, controlling for factor costs and host country size, an MNE opens its largest affiliates first. Figure 5d 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 also opens an affiliate in country k.

Even though Propositions 1 and 2 cannot be proven in the general case with arbitrary fixed costs, our quantitative analysis reveals that their implications hold in the general case where exit thresholds are active. Furthermore, in the next section, we provide empirical support for the implications of both propositions.

[27]This prediction is consistent with the findings in Egger et al. (2014) for German MNEs, and with the evidence in Appendix Table B.6.

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 then proceed to show the fit of the model to the data.

4.1 Procedure

We start by setting the parameters capturing preferences and technology 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. 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.

We set the elasticity of substitution to $\eta = 5$, 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 does not diverge ($\rho > \tilde{\mu}_j,\ \forall j$).[28] 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 Simonovska and Waugh (2014).

We use data on expenditure-based real GDP growth across countries from the Penn World Tables (9.1) to calibrate the composite shock process for each country in our sample. The composite shock 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. For $\tilde{\sigma}_j$, we first set the standard deviation of US aggregate productivity, σ, to match the standard deviation of labor productivity among US firms, and next, 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 (6) 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, $\forall j$. We take wages as exogenous

[28]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 chose not to add unnecessary parameters and rather to assume a high value for the time preference rate.

and we set them to match real GDP per unit of equipped labor, from Klenow and Rodriguez-Clare (2005), an average over the period 1995-2000. Appendix Table D.1 shows the parameters that are calibrated directly from the data, for each of the top ten host countries for US FDI.

It remains to calibrate the preference shifters, λ_{ij}, and the parameters related to the costs of MNE expansion: the fixed and sunk costs of affiliate opening, f_j^h and F_j^h, for $j = 1, ...10$; and the variable, fixed and sunk costs of exports, both from the US and from foreign affiliates, τ_{jk}, f_{jk}^e and F_{jk}^e, for $j = US, 1, ...10$, $k \neq j$.

Due to data limitations, we make some symmetry assumptions. First, we assume that $\lambda_{ii} = 1$ and $\lambda_{ij} = \lambda_j \neq 1$, for $i = US$ and for all $j \neq i$. 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 $j, k = 1, ...10$, $k \neq j$, and $k \neq US$. Third, we assume that iceberg trade costs for destinations where 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.[29] Fourth, we assume that the fixed and sunk costs of US exports are symmetric across destination countries: $f_{US,k}^e = f_{US}^e$ and $F_{US,k}^e = F_{US}^e$, for $k = 1, ...10$. Notice that if fixed and sunk costs were constant, together with constant wages and growing demand, they would imply that frictions to MNE activities would be irrelevant in the long run and all firms would become MNEs with affiliates in every country. To avoid this counterfactual situation, we assume that the fixed and sunk costs of MNE activities in each host country j also grow at the same rate as aggregate demand, μ_j. Hence, we need to calibrate the initial values of the fixed and sunk costs for each host country.

We are left with 122 parameters to calibrate, for which we target 122 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. Specifically, the intensive margin of exports, given by the 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 costs, F_j^h and f_j^h, respectively. 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, affiliate horizontal sales in country j, relative to US parent sales, help identify the taste shifter λ_j.

We choose the values of the parameters to best fit the data moments, for each country. To

[29] 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, consistent with the calibrated value of η.

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 fit

Table 2 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 for each host country, including direct exports from the US parents. Appendix Tables D.2 and D.3 show the calibrated parameters by host country.

Table 2 shows that the model matches quite well both the static and dynamic targeted moments. We also include in the table three sets of non-targeted moments: moments related to affiliate size advantage, MNE sorting patterns, and exit moments.

The moments capturing the affiliate size advantage are related to the analytical predictions of the model described in Section 3.4. First, 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 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. To provide more detail of this prediction in the data, Figure 6a plots the distribution of log horizontal sales for two subsets of affiliates in our sample: affiliates that are born with only horizontal sales, and affiliates that are born also with exports. The figure clearly shows that affiliates that export are on average larger than affiliates with only horizontal sales at birth. 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 exports in its first year of life are 3.7 times larger than the average horizontal sales of an affiliate that starts exporting after its first year of life, averaging across affiliates' host markets. Our calibrated model generates an early-exporter premium of 5.5. Figure 6b further shows that affiliates that start exporting in the first year of life are larger in their host country compared to affiliates that start exporting later.[30]

Additionally, the model has predictions about the expansion patterns of an MNE. Proposition 2

[30]The relation between size in the host market, export status, and age at first export survives the inclusion of affiliate age, country-year and industry fixed effects (results available upon request).

Table 2: 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	0.146
1.3 Affiliate sales share to third countries	0.276	0.262
1.4 Affiliate sales share to Canada	0.015	0.013
1.5 Affiliate sales share to the U.K.	0.069	0.087
1.6 Affiliate sales share to Japan	0.033	0.023
2. Static moments: extensive margin		
2.1 Share of MNEs with affiliates in j	0.287	0.285
2.2 Share of affiliates in j exporting to US	0.566	0.567
2.3 Share of affiliates in j exporting to third countries	0.650	0.649
3. Dynamic moments: entry		
3.1 Share of MNEs opening affiliates in j	0.035	0.018
3.2 Share of affiliates in j that start exporting to the US	0.030	0.023
3.3 Share of affiliates in j that start exporting to third countries	0.031	0.025
Non-Targeted Moments		
4. Static moments: affiliate size advantage		
4.1 Exporter size advantage	6.27	7.04
4.2 Early exporter size advantage	3.68	5.58
4.3 First affiliate size advantage	2.57	2.13
5. Static moments: MNE sorting		
5.1 Elasticity of average sales in US w.r.t. # of markets entered	0.736	0.816
5.2 Elasticity of average sales in US w.r.t # of firms entering multiple markets	-0.424	-0.722
6. Dynamic moments: exit		
6.1 Share of MNEs shutting down affiliates in j	0.113	0.066
6.2 Share of affiliates in j that stop exporting to the US	0.025	0.040
6.3 Share of affiliates in j that stop exporting to third countries	0.026	0.019

Note: Averages across host countries and years. Data moments for Japan, Canada, and the United Kingdom are averages over benchmark year surveys only. Denominators are: in 1.1, US parent's sales; 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; in 3.2 (3.3), total number of affiliates in j in period before export entry into US (third countries); in 6.1, total number of affiliates in j in period before exit; and in 6.2 (6.3), the total number of affiliates in j that export to the US (third countries) in the period before export exit. 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, early exporter affiliate size advantage refers to the average size of exporting MNE affiliates that start exports in their first year of life relative to the average size of exporting MNE affiliates that start exports after their first year. In 4.3, first affiliate size advantage refers to the ratio of the size of the first foreign affiliate of an 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. The elasticities in 5. are computed by OLS, aggregating the firm-level observations of MNEs that enter the same number of countries.

Figure 6: Affiliate size, export status, and the timing of export entry.

(a) Non-exporters vs exporters at birth (b) Early vs late exporters

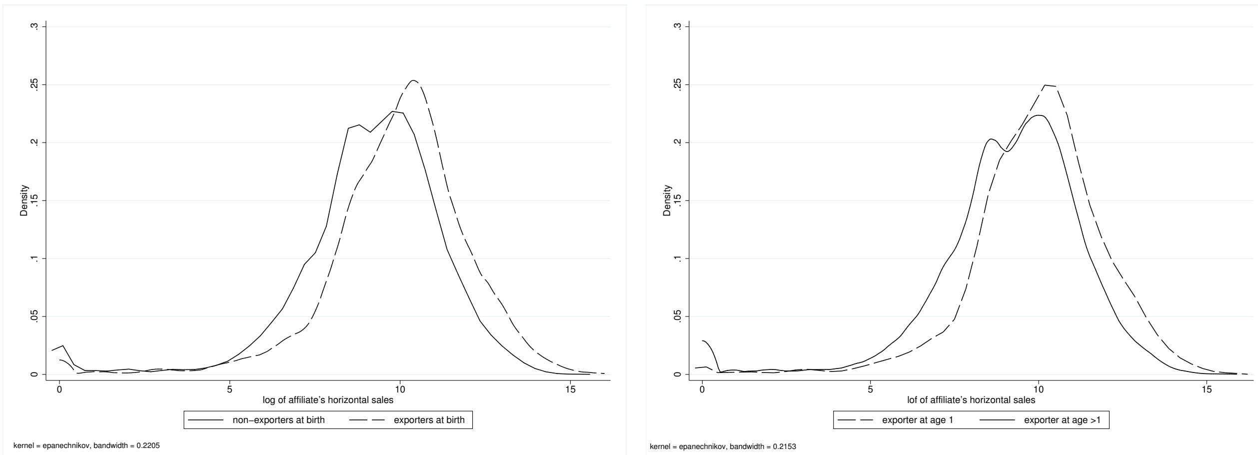

Notes: Sample of new majority-owned affiliates that survive for at least ten consecutive years, in manufacturing. Kernel density of log horizontal sales for affiliates that: are born with exclusively horizontal sales (non-exporters) and those with exports (exporters), in (6a); start exporting in their first year of life and those that start older, in (6b).

implies that MNEs open their largest affiliates first. Table 2 shows that, in the data, the horizontal sales of a first affiliate of an MNE are, on average, 2.6 times larger than the horizontal sales of the MNE's subsequent affiliates. The model generates a first-affiliate size premium of 2.1.

Similarly to Eaton et al. (2011) for exporters, the model predicts that the MNEs with the largest US parent sales should enter more markets, including less popular ones. To support this implication, in Figure 7a, we plot the average domestic sales of the US parents against the minimum number of markets where those MNEs have affiliates. The positive correlation between these two variables indicates that –consistent with Proposition 1– MNEs with larger US sales have affiliates in more foreign markets. In turn, Figure 7b confirms that not only are the MNEs with the largest parent sales the ones that enter more markets, but also that they are the MNEs entering "less popular" markets (see Corollary 2 of Proposition 1 in Appendix C.3). Table 2 shows that the calibrated model captures the magnitude of the elasticity of US parent sales with respect to the number of countries where MNE affiliates operate very accurately. In contrast, the calibrated model overestimates the elasticity of parent sales with respect to the number of firms entering multiple markets.

Finally, as shown in the last panel of Table 2, the exit rates in the calibrated model are close to the ones we observe in the data. The model slightly under-predicts affiliate exit, and over-predicts exit from export markets.

Figure 7: Parent size in the United States and MNE entry.

(a) Number of markets penetrated

(b) Number of firms penetrating multiple markets

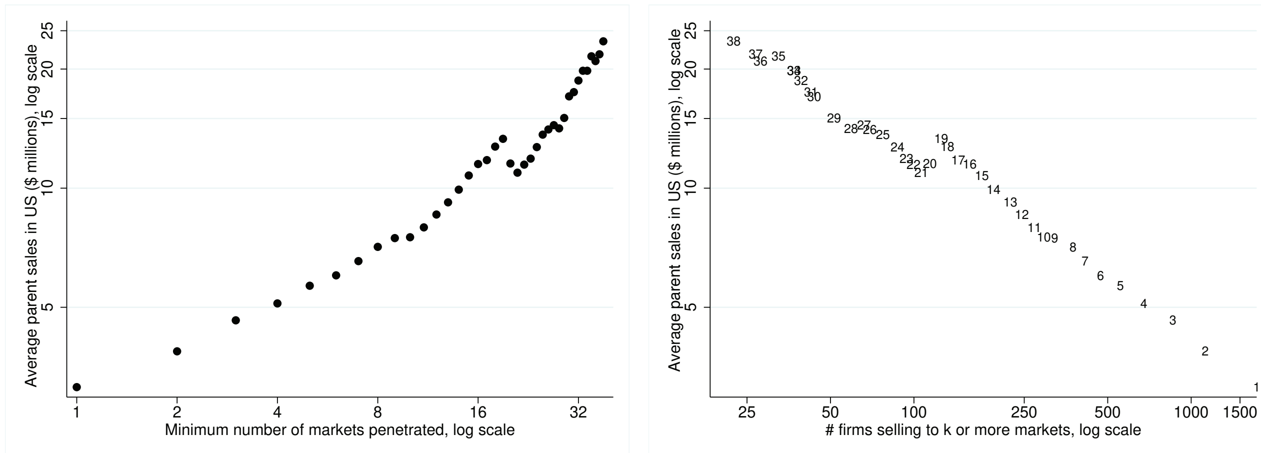

Notes: Observations at the country-year level. Slopes are: (7a) 0.566 (s.e. 0.034); (7b) -0.442 (s.e. 0.008).

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 3 we express the calibrated MNE costs as shares of firm revenues and in monetary values. Appendix Tables D.10 and D.11 report these results by host country.

On average, opening an affiliate involves spending a very low share of the 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 are less costly than exports to other destination markets, both in terms of fixed and sunk costs. This result is intuitive as the United States is the origin country of the affiliates.[31]

Table 4 highlights interesting variation in the calibrated frictions across selected host countries. Favorable "tax-haven"-like policies that attract FDI make opening an affiliate in Ireland much less costly compared to other countries like Brazil, France, and Japan, as reflected in the calibrated values of sunk MNE costs, F_j^h, expressed as a share of US parent revenues. Conversely, the operating

[31] To put our numbers in perspective, Das et al. (2007) estimate the export entry costs of Colombian exporters at around 400 thousand 1986 US dollars. Despite being set to match quantitatively similar entry rates, our numbers for affiliate exports are much smaller than theirs, consistent with the observation in Das et al. (2007) that larger firms tend to have lower export entry costs—and foreign affiliates of US MNEs are much larger than Colombian exporters.

Table 3: 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	0.003	168.6
Fixed affiliate entry cost f_j^h, as % of horizontal sales	18.7	12.2	2.35	1,642
Sunk export cost F_{jk}^e, as % of horizontal sales				
To United States	0.33	0.16	0.03	16.3
To other destinations	2.84	0.96	0.18	82.0
Fixed export cost f_{jk}^e, as % of exports sales				
To United States	12.7	10.1	1.67	98.5
To other destinations	13.2	11.3	3.22	783.8

Note: US parent sales in the year of affiliate entry. Horizontal sales evaluated in the year the affiliate first exports to the destination, for F_{jk}^e, and averaged across years, for f_j^h. Export sales to a destination are averaged across years. 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, averaged across host countries. For confidentiality purposes, median sales are an average of the 9 observations around the median.

Table 4: Calibrated MNE costs as shares of sales: by type, destination, and selected host country.

as % of sales	Sunk MNE cost	Fixed MNE cost	Sunk export cost		Fixed export costs	
			to US	to other	to US	to other
	F_j^h	F_j^h	$F_{j,US}^e$	F_{jk}^e	$f_{j,US}^e$	f_{jk}^e
Brazil	0.068	9.96	0.003	0.599	9.57	11.76
France	0.035	15.81	0.003	1.667	9.90	11.69
Ireland	0.006	10.85	0.000	3.706	10.88	10.01
Japan	0.066	16.30	0.000	0.020	11.10	13.14

Note: Costs are reported as percentage of sales, using sales values for the median affiliate in each of the top-ten host countries included in the calibration. For confidentiality purposes, median sales are an average of the 9 observations around the median. Sunk MNE cost, F_j^h, reported as a share of US parent sales in the year of affiliate entry. Fixed MNE cost, f_j^h, reported as a share of average horizontal sales across years. Sunk export cost, F_{jk}^e reported as a share of horizontal affiliate sales in the year the affiliate first exports to the destination. Fixed export cost, f_{jk}^e, reported as a share of average affiliate export sales to the destination across years.

costs of maintaining affiliates, f_j^h, as a share of affiliate revenues, reflect variation in factor costs: higher for France and Japan, and lower for Ireland. The sunk export costs to the US, $F_{j,US}^e$, as a share of the affiliate horizontal sales in the year of export entry, are negligible across host countries, highlighting the fact that it is almost costless to start an export channel to the Home country of the MNE. Meanwhile, the sunk export costs to other countries, F_{jk}^e, are sizable. The fixed export costs, both to the US ($f_{j,US}^e$) and to third countries (f_{jk}^e), as a share of the affiliate exports to those destinations, display less variation across host countries.

4.4 The proximity-concentration tradeoff

An important feature of the data that our quantitative model reproduces is the proximity-concentration tradeoff, or the observation, first documented by Brainard (1997), that the ratio of a country's exports to horizontal affiliate sales to a foreign country is decreasing in transportation costs and increasing in plant-level scale economies. To show that our model generates a proximity-concentration tradeoff, we focus on the United States and regress the ratio of US exports to country j to the horizontal sales in country j of affiliates of US MNEs, as implied by the calibrated model, on our calibrated variable export costs from the United States to country j, the calibrated ratio of fixed MNE costs incurred in country j to fixed export costs from the United States to country j at time t, and the calibrated size of aggregate demand in country j at time t:

$$\log\left(\frac{Exports_{us,j}(t)}{Horizontal\ sales_j(t)}\right) = \beta_\tau \log(\tau_{us,j}) + \beta_f \log\left(\frac{f_j^h(t)}{f_{us,j}^e(t)}\right) + \beta_Q \log(Q_j(t)). \tag{23}$$

The first column of Table 5 shows the results. Our calibrated model delivers the proximity-concentration tradeoff at the aggregate level: US exports relative to horizontal FDI to a country are decreasing in $\tau_{us,j}$ and increasing in $f_j^h(t)/f_{us,j}^e(t)$. The second column of Table 5 shows that this result is robust to the inclusion of annualized sunk costs.

5 Quantitative Analysis: Brexit

Armed with the calibrated model, we explore the implications of different 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 quantitative exercise the United Kingdom abandoning the European Union (EU), "Brexit." Different implementations of Brexit have as common element the increase in export costs between the United Kingdom and countries in the

Table 5: The proximity-concentration trade off, United States.

Dep var	$log \frac{Exports_{us,j}(t)}{Horizontal\ sales_j(t)}$	
	(1)	(2)
$log\ \tau_{us,j}$	-4.342***	-4.361***
	(0.189)	(0.187)
$log\ \frac{f_j^h(t)}{f_{us,j}^e(t)}$	0.695***	
	(0.021)	
$log\ \frac{f_j^h(t)+\rho F_j^h(t)}{f_{us,j}^e(t)+\rho F_{us,j}^e(t)}$		0.692***
		(0.027)
$log\ Q_j(t)$	-0.674***	-0.672***
	(0.117)	(0.012)
R-squared	0.983	0.983
Observations	300	300

Note: OLS estimates of (23). Model-simulated data and calibrated parameters. All specifications include time fixed effects. Robust standard errors in parenthesis. Levels of significance are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

EU. In the model, this increase can be captured by increases in the iceberg trade cost, the fixed export cost, or the sunk export cost. We analyze scenarios where we increase trade frictions one type at the time, and a scenario where we increase all trade frictions simultaneously.

In a different exercise, we evaluate quantitatively the role played by the compound option structure of our model. To such end, we evaluate the response of an increase in the cost of US MNE activities in the United Kingdom, captured in our model by the per-period cost of MNE operations, f_{UK}^h, and by the sunk entry cost, F_{UK}^h, in the calibrated model with and without export platforms.

Overall, our quantitative exercises point to the importance of including dynamics, together with a rich spatial structure, when evaluating the response of MNE expansion to globalization shocks.

5.1 Increases in trade costs

We simulate the model for 30 periods and impose a permanent change in one or more of the cost parameters at $t = 15$. First we increase, one at a time, the barriers to export between the United Kingdom and the EU countries in our sample: $\tau_{UK,j}$, $f_{UK,j}^e$, and $F_{UK,j}^e$ (j = Ireland, Germany, and France). For comparability, we increase each friction by an amount that results in a 20 percent increase in the total per-period cost of FDI, $\left(f_{UK,j}^e + \rho F_{UK,j}^e\right) \tau_{UK,j}^{\eta-1}$. When we increase all export barriers at once, we increase each barrier by the amount calculated for the first exercise.

Increasing trade barriers between the United Kingdom and EU countries has three main effects. First, when exporting from the United Kingdom to the EU becomes more costly, 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 export destinations. Analogously, exporting from the EU to the United Kingdom also becomes more costly, and export sales from EU-n to UK-based affiliates also decline, decreasing the incentive to open affiliates in those countries. These are static, partial equilibrium effects. In addition, increases in trade costs affects the affiliate export band of inaction, which affects affiliate export entry and exit —an effect only present in dynamic models. Finally, increases in trade frictions raise the price index in the destination countries, encouraging more export entry.

The quantitative results in Figure 8 combine the effects of these three forces. Results are presented as deviations from the baseline scenario. Before the shock, about half of US MNEs had affiliates located in the United Kingdom (see Appendix Table D.6). Increasing any of the trade costs causes a permanent decrease in the share of UK-based affiliates of US MNEs of around three percent, while the increase of all trade frictions at once drives a decrease in this share of five percent (Figure 8b). Naturally, the affiliates that exit are the smallest and account for only a few percentages of affiliate sales. This observation, together with the increase in the UK price index (see Appendix Figure E.1a), explains the somewhat counterintuitive result of Figure 8a, where affiliate sales to the United Kingdom increase when variable trade costs increase. The effects on horizontal activities are small, since trade frictions affect those activities only indirectly through the compound option.

Figures 9c and 9d show the effects of increasing trade costs on export activities of UK-based affiliates of US MNEs directed to EU countries. 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. Specifically, the decline in export sales when either $f^e_{UK,j}$ or $\tau^e_{UK,j}$ increase comes from affiliates that stop exporting. For the case of a shock to $\tau_{UK,j}$, export sales also decline due a decline in the intensive margin, so that changes in this friction have the highest impact on affiliate export sales. 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, much larger than the one produced by increases in fixed export costs. Conversely, 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 small increase in affiliate export sales in Figure 8c and in the share of exporting affiliates in Figure 8d. Except for the case of an increase in the sunk export cost, the Brexit shock reduces the share of affiliates that export. The increase in the fixed export cost

Figure 8: Brexit: US MNE affiliates in the United Kingdom.

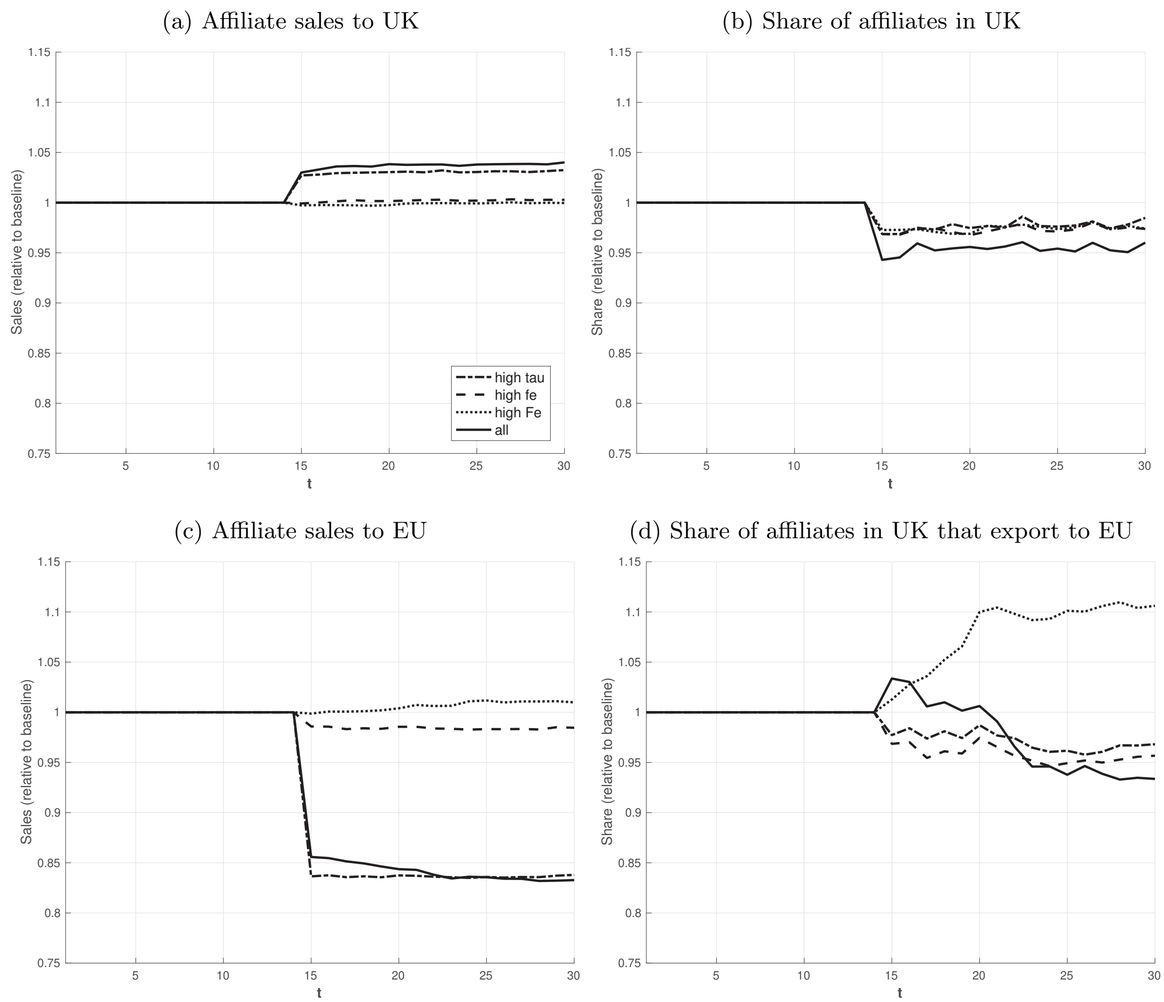

Note: "high X" refers to an increase in the barrier X from/to the United Kingdom to/from country k. "All" refers to increasing all three export barriers from/to the United Kingdom to/from country k at once. Country k refers to Ireland, Germany, and France.

produces the largest decline in export participation because this cost is intimately related to the affiliates' decision to exit a market.

The difference in the magnitude and direction of the effects when either sunk, fixed, or variable costs change highlights the importance of including dynamics in models of MNE behavior. Static models cannot distinguish between sunk and fixed costs; yet, MNEs respond very differently to changes in these types of costs.

The scenario where we increase all the trade frictions simultaneously highlights the differences between short and long run dynamics after a shock. In particular, in the short run, the effect of the increase in sunk export costs dominates and the share of UK-based affiliates 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 US affiliates located in France, Ireland, and Germany present some important quantitative differences from the results for the United Kingdom. Figure 9 shows that 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.[32] As for the United Kingdom, when $\tau_{UK,FR}$ increases, this decline in the extensive margin is accompanied by an increase in the intensive margin of horizontal sales in France. This increase is driven by the increase in the French price index. Naturally, the share of US affiliates in France that export to the United Kingdom, as well as their export sales, drop after the Brexit shock in almost all specifications, with an increase in the sunk export cost generating a non-monotonic response (the share of exporting affiliates first increases due to a decline in exit rates and then decreases).

Increasing variable, fixed, or sunk cost of exporting has different *quantitative* effects on aggregate firm dynamics. Even though the increase in the per-period cost of FDI is the same in all cases, the type of trade barrier that changes matters for both the extensive and intensive margins of MNEs' decisions, and in turn, for aggregate dynamic responses. Put differently, considering the global structure of the MNE in time *and* space is important for accurately assessing the consequences of shocks to the costs of MNE expansion.

5.2 Reallocation towards US parent exports

Our analysis of Brexit so far has focused on the foreign affiliates of US MNEs. How do US parent exports change following increases in trade costs across Europe? As the cost of intra-Europe exports increases, US parent exports to the United Kingdom and to EU become a more profitable option. Table 6 shows that those export increase following the Brexit shock. For brevity, we focus on an increase in $\tau^e_{UK,j}$ only and an increase in all trade costs simultaneously. In both scenarios, US parents' exports to the United Kingdom and to EU countries increase on impact, both on the extensive and on the intensive margin. While the increase in US exports is almost the same in the short and long run when $\tau^e_{UK,j}$ increases, this increase is steady over time in the case of the trade barriers driving dynamics, $f^e_{UK,j}$ and $F^e_{UK,j}$, also increasing.

Table 6 shows the forces of the proximity-concentration tradeoff at work in a quantitatively relevant

[32] Appendix Figures E.2 and E.3 show the results for Ireland and Germany.

Figure 9: Brexit: US affiliates in France.

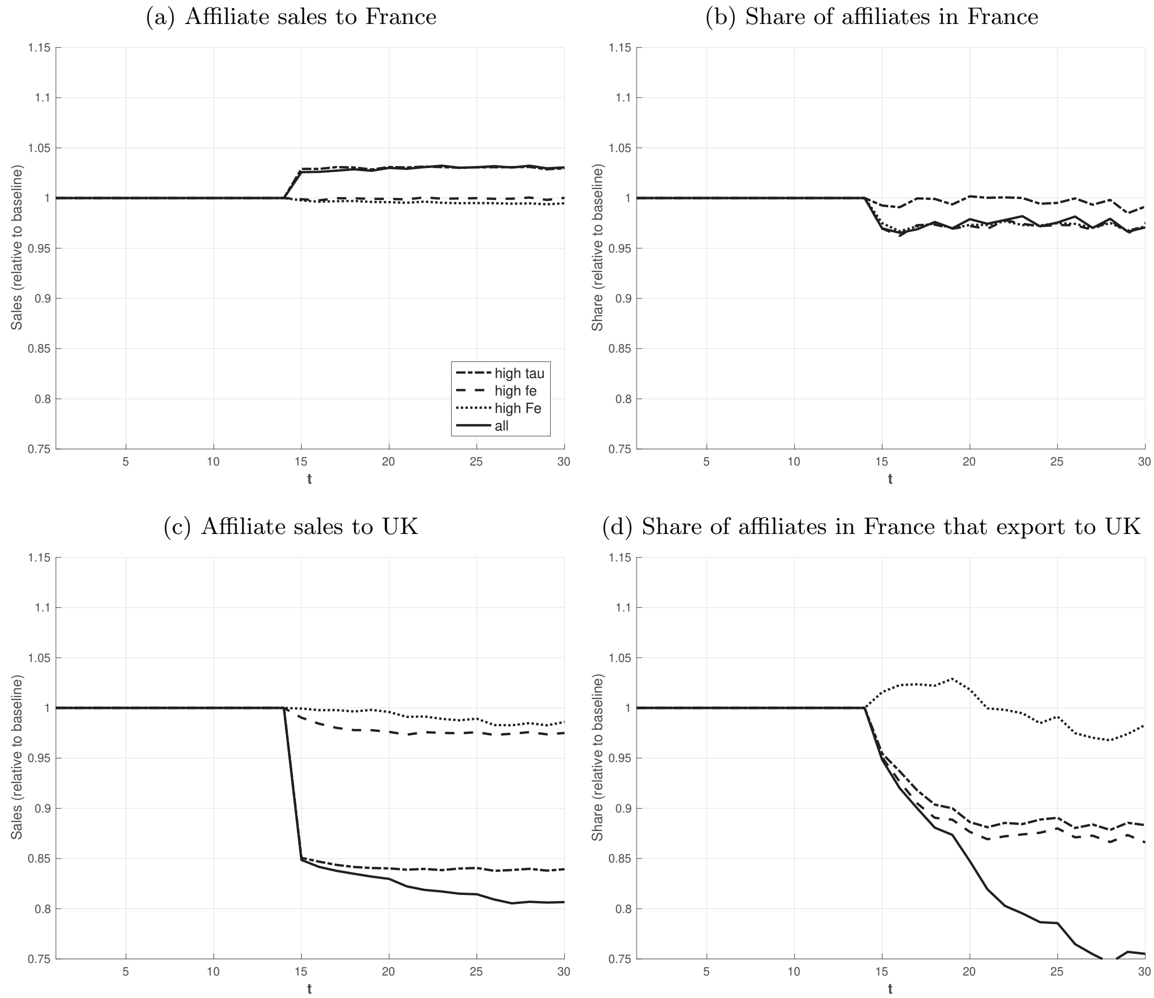

Note: "high X" refers to an increase in the barrier X from/to the United Kingdom to/from country k. "All" refers to increasing all three export barriers from/to the United Kingdom to/from country k at once. Country k refers to Ireland, Germany, and France.

scenario. US exports increase when the cost of multinational activity (here, export platforms) increases. General equilibrium effects on the price indexes affect the relative costs of production across countries, with the consequence that the increases in the price indexes across Europe following the Brexit shock make US domestic production and exports a cheaper way to serve UK and EU markets.

Table 6: US exports after Brexit: the proximity-concentration trade off.

	One period after shock		15 periods after shock	
	high τ	all	high τ	all
US parent export sales to UK	1.034	1.037	1.036	1.045
Share of US parents exporting to UK	1.008	1.008	1.017	1.021
US parent export sales to EU	1.028	1.029	1.028	1.033
Share of US parents exporting to EU	1.020	1.020	1.008	1.009

Note: Variables are relative to the baseline in the same period. "high τ" refers to an increase in τ from/to the United Kingdom to/from country k. "All" refers to increasing all three export barriers from/to the United Kingdom to/from country k at once. Country k refers to Ireland, Germany, and France.

5.3 The role of endogenous prices

Due to the tractability of the model, we can solve numerically for the aggregate price index in each period, for each country. In this way, 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*? Figure 10 shows the dynamics of aggregate outcomes after the increase in all trade barriers simultaneously between the United Kingdom and EU countries, under endogenous and exogenous prices. The price adjustments act as buffers to the decrease in affiliate horizontal and export sales and the share of US MNEs operating in the United Kingdom. In the case of horizontal sales (Figure 10a), the effect is strong enough to reverse the pattern under exogenous prices from a two-percent decrease to a three-percent increase.

We can further evaluate welfare losses from the different Brexit scenarios using the implied price changes after each Brexit shock, for each of the countries involved (see Appendix Figure E.1). The largest increase in prices is experienced, on impact, when all trade barriers are increased at once, ranging from 0.2 percent in Germany to 0.8 percent in Ireland. Price indexes continue to increase over time, leading to more than one percent decline for Ireland. These magnitudes are not small, given that our model only includes the behavior of US MNEs. The increase in trade barriers would presumably also affect local exporters and other non-US MNE exporters.

When interpreting our results, two remarks are in order. On the one hand, the compound option structure, together with Assumption 1, implies that increases in bilateral frictions only affect the incentives to operate affiliates in continental Europe through changes in the cost of accessing the export network available from a given host country. The model produces reallocations across host countries only as a response to changes in the price indexes, and this feature may have the effect of

Figure 10: Brexit: US affiliates in the United Kingdom. Endogenous vs exogenous prices.

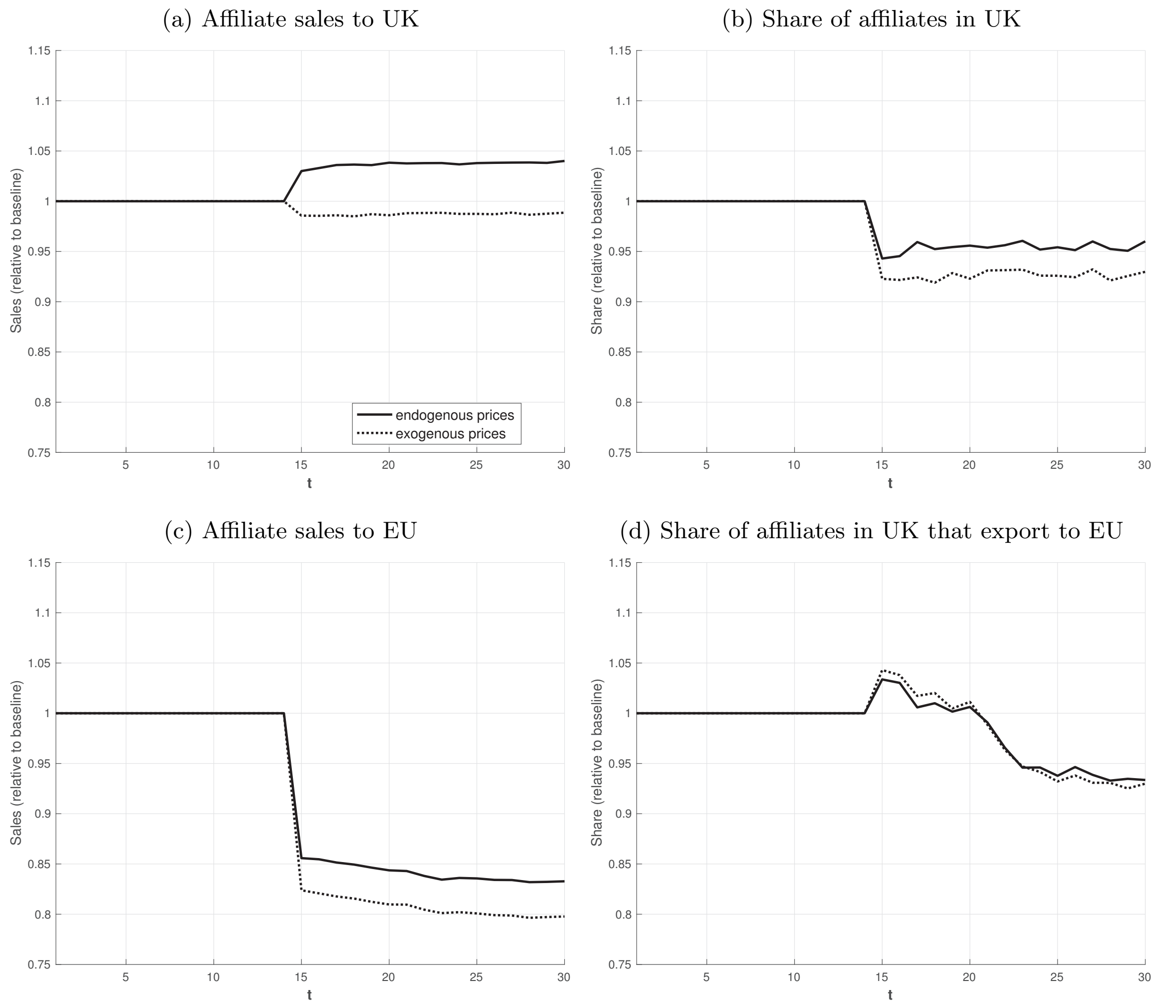

Note: The shock refers to increasing all three export barriers from/to the United Kingdom to/from country k at once. Country k refers to Ireland, Germany, and France.

overstating the losses from increasing MNE frictions. On the other hand, with exogenous wages, the aggregate price effects are stronger than in the case of endogenous wages, so that this assumption may understate the losses from increasing MNE frictions.

5.4 The role of the compound option

To evaluate quantitatively the role played by the compound option structure of our model, we analyze the effects of an increase in the barriers to MNE activities in a model with and without the compound option—that is, a model with and without affiliate exports. We increase the per-period fixed cost and the one-time sunk cost of MNE activities in the United Kingdom, f^h_{UK} and F^h_{UK}, in an amount such that, in either case, $f^h_{UK} + \rho F^h_{UK}$ increases by 20 percent. We calibrate the model with no compound option by targeting the moments in our baseline calibration that do not involve affiliate exports (moments 1.1, 2.1, and 3.1 in Table 2).[33]

Figure 11: The role of the compound option: US MNE affiliates in the United Kingdom.

(a) Affiliate sales to UK (b) Share of affiliates in UK

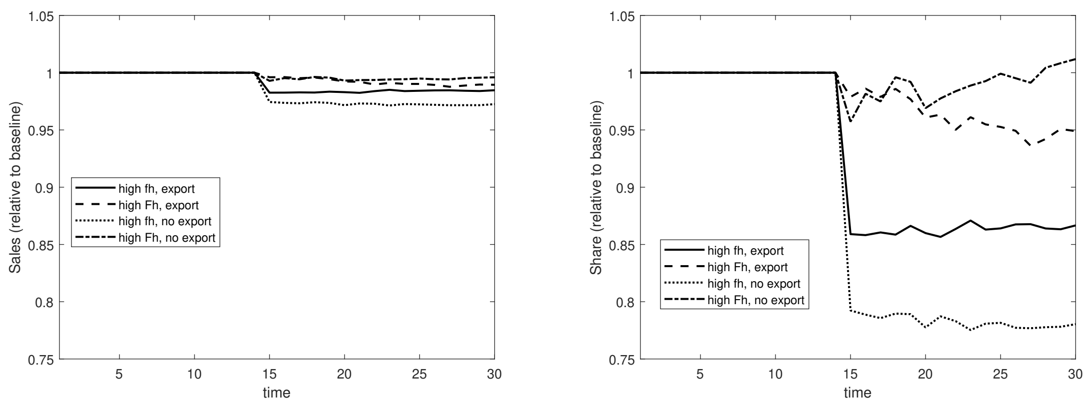

Note: At $t = 15$, the per-period fixed cost (one-time sunk cost) of MNE activities in the UK, f^h_{UK} (F^h_{UK}), increases by an amount such that $f^h_{UK} + \rho F^h_{UK}$ increases by 20 percent. "Exports" refers to the full model where MNE affiliates can export, while "no exports" refers to the calibrated model where MNE affiliates cannot export. Results are shown as deviations from the calibrated models with and without export platforms.

Figure 11 shows the dynamics of horizontal affiliate sales and of the share of US affiliates in the United Kingdom after an increase in f^h_{UK} and F^h_{UK}. Results are shown as deviations from the respective calibrated models with and without the compound option. In the model with only horizontal sales, MNE affiliates 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 this possibility, an increase in the per-period costs of MNE activities in the United Kingdom would decrease the presence of US affiliates there by 25 percent more (in the short run) than in the case in which exports are possible. An increase in sunk costs widens the affiliates' band of inaction, so that fewer new affiliates enter and fewer incumbent affiliates exit. The response of these two margins is of different magnitude between the models with and without affiliate exports. In our baseline model with affiliate exports, the reduced entry margin dominates, so that in the long run the share of US MNEs in the United Kingdom declines.

[33]The calibrated model without affiliate exports matches the targeted moments slightly worse than the baseline model. However, it delivers estimates of sunk (fixed) MNE costs that are very different, as share of US parent (horizontal affiliate) sales. Sunk MNE costs are virtually zero for all sales' percentiles, while fixed costs are around half the magnitude the estimates for the baseline model, for all sales percentiles.

Conversely, in the model without affiliate exports, the reduced exit margin dominates: after 15 periods from the shock, the share of US MNEs in the UK is even higher than in the baseline. The response of the intensive margin of sales in the two models mimics the extensive margin response, but at a much smaller scale.

The results in Figure 11 point not only to the importance of including a rich dynamic spatial structure when evaluating the response of MNE expansion to shocks, but once again to the importance of distinguishing between per-period and one-time costs, which cannot be done in static models.

6 Conclusions

This paper studies the expansion patterns of multinational enterprises (MNEs) in time and space. Using a long panel of US MNEs, we document a set of facts that guide the development of a dynamic model of MNEs that is tractable and, at the same time, rich enough to capture the spatial complexity of MNE activities observed in the data. The model features heterogeneity in firm productivity, persistent aggregate shocks, and a realistic 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. We introduce a compound option formulation for the dynamic problem of the MNE, which is novel to the literature. Our quantitative exercises reveal that the compound option structure is important for understanding the reallocation of MNE activity in time and space after a shock. These exercises also reveal that the nature of the frictions to MNE activities (variable, fixed, or sunk) is important for understanding aggregate firm dynamics.

The problem of the MNE is a natural environment to apply a compound-option structure, since MNE location and export choices happen over time and they are likely to be affected by uncertainty in demand and other market characteristics. This structure, however, can prove useful for problems in other contexts. Problems related to global sourcing decisions, which are likely to occur over time and under uncertainty, are good candidates. This is an avenue for future research.

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.

Alviarez, V. (2019). Multinational production and comparative advantage. *Journal of International Economics 119*, 1–54.

Antrás, 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.

Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers 16*(1), 159–178.

Baldwin, R. and P. Krugman (1989). Persistent trade effects of large exchange rate shocks. *Quarterly Journal of Economics 104*(4), 635–654.

Brainard, S. L. (1997). An empirical assessment of the proximity-concentration trade-off between multinational sales and trade. *The American Economic Review 87*(4), 520–544.

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.

Eaton, J., S. Kortum, and F. Kramarz (2011). An anatomy of international trade: Evidence from french firms. *Econometrica 79*(5), 1453–1498.

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. (2020). 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 (2019). 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.

Gumpert, A., H. Li, A. Moxnes, N. Ramondo, and F. Tintelnot (2020). The life-cycle dynamics of exporters and multinational firms. *Journal of International Economics 126.*

Head, K. and T. Mayer (2019). Brands in motion: How frictions shape multinational production. *American Economic Review 109*(9), 3073–3124.

Helpman, E., M. J. Melitz, and S. R. Yeaple (2004). Exports versus fdi with heterogeneous firms. *American Economic Review 94*(1), 300–316.

Impullitti, G., A. A. Irarrazabal, and L. D. Opromolla (2013). A theory of entry and exit into exports markets. *Journal of International Economics 90*(1), 75–90.

Irarrazabal, A., A. Moxnes, and L. D. Opromolla (2013). The margins of multinational production and the role of intra-firm trade. *Journal of Political Economy 121*(1), 74–126.

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 (2019). Extended gravity. *The Review of Economic Studies 86*(6), 2668–2712.

Ramondo, N., V. Rappoport, and K. J. Ruhl (2013). The proximity-concentration tradeoff under uncertainty. *Review of Economic Studies 80*(4), 1582–1621.

Ramondo, N. and A. Rodríguez-Clare (2013). Trade, multinational production, and the gains from openness. *Journal of Political Economy 121*(2), 273–322.

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), 703–725.

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.

Yeaple, S. R. (2013). The multinational firm. *Annual Review of Economics 5*, 193–217.